aruba
a Hewlett Packard Enterprise company

ArubaOS 8 Fundamentals

Authors:
Syed Ahmed
Jerrod Howard
Kevin Marshall
Mak Moussa
Andrew Tanguay

Contributors:
Shiv Mehra
Dipen Vardhe

Fundamentals Guide

Copyright Information

Open Source Code

This product includes code licensed under the GNU General Public License, the GNU Lesser General Public License, and/or certain other open source licenses. A complete machine-readable copy of the source code corresponding to such code is available upon request. This offer is valid to anyone in receipt of this information and shall expire three years following the date of the final distribution of this product version by Hewlett Packard Enterprise Company. To obtain such source code, send a check or money order in the amount of US $10.00 to:

Hewlett Packard Enterprise Company
Attn: General Counsel
3000 Hanover Street
Palo Alto, CA 94304
USA

aruba
a Hewlett Packard
Enterprise company

www.arubanetworks.com

3333 Scott Blvd

Santa Clara, CA 95054

Phone: 1-800-WIFI-LAN (+800-943-4526)

Fax 408.227.4550

Revision History

The following table lists the revisions of this document:

Revision	Date	Change Description
1.1.0	6/1/2018	Edits made to Controller Mode Comparison, Change of Authorization, and Licensing Concepts
1.0.0	5/28/2018	Initial Publication

Table 1 *Revision History*

About This Guide

Overview

Aruba Deployment Guides are best practice recommendation documents specifically designed to outline how Aruba technology works and to enable customers who deploy Aruba solutions to achieve optimal results. This document is not only intended to serve a deployment guide but also to provide descriptions of Aruba technology, recommendations for product selections, network design decisions, configuration procedures, and best practices. Together, Aruba documentation suite for ArubaOS 8 comprises a reference model for understanding Aruba technology and designs for common customer deployment scenarios. Our customers rely on these proven designs to rapidly deploy Aruba solutions in their production environments with the assurance that they will perform and scale as expected.

Intended Audience

This guide is intended for administrators who are responsible for deploying and configuring AOS 8 solutions on customer premises. Readers should have at least a basic understanding of WLAN concepts. This is a base design guide for ArubaOS and it is assumed that readers have at least a working understanding of fundamental wireless concepts as well as Aruba technology.

Related Documents

The following documents may be helpful as supplemental reference material to this guide:
ArubaOS 8 User Guide
ArubaOS 8 CLI Reference Guide
Aruba Solution Exchange

Conventions

Typographical Conventions

The following conventions are used throughout this manual to emphasize important concepts:

Style Type	Description
Italics	Italics are used to emphasize important terms and to mark the titles of documents.
Bolded words	Bolded words indicate an option that should be selected in the Graphical User Interface (GUI). The angled brackets indicate that the choices are part of a path in the GUI.
`Command Text`	Command text in this font will appear inside of a box and indicates commands that can be entered into the Command Line Interface (CLI).
`<`*`Arguments`*`>`	In the command examples, italicized text within single angle brackets represents items that should be replaced with information appropriate to your specific situation. For example: `# send <`*`text message`*`>` In this example, you would type “send” at the system prompt exactly as shown, followed by the text of the message you wish to send. Do not type the angle brackets.
`[Optional]`	Command examples enclosed in brackets are optional. Do not type the brackets.
`{Item A | Item B}`	In the command examples, items within curly braces and separated by a vertical bar represent the available choices. Enter only one choice. Do not type the braces or bars.

Table 2 *Typographical Conventions*

Informational Icons

The following informational icons are used throughout this guide:

NOTE Indicates helpful suggestions, pertinent information, and important things to remember.

CAUTION Indicates a risk of damage to your hardware or loss of data.

WARNING Indicates a risk of personal injury or death.

Graphical Icons

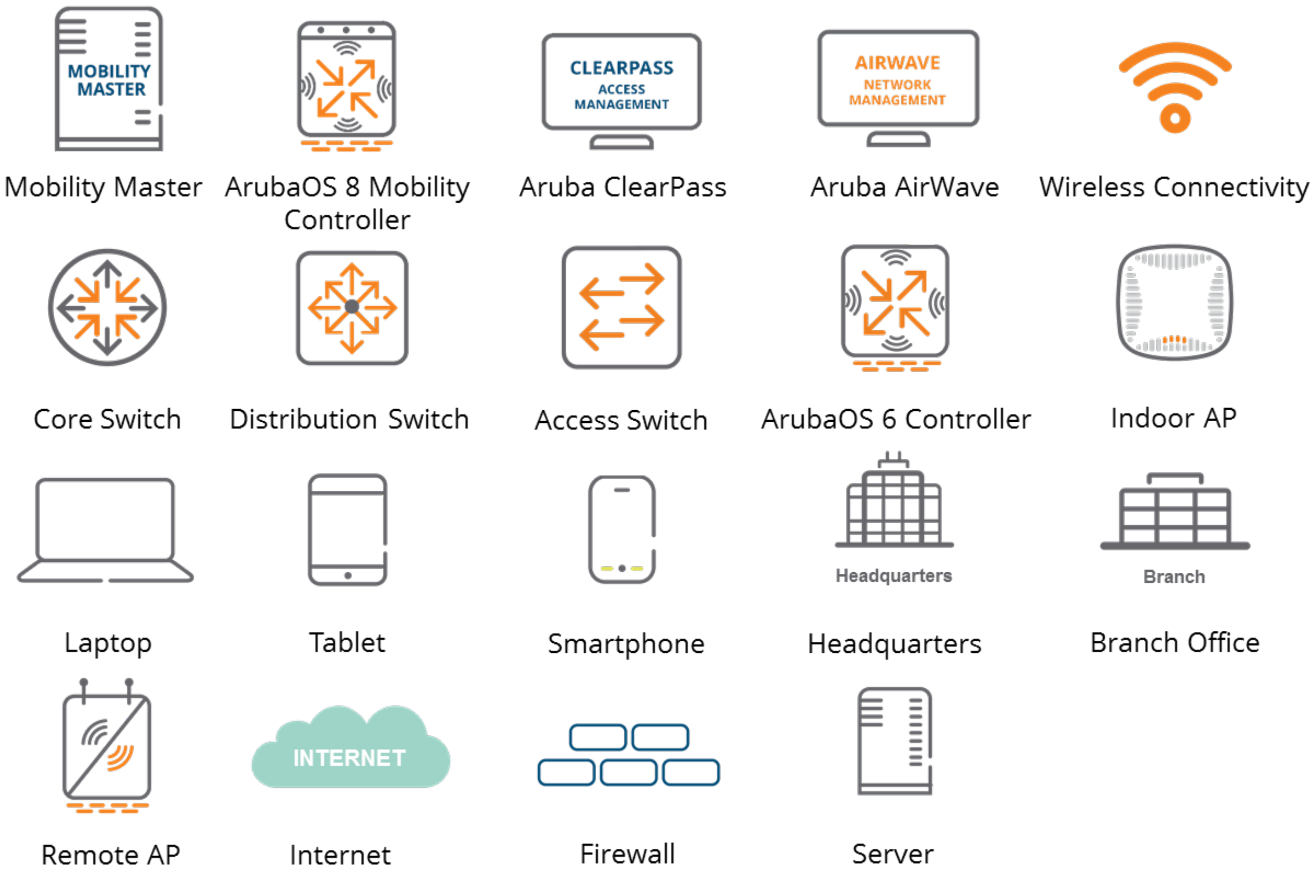

Figure 1 *VRD Icon Set*

Acronym List

Acronym	Definition
A-MPDU	Aggregated Media Access Control Packet Data Unit
A-MSDU	Aggregated Media Access Control Service Data Unit
AAA	Authentication, Authorization, and Accounting
AAC	AP Anchor Controller
ACL	Access Control List
ACR	Advanced Cryptography
AM	Air Monitor
AP	Access Point
API	Application Programming Interface

ARM	Adaptive Radio Management
ASP	Aruba Support Portal
BLMS	Backup Local Management Switch
CoA	Change of Authorization
CLI	Command Line Interface
CPSec	Control Place Security
CPU	Central Processing Unit
DC	Data Center
DLNA	Digital Living Network Alliance
DNS	Domain Name Service
DHCP	Dynamic Host Configuration Protocol
DMZ	Demilitarized Zone
DPI	Deep Packet Inspection
FQDN	Fully-qualified Domain Name
GRE	Generic Routing Encapsulation
GUI	Graphical User Interface
HA	High Availability
HMM	Hardware MM
HTTP	Hypertext Transfer Protocol
HTTPS	HTTP Secure
IAP	Instant Access Point
IDF	Intermediate Distribution Frame
IKE	Internet Key Exchange
IP	Internet Protocol
IPsec	Internet Protocol Security
IPv4	Internet Protocol Version 4
IPv6	Internet Protocol Version 6
JSON	JavaScript Object Notation
LACP	Link Aggregation Control Protocol
LAN	Local Area Network
LAG	Link Aggregated Connection

LMS	Local Management Switch
LSM	Loadable Service Module
MAC	Media Access Control
MC	Mobility Controller
MCM	Master Controller Mode
mDNS	multicast Domain Name Service
MD	Managed Device
MD	Mobility Device
MDF	Main Distribution Frame
MM	Mobility Master
MM-HW	Mobility Master - Hardware
MM-VA	Mobility Master – Virtual Appliance
MN	Managed Node
MNP	MyNetworking Portal
NAS	Network Access Server
NAT	Network Address Translation
NBAPI	Northbound Application Programming Interface
NVF	Network Virtualization Functionality
PAPI	Proprietary Access Protocol Interface
PAT	Port Address Translation
PEF	Policy Enforcement Firewall
PSK	Pre-shared Key
QoS	Quality of Service
RADIUS	Remote Authentication Dial In User Service
RAM	Random Access Memory
RAP	Remote Access Point
RF	Radio Frequency
RFP	RF Protect
S-AAC	Standby AP Anchor Controller
S-UAC	Standby User Anchor Controller
SDN	Software Defined Network

SfB	Skype for Business
SIP	Session Initiation Protocol
SSID	Service Set Identifier
SVI	Switch Virtual Interface
UAC	User Anchor Controller
UCC	Unified Communications and Collaboration
UCM	Unified Communication Manager
UDLD	Unidirectional Link Detection
URL	Uniform Resource Locator
VAP	Virtual Access Point
VIA	Virtual Internet Access
VIP	Virtual Internet Protocol address
VLAN	Virtual Local Area Network
VM	Virtual Machine
VMC	Virtual MC
VMM	Virtual MM
VPN	Virtual Private Network
VPNC	Virtual Private Network Concentrator
VRRP	Virtual Router Redundancy Protocol
VSF	Virtual Switching Framework
WLAN	Wireless Local Area Network
XML	Extensible Markup Language
ZTP	Zero-touch Provisioning

Architecture

Nearly all end user devices in modern production networks are wireless devices, including laptops which are shipped without an Ethernet port. Even wired phones are being replaced with unified communication applications such as Skype for Business (SfB). These trends are forcing enterprises to be increasingly reliant on wireless Local Area Networks (LANs) to address their business needs. With such critical dependencies on wireless LANs, network administrators are required to design complex networks to support various types of applications, users, and devices without compromising security. This deployment guide will outline all of the features enabled through Aruba's state-of-the-art ArubaOS 8 which help address these challenges being encountered in modern production networks.

ArubaOS is the operating system for all Aruba Mobility Controllers (MCs) and controller-managed wireless access points (APs). With an extensive set of integrated technologies and capabilities, ArubaOS 8 delivers unified wired and wireless access, seamless roaming, enterprise grade security, and a highly available network with the required performance, user experience, and reliability to support high density environments.

Product Portfolio

The following table lists the controllers supported in ArubaOS 8 and their capabilities:

Controller Series	Controller Model	Number of APs	Number of Users	Firewall (Gbps)
70xx	7005	16	1,024	2
	7008	16	1,024	2
	7010	32	2,048	4
	7024	32	2,048	4
	7030	64	4,096	8
72xx	7205	256	8000	12
	7210	512	16000	20
	7220	1,024	24000	40
	7240	2,048	32000	40
	7280	2,048	32000	100

Table 3 *ArubaOS 8 Product Portfolio*

ArubaOS 8.x does not support 3000 or 600 series controllers.

ArubaOS 6 Controller Modes

ArubaOS 6 supports the following controller modes:

Master Mode

If deploying a controller in Master Mode an administrator is responsible for all configuration including IP addresses, licenses, WLANs, AP groups, user roles, etc. The administrator applies the configuration by connecting to the console of the controller and navigating through the Command Line Interface (CLI) wizard.

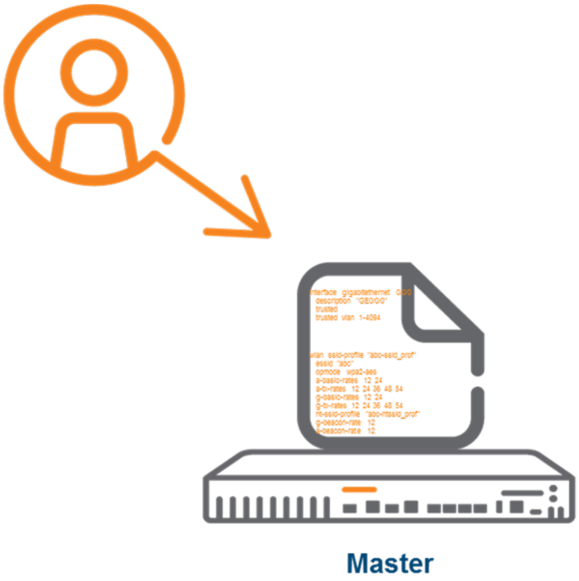

```
interface gigabitethernet 0/0/0
 description "GE0/0/0"
 trusted
 trusted vlan 1-4094

wlan ssid-profile "abc-ssid_prof"
 essid "abc"
 opmode wpa2-aes
 a-basic-rates 12 24
 a-tx-rates 12 24 36 48 54
 g-basic-rates 12 24
 g-tx-rates 12 24 36 48 54
 ht-ssid-profile "abc-htssid_prof"
 g-beacon-rate 12
 a-beacon-rate 12
```

Figure 2 *ArubaOS 6 Controller in Master Mode*

A Master Mode controller can partially manage a Local controller as well as fully manage a Branch controller. All 70xx series controllers and 72xx series controllers are capable of operating in Master Mode, however a 70xx series master controller cannot manage another 70xx series branch controller.

Local Mode

Similar to Master controllers, Local controllers need to be initially configured by an administrator using the serial console to assign IP addresses, Virtual Local Area networks (VLANs), and other network related parameters. However, global configuration parameters such as Wireless Local Area Networks (WLANs) and AP Groups are inherited from the Master controller which serves as the single pane of glass for global configuration. Global configurations inherited from a Master controller cannot be modified through the Local controller graphical user interface (GUI) and will appear to be greyed out. All 70xx series controllers and 72xx series controllers are capable of operating in Local Mode.

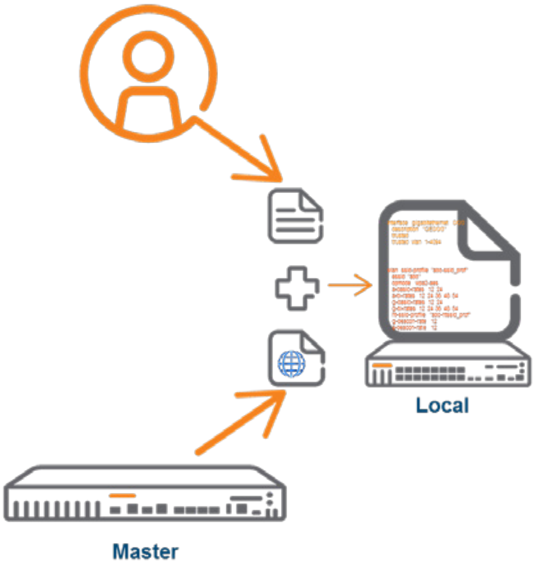

Figure 3 *ArubaOS6 Controller in Local Mode*

Branch Mode

Branch Mode is relatively new controller operating mode which was introduced in the 6.4.3 release. It is typically used in distributed enterprises with geographically separated branch and remote offices. Branch mode is enabled by default unlike Master Mode and Local Mode, however it is only supported 70xx series controllers. All 70xx series controllers ship with Dynamic Host Configuration Protocol (DHCP) client functionality enabled their last port on VLAN 4096 by default. A Master controller IP can be assigned to the Branch controller using DHCP option 43. Alternatively, a Branch controller can use Zero Touch Provisioning (ZTP) through Aruba Activate to obtain the Master Controller IP.

In contrast to Local controllers, Branch controllers receive their entire configuration from their Master controller. There are no write permissions permitted in the Branch controller GUI or CLI. A special GUI called Smart Config is enabled on the Master controller for purposes of Branch controller management and configuration. Controllers operating in Branch Mode have all the same features of other controllers however only the features exposed in Smart Config can be enabled. The figure below illustrates how configuration a Branch controller is configured.

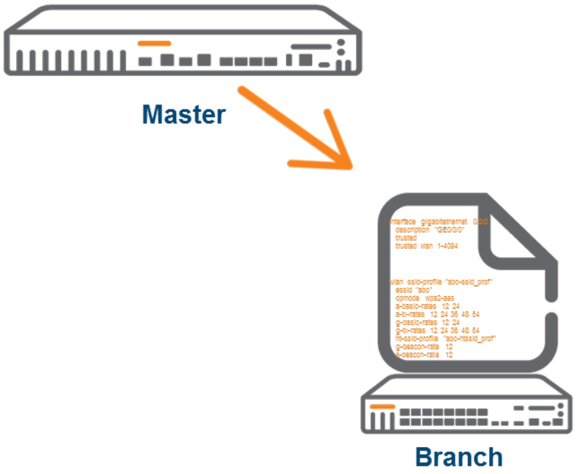

```
interface gigabitethernet 0/0/0
  description "GE0/0/0"
  trusted
  trusted vlan 1-4094

wlan ssid-profile "abc-ssid_prof"
  essid "abc"
  opmode wpa2-aes
  a-basic-rates 12 24
  a-tx-rates 12 24 36 48 54
  g-basic-rates 12 24
  g-tx-rates 12 24 36 48 54
  ht-ssid-profile "abc-htssid_prof"
  g-beacon-rate 12
  a-beacon-rate 12
```

Figure 4 *ArubaOS 6 Controller in Branch Mode*

Standalone

Standalone controllers are employed when an architectural design requires a deployment consisting of single controller. Standalone controllers are incapable of managing Local or Branch controllers and are very rarely used in production environments due to their inherent lack of redundancy. They have all the same functionality of a Master controller with the notable exception that they cannot manage other controllers.

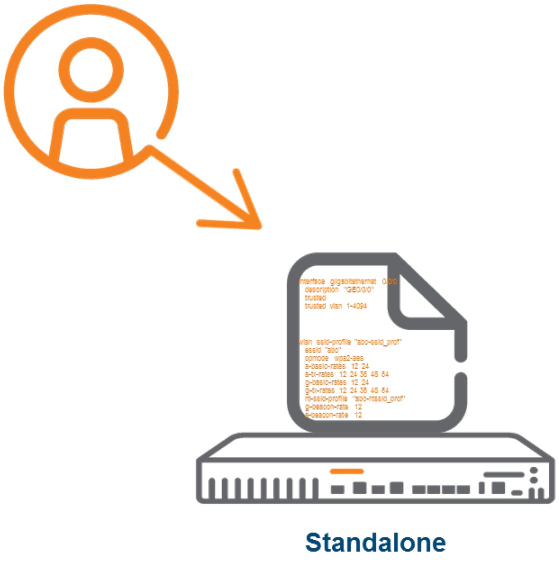

```
interface gigabitethernet 0/0/0
  description "GE0/0/0"
  trusted
  trusted vlan 1-4094

wlan ssid-profile "abc-ssid_prof"
  essid "abc"
  opmode wpa2-aes
  a-basic-rates 12 24
  a-tx-rates 12 24 36 48 54
  g-basic-rates 12 24
  g-tx-rates 12 24 36 48 54
  ht-ssid-profile "abc-htssid_prof"
  g-beacon-rate 12
  a-beacon-rate 12
```

Figure 5 *6.x Controller in Standalone Mode*

ArubaOS 6 Topology

A typical ArubaOS 6 controller topology is characterized by two Master controllers at a headquarters campus location in a Master/Standby configuration. The Master controllers are connected to several Local controllers at the headquarters campus as well as any Branch controllers deployed at remote locations. AirWave and ClearPass as deployed as well for network management and network access control purposes.

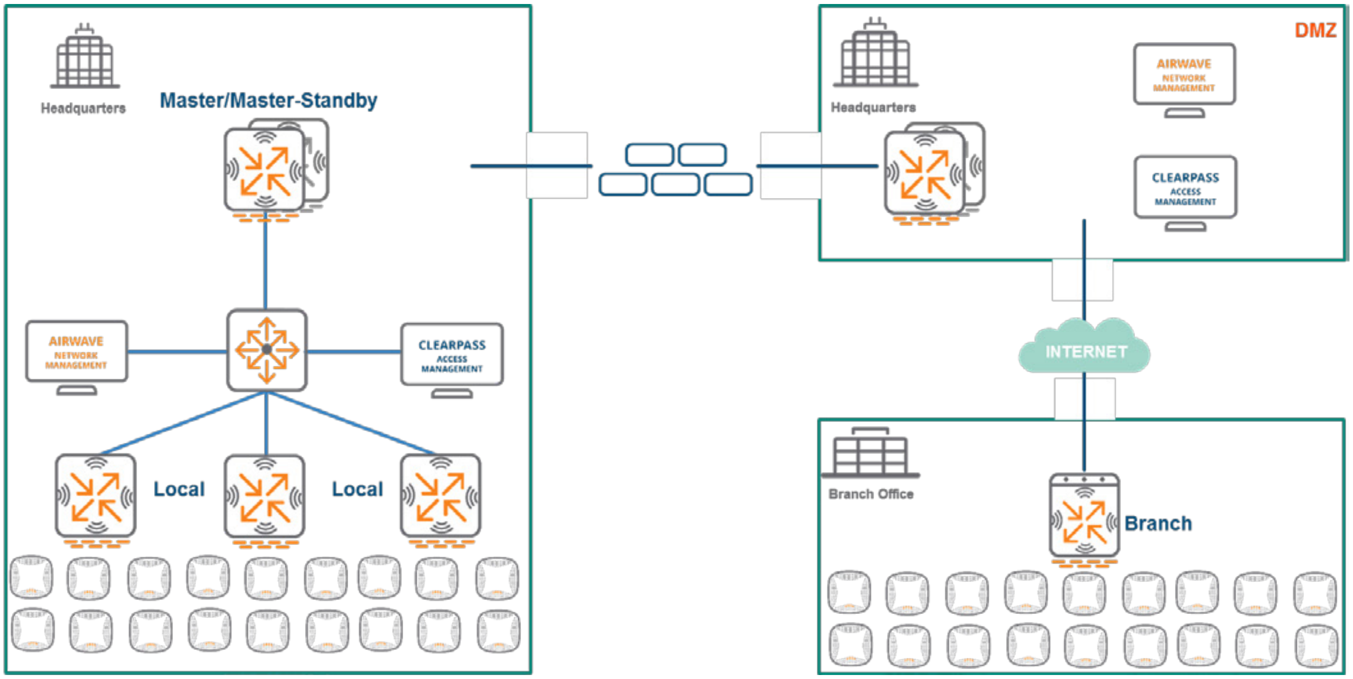

Figure 6 *Typical ArubaOS 6 Topology*

ArubaOS 8 Controller Modes

ArubaOS 8 supports the following controller modes:

Mobility Master

The concept of the Mobility Master (MM) is new to ArubaOS 8. MMs come in two variations: as Virtual MM (VMM) or as a Hardware MM (HMM). The MM is designed to run on an x86-based platform as many of the features introduced in ArubaOS 8 require random access memory (RAM), central processing unit (CPU), and storage space that are not supported by physical controllers. The MM must be fully configured by an administrator similar to how a Master controller would be configured in ArubaOS 6. Its primary role is to serve as the single point of configuration and image management for the network. In addition, the MM can be configured via a Northbound Application Programmable Interface (NBAPI). A VMM can be installed on VMWare, KVM, or Hyper-V according to what is more suitable for where it will be deployed. HMMs and VMMs may alternatively be referred to as MM-HW and MM-VA meaning MM-Hardware and MM-Virtual Appliance, respectively.

Access points cannot be terminated on a MM.

```
Interface  gigabitethernet  0/0/0
  description  "GE0/0/0"
  trusted
  trusted  vlan  1-4094


wlan  ssid-profile  "abc-ssid_prof"
  essid  "abc"
  opmode  wpa2-aes
  a-basic-rates  12  24
  a-tx-rates  12  24  36  48  54
  g-basic-rates  12  24
  g-tx-rates  12  24  36  48  54
  ht-ssid-profile  "abc-htssid_prof"
  g-beacon-rate  12
  a-beacon-rate  12
```

Figure 7 *ArubaOS 8 MM*

Master Controller Mode

ArubaOS 8 also introduces the concept of Master Controller Mode (MCM) to enable a seamless transition ArubaOS 6 without requiring an x86-based appliance (HMM) or VMM. An MCM is capable of managing other MCs, however only a subset of MM features is available and APs cannot be terminated as they would be with an MM. Only the 7030 model or 72xx series controllers support MCM.

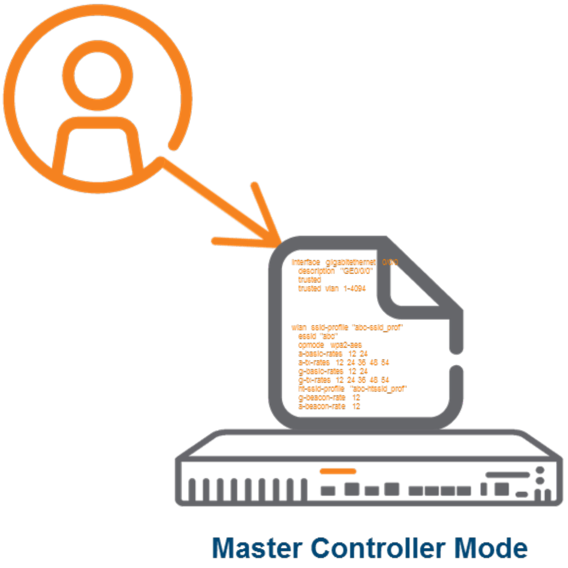

```
interface gigabitethernet 0/0/0
  description "GE0/0/0"
  trusted
  trusted vlan 1-4094


wlan ssid-profile "abc-ssid_prof"
  essid "abc"
  opmode wpa2-aes
  a-basic-rates 12 24
  a-tx-rates 12 24 36 48 54
  g-basic-rates 12 24
  g-tx-rates 12 24 36 48 54
  ht-ssid-profile "abc-htssid_prof"
  g-beacon-rate 12
  a-beacon-rate 12
```

Figure 8 *ArubaOS 8 Master Controller Mode*

The following table outlines which features are supported by the MCM and which require an MM:

Supported Features	Unsupported Features
New GUI, Workflows and Hierarchical Configuration	Clustering
Multizone	AirMatch
Multi-threaded CLI with auto-completion	Centralized App Support (UCC, AppRF)
WAN Link Bonding and Load Balancing	Live Upgrade
Distributed UCC, AppRF, ARM and AirGroup	Centralized Visibility

Table 4 *Master Controller Mode Feature Matrix*

Mobility Controller

The concept of the Mobility Controller or MC is also new to ArubaOS 8. An MC, in the past, has also been known as Managed Node (MN), Managed Device (MD), or Mobility Device (MD) in some Aruba documentation. An MC is similar to a Branch controller in ArubaOS 6 in the sense that it can be adopted using ZTP and Aruba Activate. The last port of an MC is enabled as a DHCP client on VLAN 4094 in its factory default configuration.

NOTE

MMs and MCMs cannot adopt an MC using DHCP Option 43 since MM or MCM certificate distribution is not supported with DHCP Options.

Unlike the ArubaOS 6 Local controllers, MCs can be fully managed by an MM or MCM. In addition, unlike ArubaOS 6 Branch controllers, an administrator can configure every feature of an MC. All 70xx series controllers and 72xx series controllers are shipped as MCs. ArubaOS 8 also supports

Virtual MCs (VMC). A VMC can be deployed either on VMWare, KVM, or Hyper-V. MCs can be configured as Virtual Private Network Concentrators (VPNCs).

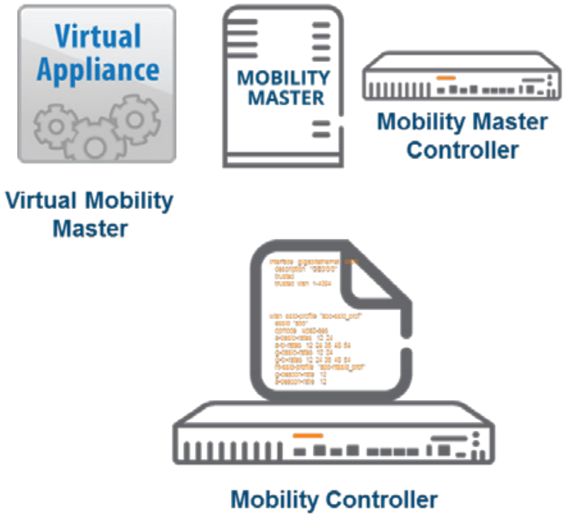

Figure 9 *MC Management*

Standalone Controller

ArubaOS 8 includes the ability to configure a Standalone controller. A Standalone controller cannot be managed by a MM and cannot be clustered with other Standalone controllers. It is similar in functionality to Standalone controllers in ArubaOS 6 and supports the Multizone feature.

AirMatch and clustering are not enabled on Standalone controllers because they can only be implemented with the assistance of a virtual machine (VM). ARM is the sole RF optimization method available. Other features such as WebCC, AppRF, UCC, AirGroup, Northbound API, UCM, and WMS will all function as they do on an ArubaOS 6 Local controller.

Standalone Controller

Figure 10 *ArubaOS 8 Standalone Controller*

ArubaOS 8 Topology and Feature Summary

Migration from a an ArubaOS 6 topology to an ArubaOS 8 topology primarily consists of replacing the Master controller with an MM and replacing all Local and Branch controllers with MCs.

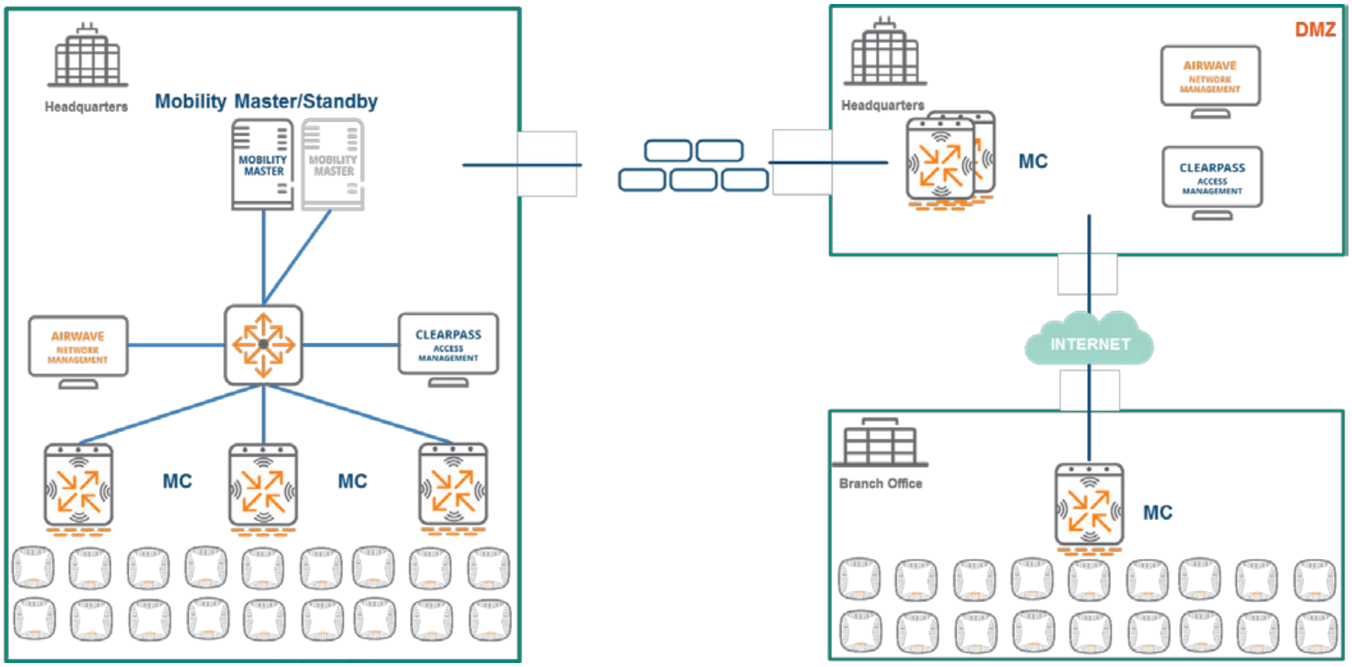

Figure 11 *ArubaOS 8 Topology*

The ArubaOS 8 code base includes the following points of differentiation compared to an ArubaOS 6 deployment:

- Introduces the VM-based MM as a single point of configuration and image management
- Introduces MCs which are completely managed by MMs using ZTP
- The MM does not terminate any APs
- All 72xx/70xx Controllers can be set up as MCs or Standalone controllers
- Introduces Master Controller Mode as a migration path

Controller Mode Comparison

The following table outlines how the previous designations for controller devices in an ArubaOS 6-based architecture have changed and identifies their appropriate counterparts in an ArubaOS 8-based deployment:

ArubaOS 6	ArubaOS 8
Master Controller	MM (VM or hardware) or MCM (72xx and 7030 Only)
Local Controller	MC
Branch Controller	MC
Standalone Controller	Standalone Controller

Table 5 *Controller Mode Summary*

The following key points should be noted regarding controller modes in ArubaOS 8:

1. ArubaOS 6 Master controllers can partially manage Local controllers or fully manage Branch controllers
2. MMs in ArubaOS 8 can manage all types of controllers regardless of where they are deployed.
3. The key distinction between an MM in ArubaOS 8 and a Master controller in ArubaOS 6 is that an MM can neither adopt APs nor point an AP to an MC
4. MCM was introduced as a migration path to ArubaOS 8 as it doesn't require a Virtual MM.
5. Standalone Mode functions in the same manner in ArubaOS 8 as it does in ArubaOS 6. It is supported on hardware-based controllers only
6. ArubaOS 6-based Local controllers only receive a partial configuration from their Master and do not support ZTP. All ArubaOS 8 hardware controllers support ZTP
7. Branch controllers from ArubaOS 6 are replaced by MCs and have full configuration capabilities unlike the limited functionality of Smart Config in ArubaOS 6.

AP Modes

Campus APs

In most ArubaOS 8 topologies campus APs typically operate in one of two modes when communicating with their MC:

- Tunnel Mode
- Decrypt Tunnel Mode

The advantage to both these modes is that the user VLANs reside on the controller and do not have to be managed at the edge. Additional VLANs can be added to the core switch where the MC uplink is connected if necessary. There is no need to add them to the edge switch where the APs terminate. Both of these operating modes simplify network design and allow for flexibility in terminating users.

Tunnel Mode

When operating in the Tunnel forwarding mode, the AP handles all 802.11 association requests and responses, but sends all 802.11 data packets, action frames, and EAPOL frames over a GRE tunnel to the MC for processing. The MC then removes or adds the GRE headers, decrypts or encrypts 802.11 frames, and applies firewall rules to the user traffic as usual.

To achieve maximum performance benefits with Tunnel Mode, end-to-end jumbo frame support should be enabled on a wired switch due to the increased aggregation introduced with the IEEE 802.11ac standard. Using Control Plane Security (CPSec) in Tunnel Mode is not mandatory. The majority of production deployments utilize Tunnel Mode for AP forwarding where the AP sends 802.11 traffic to the controller. Control and data plane traffic between the AP and the MC is always encrypted. Aruba recommends using Tunnel Mode as a best practice as the majority of traffic fits in a standard 1500 byte Ethernet frame and no special handling is required on the wired network to achieve maximum aggregate performance.

When using jumbo frames with Tunnel Mode the network should support a maximum transmission unit (MTU) size of at least 4500 bytes. If the network cannot support an MTU of this size then the benefits of aggregation efficiency over the air will be lost due to fragmentation. Without end-to-end jumbo frames on the wired network, 802.11ac networks can experience performance degradation of up to 30% in some cases. Although, it should be noted that this adverse impact to performance is only noticed when the peak network performance is measured during technology demonstrations. The day-to-day operations in real world production networks are typically unaffected without jumbo frames turned on.

NOTE

Aruba recommends enabling jumbo frames end-to-end as a best practice.

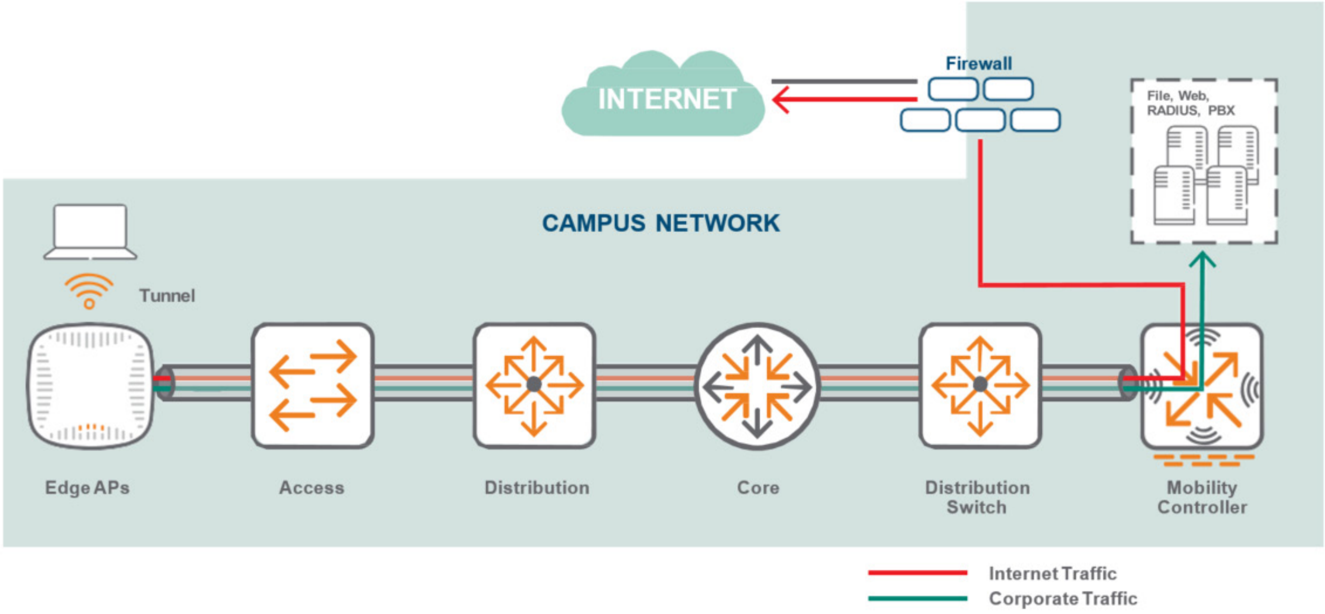

Figure 12 *Tunnel Forwarding Mode*

Decrypt-Tunnel Mode

Decrypt-tunnel mode allows an AP-client pair to take full advantage of Aggregated-Media Access Control (MAC) Service Data Units (A-MSDUs) and Aggregated-MAC Packet Data Units (A-MPDUs) without requiring the wired network to transport jumbo frames. APs perform decryption and de-aggregation on themselves locally. It is mandatory to enable control plane security (CPSec) between APs and Controllers when using Decrypt-tunnel Mode.

CAUTION

Decrypt-tunnel Mode does not provide end-to-end encryption. Only the control plane traffic between APs and MCs is encrypted in Decrypt-tunnel Mode.

In Decrypt-tunnel mode the AP acts as a bridge between clients and the controller in addition to performing encryption and decryption. The MC still acts as the aggregation point for terminating data traffic. This allows the AP-Client pair to take advantage of A-MSDU and A-MPDU on the WLAN radio side without requiring the wired network to transport the jumbo frames since the AP performs all assembly aggregation and de-aggregation locally. The payload is then sent to the controller for firewall processing and L2/L3 forwarding.

Decrypt-tunnel Mode is functionally equivalent to Tunnel Mode with jumbo frames enabled and is typically used for technology demonstrations. It is important to keep in mind that the AP wireless chipset performs cryptography for up to 50 clients which is offloaded to the AP hardware. Scenarios involving more than 50 clients will likely experience minor performance degradation due to this offload process.

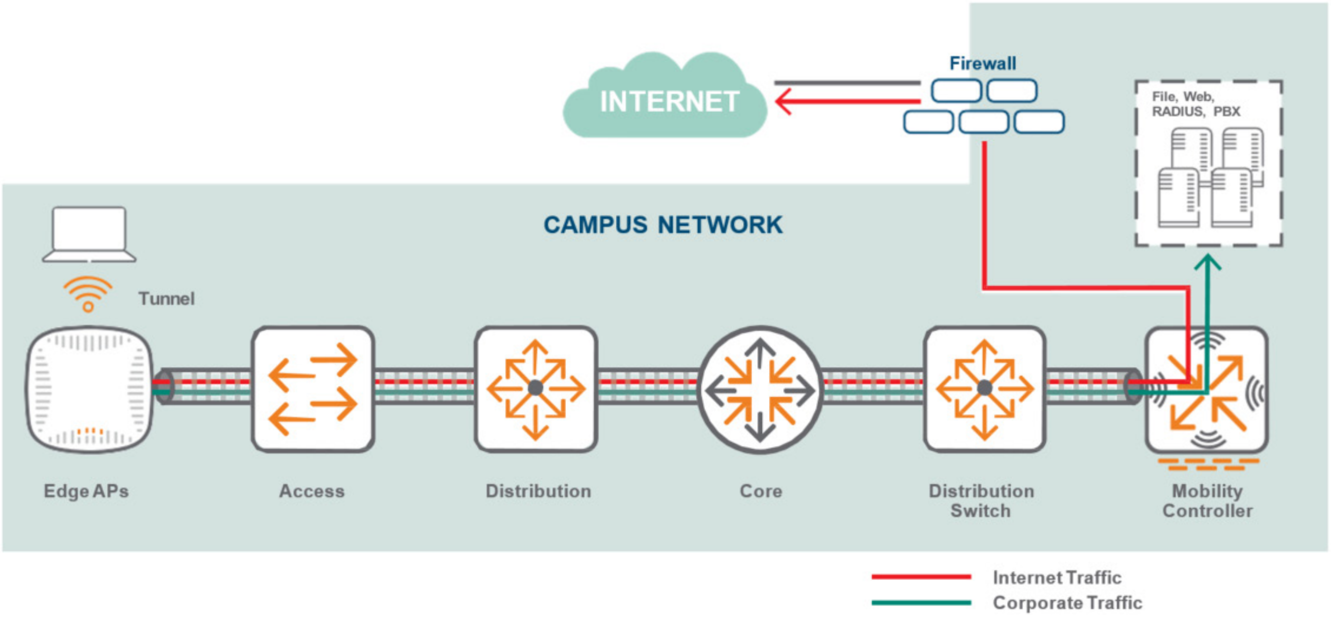

Figure 13 *Decrypt Tunnel Forwarding Mode*

Control Plane Security

The CPSec feature has two main functions:

1. Securing the control channel between Aruba MCs and their attached APs
2. Preventing unauthorized APs from joining the Aruba WLAN network

The aforementioned goals are achieved in the following manner:

- The control traffic transported using Proprietary Access Protocol Interface (PAPI) is secured using a certificate-based Internet Protocol Security (IPsec) tunnel in transport mode
- A CPSec whitelist database holds the list of APs authorized to connect to the Aruba controllers and join the WLAN network

Since CPSec is enabled by default, the MM certifies its MCs using its generated factory certificate after booting up. MCs in turn certify their APs by signing their factory default certificates. Once the APs are authorized through the CPSec whitelist and enter the *certified-factory-cert* state they will initiate secure PAPI (UDP 8209 inside IPsec) communication with the controller, synchronize their firmware, and download their configuration.

Boot Process with Control Plane Security

The figure below illustrates the steps involved in the campus AP boot process with CPSec:

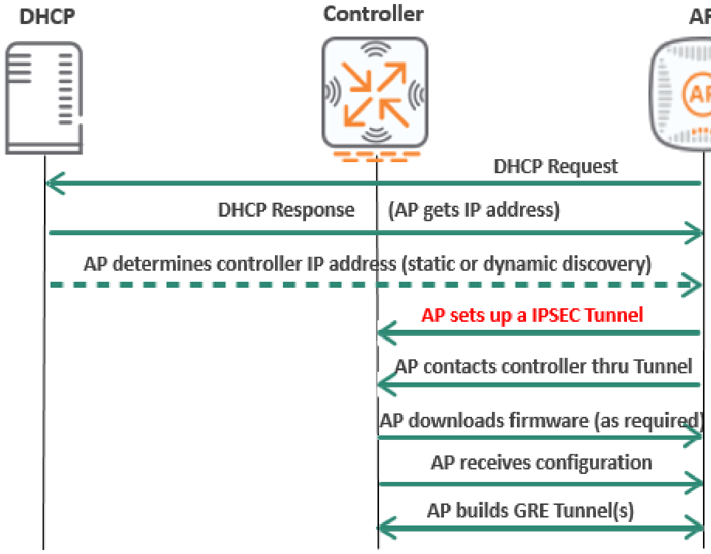
DHCP
Controller
AP
DHCP Request
DHCP Response
(AP gets IP address)
AP determines controller IP address (static or dynamic discovery)
AP sets up a IPSEC Tunnel
AP contacts controller thru Tunnel
AP downloads firmware (as required)
AP receives configuration
AP builds GRE Tunnel(s)

Figure 14 *AP Boot Process with CPSec*

The process includes the following steps:

1. AP sends a DHCP Request
2. AP receives an IP address in the DHCP Response
3. AP determines its controller's IP address either statically or dynamically
4. AP establishes an IPsec tunnel with the controller
5. AP exchanges PAPI (UDP 8209) over the IPsec tunnel with the controller
6. If required, the AP downloads firmware from the newly discovered AP master to ensure version consistency
7. AP receives the configuration from the controller
8. AP creates a Generic Routing Encapsulation (GRE) tunnel for user traffic

Local Management Switch

In multi-controller networks, each controller acts as a local management switch (LMS) by terminating user traffic from the APs, processing, and forwarding the traffic to the wired network. An LMS and a backup local management switch (BLMS) are the primary and secondary connection points for an AP. APs rely on heartbeat timeouts with the LMS controller to failover to a preconfigured BLMS controller.

When controllers are in separate L3 networks, Virtual Router Redundancy Protocol (VRRP) cannot be used for redundancy. In such a case, the LMS and BLMS should be used for redundancy.

In the most basic scenario of two L3 separated controllers:

- The AP finds aruba-master and obtains the LMS and BLMS IPs as part of its configuration
- The AP terminates on the LMS controller
- If the LMS controller fails, eight consecutive missed heartbeats will trigger an AP failover
- The AP comes up on the BLMS

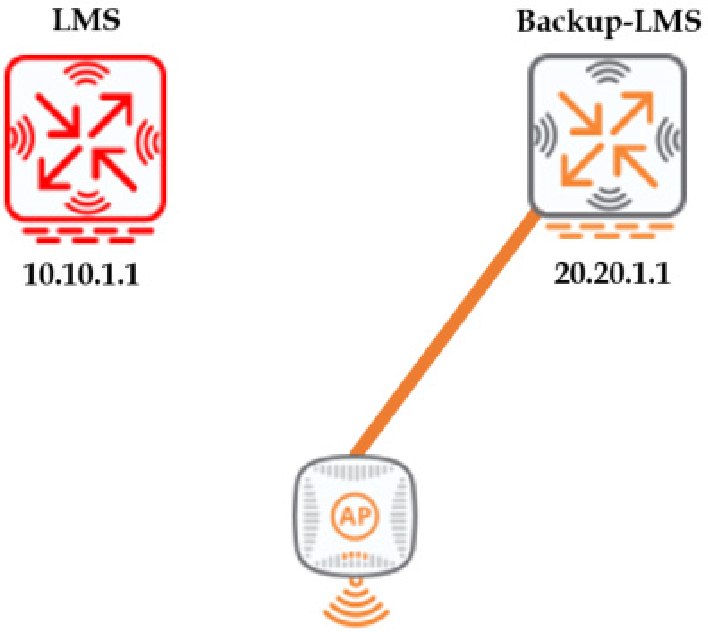

Figure 15 LMS and Backup LMS

Another scenario could be an AP terminated on an LMS that is a cluster of controllers.

An AP finds the aruba-master and obtains the IP addresses of its LMS and BLMS as part of its configuration. If the LMS is located in a cluster of controllers and the LMS fails, any APs terminated on that LMS will attempt to failover to the other members of the cluster. The AP will only failover to its BLMS if all of the other members in the cluster have failed. The BLMS could either be a single controller or a member of a cluster.

NOTE

The concept of clustering is covered in detail in the in the Clustering chapter of this document.

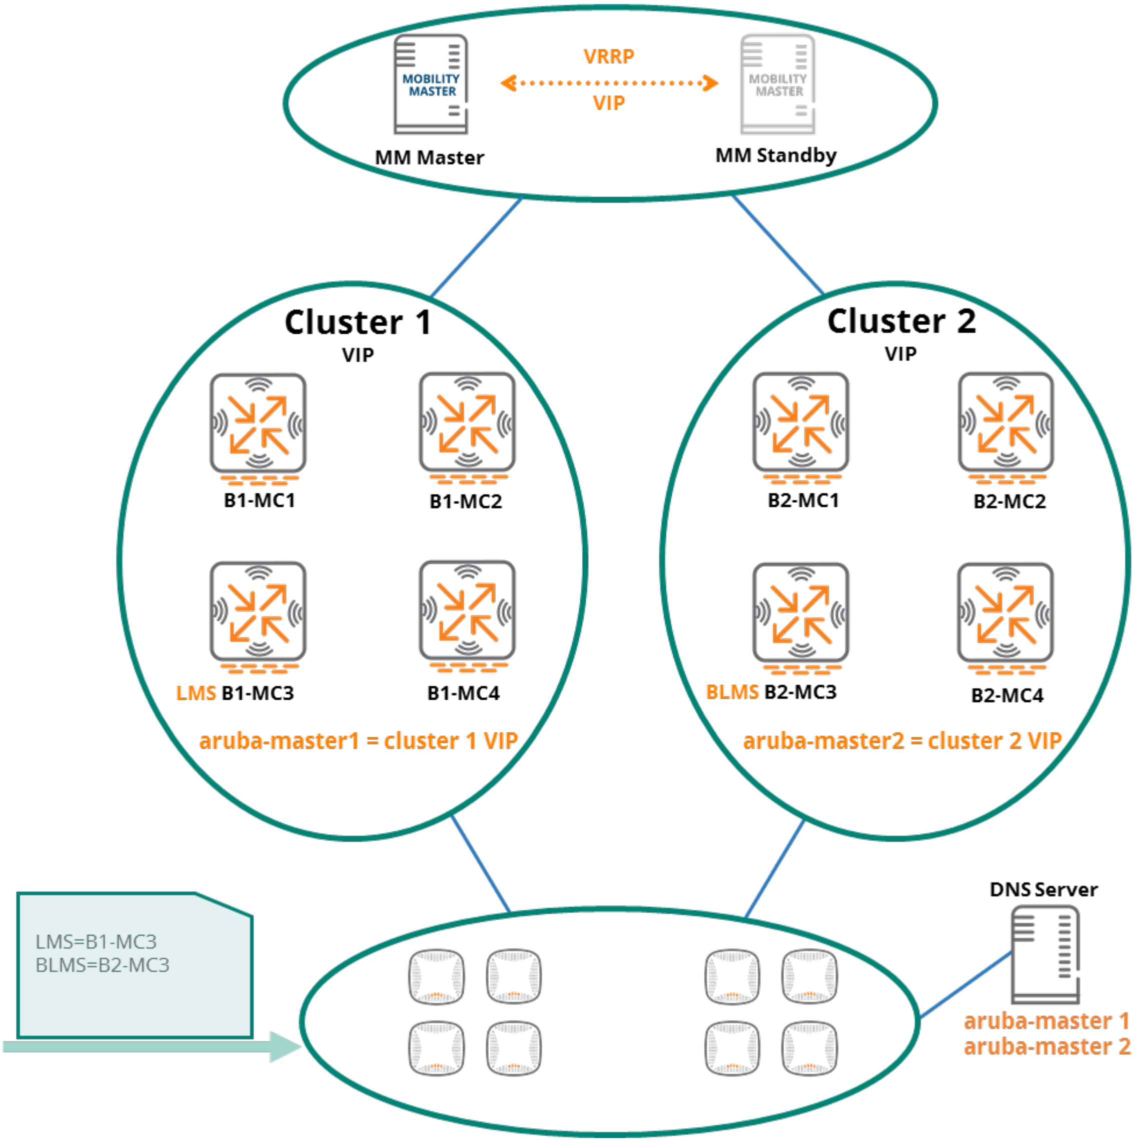

Figure 16 *LMS and BLMS Architecture with Clusters*

Remote APs

Remote Access Points or RAPs are purpose built APs for remote access use cases. Remote users typically work from home offices, small satellite offices, medium-sized branch offices, or on the road from hotels, hot spots, or customer locations. Each of these remote locations has different connectivity, capacity, and usage requirements.

IT organizations have traditionally served each category using a different remote network architecture. E.g., micro branches used a branch office router to interconnect an IP subnet at the remote site to the corporate network core, while telecommuters with only a single PC or laptop could be served with a software VPN client.

Aruba RAPs offer a solution for remote corporate users working from home or remote branches. Such users are granted access to the same wireless network they would access at the main corporate office from wherever they happen to be located.

Tunnel Mode

When RAPs operate in Tunnel Mode all traffic is tunneled back to the corporate network. There is wireless encryption on the client and controller as well as wired encryption on the RAP and controller. There is no access to local traffic e.g., a printer, home desktop, etc.

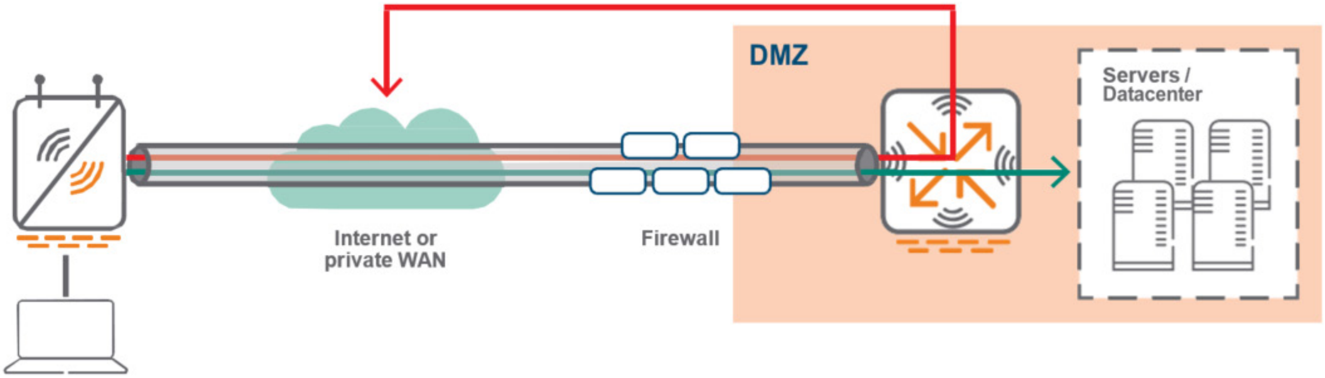

Figure 17 *RAP Tunnel Mode*

Split Tunnel Mode

Split Tunnel Mode allows non-corporate traffic to be bridged out locally to the Internet which reduces the bandwidth in the tunnel between the RAP and the controller that is transporting corporate traffic. In split-tunnel mode, there is wireless (L2) encryption and decryption on the client and RAP.

Corporate traffic is tunneled to the controller in the demilitarized zone (DMZ) and the rest of the corporate network. Traffic is encapsulated using GRE to preserve VLAN tags. The tunnel is trusted and shared by all Virtual Access Point (VAP) and wired ports. Traffic between the RAP and the controller is encrypted using IPsec.

Local traffic is source NATed (to enet0 address) and forwarded on both uplink and downlink wired interface ports according to user role and session access control list (ACL).

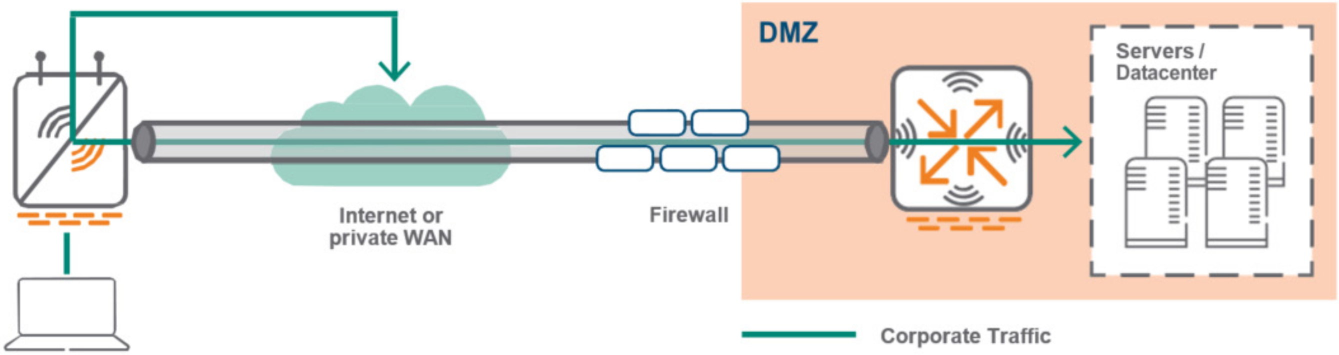

Figure 18 *RAP Split Tunnel Mode*

Bridge Mode

Bridge mode is mainly used on RAPs and Instant Access Points (IAPs). In bridge mode, there is no access to corporate traffic. There is a user traffic bridge to the local network on the AP uplink. Traffic is not sent to the controller. User VLANs have to exist on the edge of the network and authenticated traffic is tunneled to the controller. CPSec is required. DHCP, Network Address Translation (NAT), and Port Address Translation (PAT) are provided either by the RAP or an external router.

Bridge mode is typically used so that non-corporate devices such as printers or family owned devices can access the Internet directly via RAP uplink (similar to a home wireless router operation). This mode is not recommended for campus AP deployments as fewer features are supported in Bridge Mode.

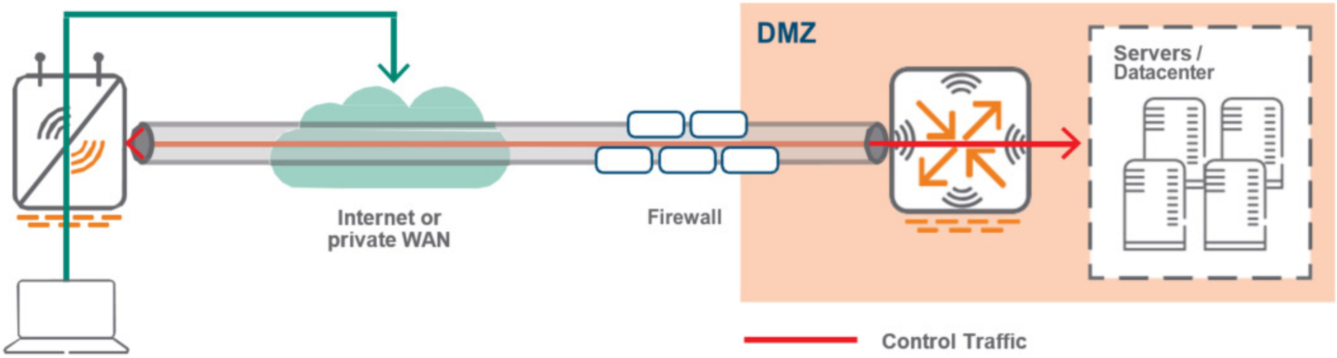

Figure 19 *RAP Bridge Mode*

Secure Jack

On any Aruba RAP that offers at least two Ethernet ports, the additional port can be configured for bridging or secure jack operation. This configuration provides maximum flexibility and allows for local wired access at remote sites. The additional Ethernet ports on a RAP can be configured for all the authentication types and forwarding modes available similar to a wireless service set identifier (SSID). A single SSID cannot be configured to provide 802.1X and MAC authentication simultaneously, however a wired port does not have the same limitation.

CAUTION

802.1X and MAC authentication should never be configured on the same SSID as doing so will cause client connectivity issues.

RAP Bootstrap Process

There are several different phases of the RAP bootstrapping process:

1. The RAP first obtains an IP address on the wired interface (Eth 0) by using DHCP. In remote deployment scenarios, the IP address is typically provided by the Internet service provider (ISP) when it is directly connected to the Internet
2. The RAP can be provided with a fully-qualified domain name (FQDN) or a static IP of the MC. If an FQDN is used, the RAP resolves the host name by using the DNS service provided by the ISP.
3. The RAP attempts to form an IPsec connection to the MC through the Ethernet interface. Depending on the provisioning type, either the RAP's certificate or Internet Key Exchange (IKE) Pre-shared Key (PSK) is used to complete phase 1 negotiation XAuth (an extension to IKE phase 1) is used to authenticate the RAP.
    - If IKE PSK is used then XAuth will authenticate the RAP with username and password.
    - If a certificate is used, XAuth authenticates the MAC address in the certificate against the RAP whitelist.
4. An IPsec security association (SA) is then established between the RAP and the controller
5. The MC provides the RAP with the IP addresses of the controller (LMS and BLMS IP) there it will be terminated
6. One or more IPsec encrypted GRE tunnels are formed between the RAP and the designated controller depending on the configuration

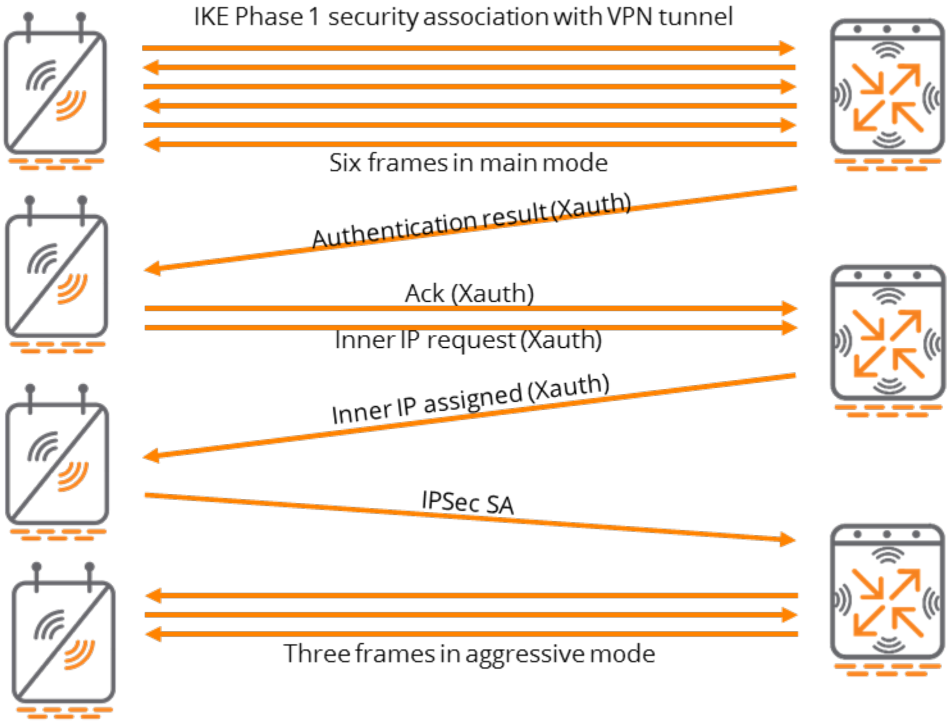
IKE Phase 1 security association with VPN tunnel
Six frames in main mode
Authentication result (Xauth)
Ack (Xauth)
Inner IP request (Xauth)
Inner IP assigned (Xauth)
IPSec SA
Three frames in aggressive mode

Figure 20 *RAP Bootstrapping*

Hierarchical Configuration

Hierarchical configuration was introduced in ArubaOS 8 to enhance the way configuration is applied in multi-controller networks.

ArubaOS 6 Configuration

In a typical ArubaOS 6 deployment where a Master controller is managing a set of Local controllers, each Local controller is brought up with its own base configuration (interfaces, VLANs, and IP addresses) the first time. Once the base configuration has been applied on each Local controller, the master then connects to the Locals and pushes configurations such as AP Groups, SSIDs, and user roles. Such an architecture involves numerous points of configuration for architectures with multiple controllers as configurations are not entirely centralized on the Master and a single Master can only manage a finite number of Locals.

ZTP was introduced in ArubaOS 6, however is only applicable for branch controller networks and only a subset of features can be enabled using the Smart Config interface on the Master.

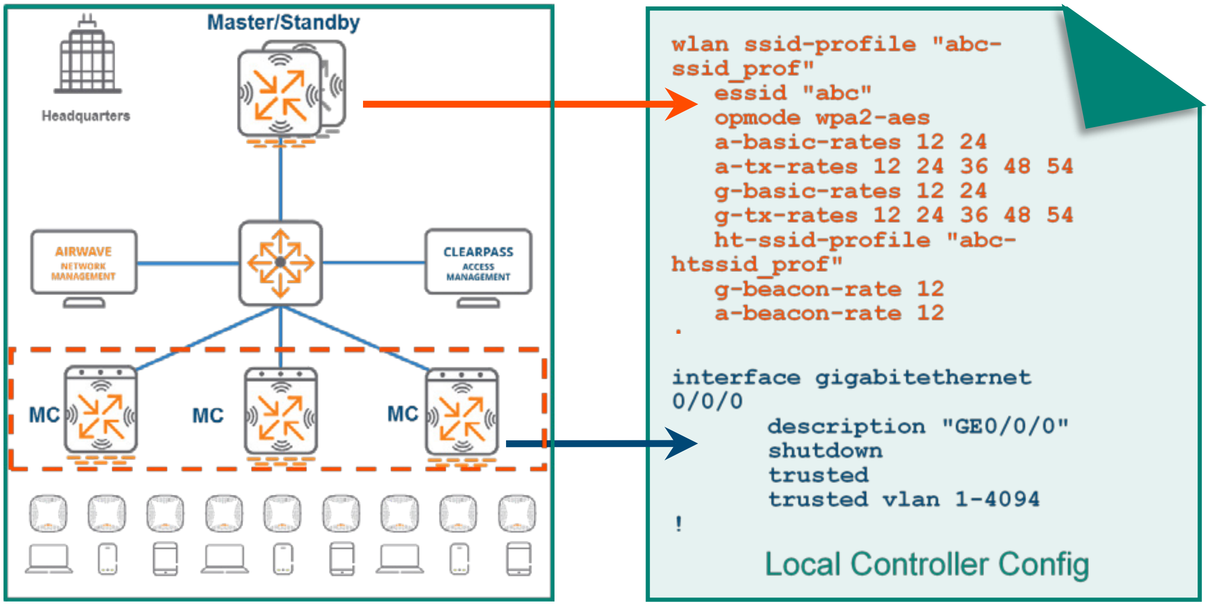

```
wlan ssid-profile "abc-
ssid_prof"
   essid "abc"
   opmode wpa2-aes
   a-basic-rates 12 24
   a-tx-rates 12 24 36 48 54
   g-basic-rates 12 24
   g-tx-rates 12 24 36 48 54
   ht-ssid-profile "abc-
htssid_prof"
   g-beacon-rate 12
   a-beacon-rate 12
.

interface gigabitethernet
0/0/0
       description "GE0/0/0"
       shutdown
       trusted
       trusted vlan 1-4094
!
```

Figure 21 *Typical AOS 6 Configuration*

The configuration pushed down from the Master is consistent across all of the Locals. The Master does not selectively push configurations for individual Local. If the Locals are geographically separated across multiple campuses and each campus requires a unique SSID, the configuration pushed to each Local will contain the SSID configurations for all the campuses. This level of SSID redundancy will likely be unnecessary for the majority of deployments. Even though the APs would subscribe to specific AP groups to broadcast the relevant SSIDs in each campus the configuration intended for other campuses is still irrelevant for the Local controllers where that SSID was not in use.

In some cases, a uniform configuration approach may also present an operational concern in the sense that the entire master configuration is exposed across all the regional locals. Since local network administrators will need access to the Master controller for configuration changes, a great amount of care needs to be taken to avoid any misconfigurations that may affect the remaining controllers being managed by the master.

AOS 8 Configuration Enhancements

ArubaOS 8 introduces true ZTP for all deployment modes as well as the concept of hierarchical configuration. New campus or branch controllers can discover the MM using DHCP options or Aruba Activate and receive their entire configuration from the MM. Regardless of the scale of controllers being managed the MM acts as a single touch point for the entire deployment.

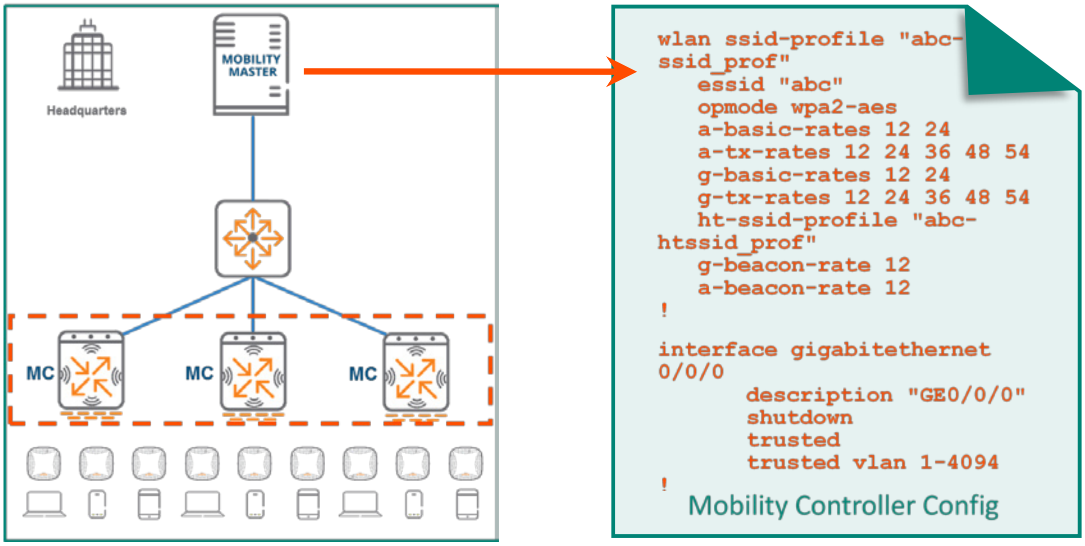

```
wlan ssid-profile "abc-
ssid_prof"
   essid "abc"
   opmode wpa2-aes
   a-basic-rates 12 24
   a-tx-rates 12 24 36 48 54
   g-basic-rates 12 24
   g-tx-rates 12 24 36 48 54
   ht-ssid-profile "abc-
htssid_prof"
   g-beacon-rate 12
   a-beacon-rate 12
!

interface gigabitethernet
0/0/0
       description "GE0/0/0"
       shutdown
       trusted
       trusted vlan 1-4094
!
```

Figure 22 *AOS 8 Configuration*

Hierarchical configuration allows configuration nodes to be created on the MM which contain common configurations for a particular region, campus, or building. Once a controller is whitelisted under a configuration node, a device-level configuration can be added on the device configuration node. When the MC contacts the MM for the first time, the group level configuration is merged with the device level configuration and then pushed down to the MC.

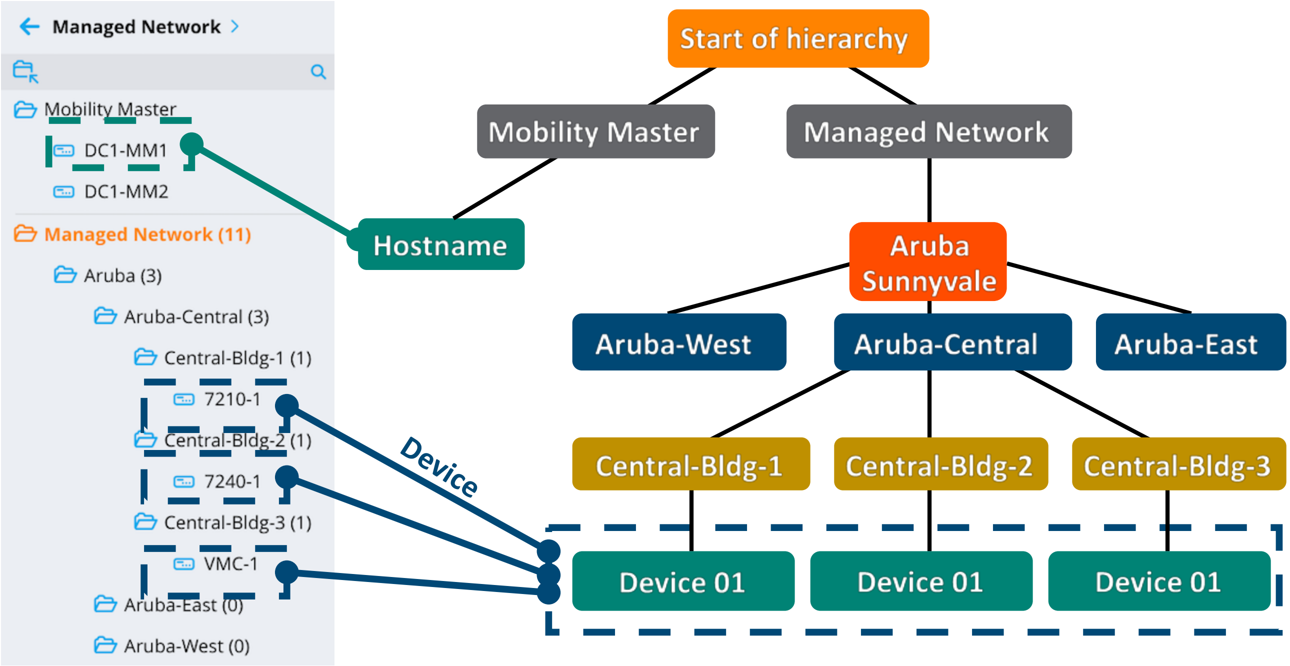

Figure 23 *Configuration Hierarchy*

The hierarchical configuration model has system-defined as well as user-defined configuration nodes.

System Nodes

System level nodes are present on GUI of the MM by default and cannot be deleted. The system nodes are as follows:

- **MM** – In the case of redundant MMs the configuration defined at this node is common for both active and standby MMs
- **Hostname (of MM)** – Holds configuration for the actual MM
- **Managed Network** – Hierarchy under which all the user-defined nodes are created and controllers are configured

User Nodes

User-defined nodes are created by administrators under the **Managed Network** system node. A node hierarchy can be created under this node where the upper nodes hold common configuration for all controllers. The configuration becomes more specific (based on region, campus, or building) at lower levels of the hierarchy. The device nodes are defined at the very bottom. The following examples demonstrate hierarchical group and device node definitions:

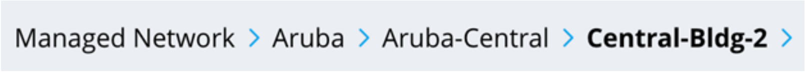

Managed Network > Aruba > Aruba-Central > **Central-Bldg-2** >

Figure 24 *Group Node Example*

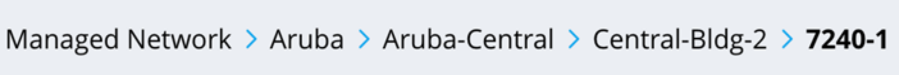

Managed Network > Aruba > Aruba-Central > Central-Bldg-2 > **7240-1**

Figure 25 *Device Node Example*

Up to four nested child nodes can be created under the **Managed Network** node. For example:

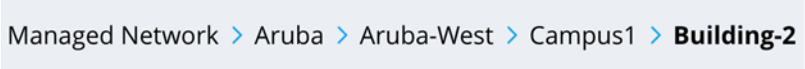

Managed Network > Aruba > Aruba-West > Campus1 > **Building-2**

Figure 26 *Four Nested Child Nodes*

Numerous child nodes can be created under the same parent node. In addition, child nodes can be freely moved to other nodes in the hierarchy as well as cloned from other nodes under the same parent node.

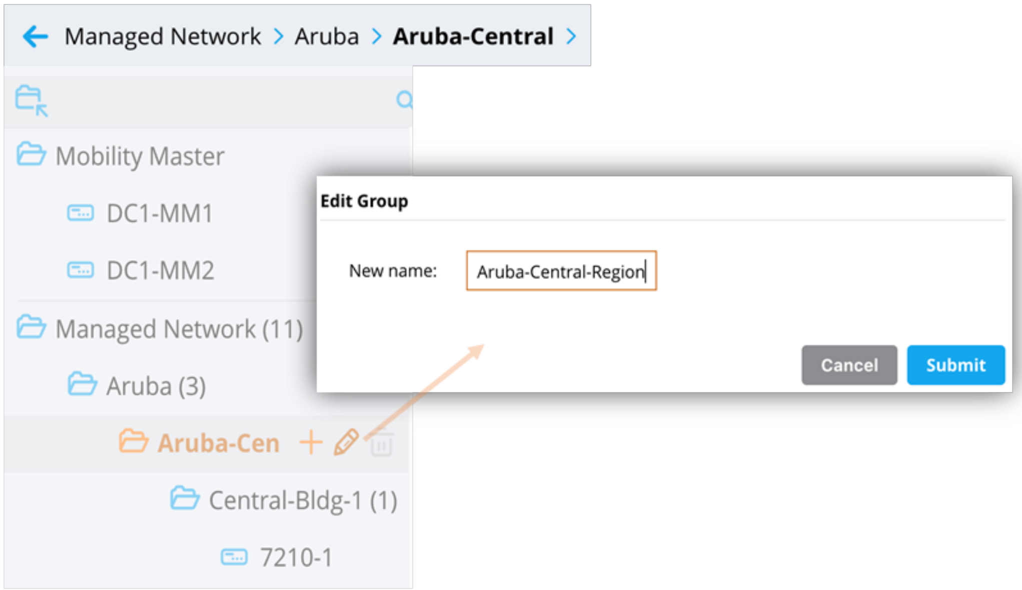

Managed Network > Aruba > **Aruba-Central** >

Edit Group

New name: Aruba-Central-Region

Figure 27 *Renaming a Node*

Configuration Inheritance

When an MC initially contacts the MM, it will merge the configuration at the device node with configurations from higher up in the hierarchy all the way up to the **Managed Network** node.

If there is a conflict or overlap in configuration on any node the configuration defined on the lower nodes will take precedence over the configuration on the higher nodes when pushing down the final configuration. The example below shows how a controller inherits its final configuration from the MM:

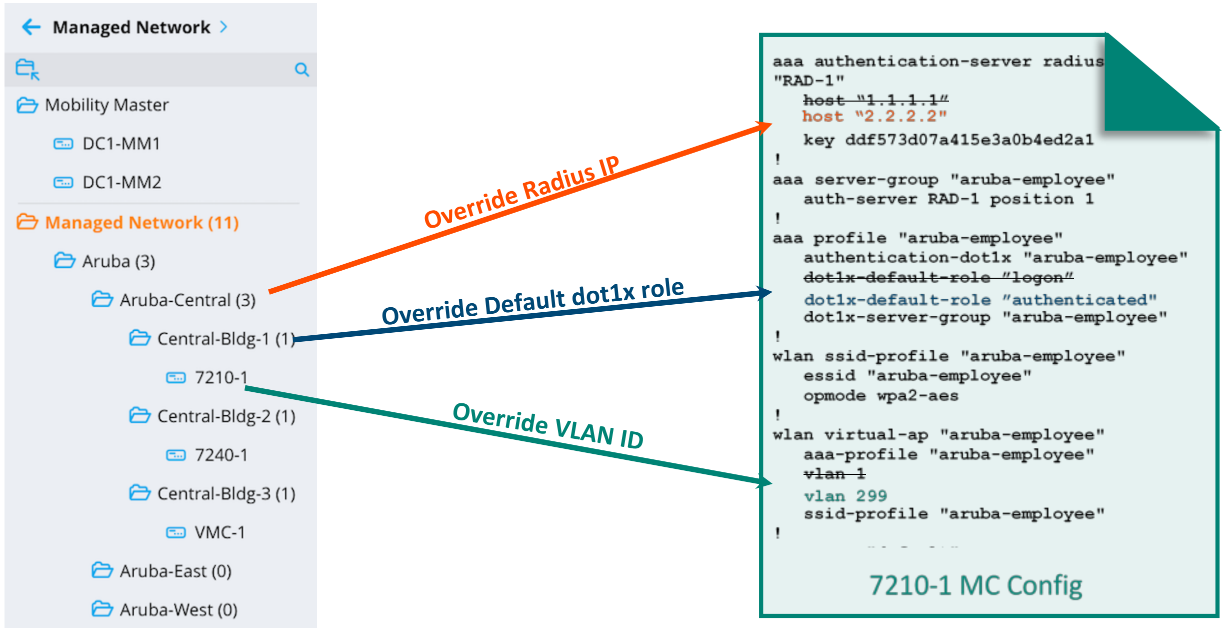

```
aaa authentication-server radius
"RAD-1"
   ~~host "1.1.1.1"~~
   host "2.2.2.2"
   key ddf573d07a415e3a0b4ed2a1
!
aaa server-group "aruba-employee"
   auth-server RAD-1 position 1
!
aaa profile "aruba-employee"
   authentication-dot1x "aruba-employee"
   ~~dot1x-default-role "logon"~~
   dot1x-default-role "authenticated"
   dot1x-server-group "aruba-employee"
!
wlan ssid-profile "aruba-employee"
   essid "aruba-employee"
   opmode wpa2-aes
!
wlan virtual-ap "aruba-employee"
   aaa-profile "aruba-employee"
   ~~vlan 1~~
   vlan 299
   ssid-profile "aruba-employee"
!
```

Figure 28 *MM Configuration Inheritance*

In the example above the initial configuration is created on the user-defined **Aruba** node which will be common to all the controllers in the organization. Since the **Aruba-Central** Node is farther down the hierarchy it will receive its initial configuration from the Aruba node and then override the RADIUS IP address. Similarly, the **Cenral-Bldg-1** node and the **7210-1** device node will override the dot1x role and the VLAN ID defined in the original configuration, respectively. The following figures display some of the key elements which were configured using the **Managed Network>Aruba** path:

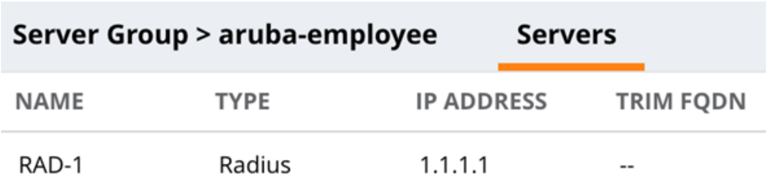

NAME	TYPE	IP ADDRESS	TRIM FQDN
RAD-1	Radius	1.1.1.1	--

Figure 29 *RADIUS Server RAD-1 with RADIUS IP "1.1.1.1"*

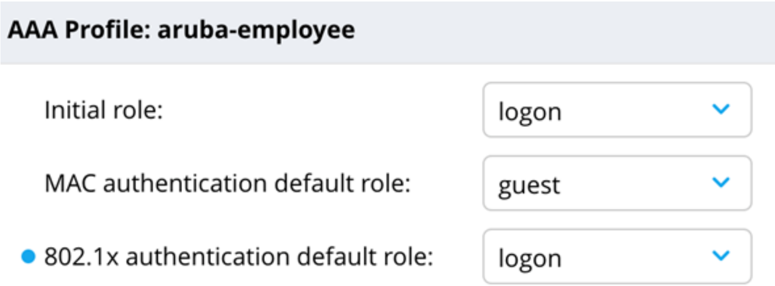

AAA Profile: aruba-employee

Initial role: logon

MAC authentication default role: guest

● 802.1x authentication default role: logon

Figure 30 *802.1X Default Role of Logon*

The presence of the blue dot next to a configuration parameter indicates the value was overridden such as in the case of a change to the configuration inherited from the parent node or a blank value that was replaced. Clicking on the blue dot displays additional details about the change and provides the option to either remove or retain the override.

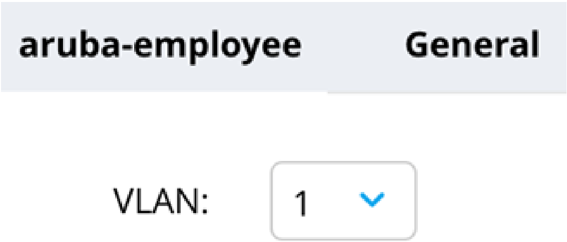

aruba-employee General

VLAN: 1

Figure 31 *VLAN "1"*

In the figure below the **Aruba-Central** node inherited this configuration, however the IP of the RADIUS server **RAD-1** was changed to "2.2.2.2":

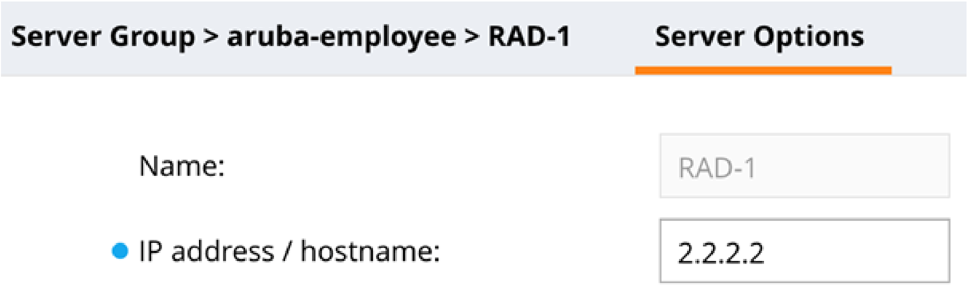

Server Group > aruba-employee > RAD-1 Server Options

Name: RAD-1

● IP address / hostname: 2.2.2.2

Figure 32 *RADIUS Server IP Override*

On the **Central-Bldg-1** node father down, the presence of the blue dot indicates that the default 802.1X role in the authentication, authorization, and accounting (AAA) profile was changed to "authenticated" and the configuration received from the parent node was overridden. **Managed Network>Aruba>Aruba-Central>Central-Bldg-1>**:

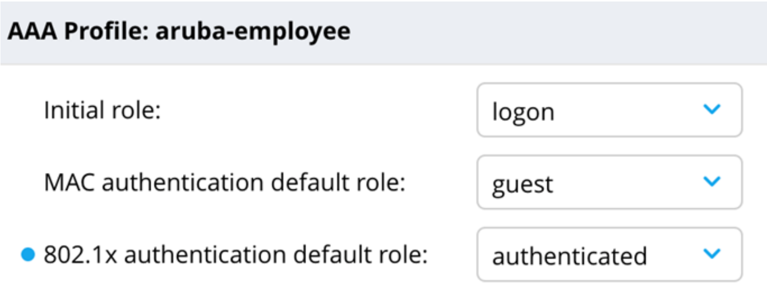

AAA Profile: aruba-employee

Initial role: logon

MAC authentication default role: guest

● 802.1x authentication default role: authenticated

Figure 33 *Authentication Default Role Override*

Lastly, the VLAN applied to the Virtual AP profile “aruba-employee” on the device node **7210-1** was changed to “299”. **Managed Network>Aruba>Aruba-Central>Central-Bldg-1>7210-1**:

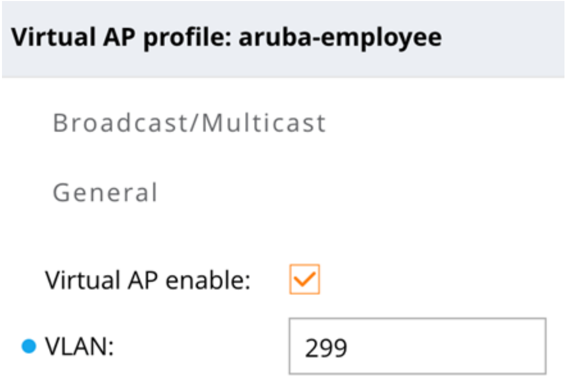

Virtual AP profile: aruba-employee

Broadcast/Multicast

General

Virtual AP enable: ☑

● VLAN: 299

Figure 34 *VLAN Changed to 299*

When the 7210 controller assigned to the **Central-Blg-1** node contacts the MM for the first time, its inherited configuration will result in the following changes:

	Original Configuration	Inherited Configuration
RADIUS Server IP	1.1.1.1	2.2.2.2
802.1X Default Role	logon	authenticated
VLAN	1	299

Table 6 *Summary of Inherited Configuration Changes*

Node Level Administration

Hierarchical configuration makes it possible to create node-level administration accounts on an MM. Network administrators can fully manage configuration for controllers at and below the configuration nodes that they have the necessary permission to access such as at a region, campus, or building level without affecting controllers elsewhere in the global hierarchy. This feature ensures that any undesirable configuration changes made at local sites are contained and do not affect the entire organization.

Proof-of-concept testing is another use case where custom ArubaOS builds and features need to be lab tested before bringing them into production. In such a scenario, test configuration nodes could be created along with node-level administration accounts. Since the nodes are created in a sandbox environment, testing may be freely performed without creating any undesirable effects higher up in the configuration hierarchy.

Licensing Pools

Licensing in ArubaOS 8 is managed centrally from the MM and the global license pool will be used by default for all controllers under its management. However, if specific license pools need to be dedicated, such as for a particular region, then custom license pools must be created on the MM with the appropriate hierarchical node and license counts definitions.

For additional information on please refer to the "Creating Licensing Pools on the MM" section.

Configuration Best Practices

Prior to deploying controllers there should be a defined plan on the configuration hierarchy for a network will look like. The following sections provides guidance for developing a configuration and deployment plan.

Node Hierarchy Design

There are multiple approaches to implementing a hierarchical design:

- A configuration hierarchy is typically created based on geographical segmentation of controllers. If an organization has multiple offices across a country then it makes sense to create configuration nodes for each region such as East, Central, and West. Each of these regions in turn may have multiple campuses, buildings, and devices which each have their own configuration node.
- An alternative way of organizing a hierarchy could be based on the type of services offered such as campus and remote with regional variations at the bottom of the tree.

Hierarchical configurations should be designed so that that configurations that are common to the organization reside on the higher level nodes. The rest of the configuration will be inherited by the lower nodes of the hierarchy as network requirements become more specific. E.g., a named VLAN can be defined at a higher level of the hierarchy and then assigned with specific VLAN IDs at the lower levels. Finally, configurations specific to individual controllers such as IP addresses, physical and virtual interfaces, and cluster membership are configured at the device level nodes. As a best practice, all configurations that are dependent on a single node should be always be defined e.g., defining VLAN ID and VLAN interface parameters together in a common node.

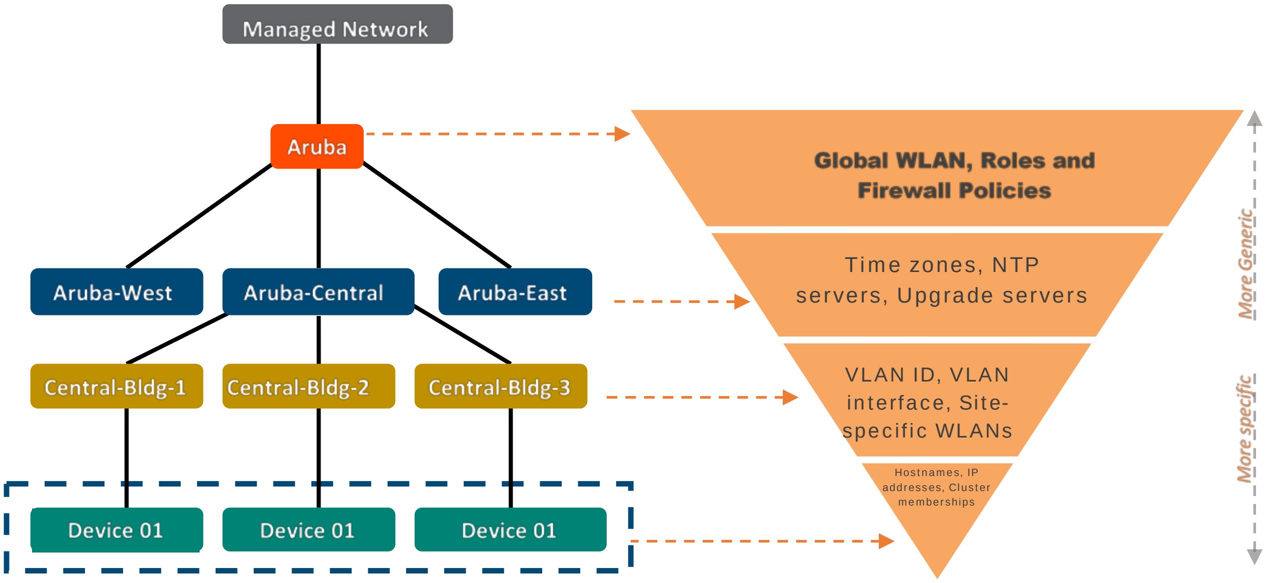

Managed Network
Aruba
Aruba-West
Aruba-Central
Aruba-East
Central-Bldg-1
Central-Bldg-2
Central-Bldg-3
Device 01
Device 01
Device 01
Global WLAN, Roles and Firewall Policies
Time zones, NTP servers, Upgrade servers
VLAN ID, VLAN interface, Site-specific WLANs
Hostnames, IP addresses, Cluster memberships
More Generic
More specific

Figure 35 *Node Hierarchy Design*

Configuration Overrides

Generally, a configuration that is inherited from higher level nodes cannot be deleted, however it can be overridden on the lower level nodes. However, there are certain configuration parameters that cannot be overridden at the lower level nodes. These parameters include the following:

- Net destinations
- IP access lists
- User roles
- AAA server groups
- AAA user derivation rules

Too many overrides across many hierarchy levels should be avoided as it can make troubleshooting challenging.

Depth of Hierarchy

Up to four nested child nodes can be created under the **Managed Network** node. However, it is recommended to create only as many nested nodes that are needed for purposes of configuration management simplification.

The Managed Network Node

As a best practice Aruba recommends defining configurations at a node below the **Managed Network** node and not on the **Managed Network** node itself. This is done to allow for sufficient network growth and scalability while simultaneously maintaining a separate configuration hierarchy for new sites. Configuration on the **Managed Network** node should be kept as minimal as possible in order to prevent the spread of issues related to misconfiguration across every other node in the hierarchy.

CAUTION

Aruba strongly discourages placing any configuration on the **/md (Managed Network)** node under any circumstances. Modifying configurations at this level will permanently alter the configuration for every child node without any ability to determine the default settings. Configurations should always begin a level below the **Managed Network** node.

Sites can be differentiated either physically or by type. In the example below, if the organization "Aruba" acquired another company "Network-Co", then we could simply define a new configuration node under **Managed Network** called **Network-Co** parallel to the **Aruba Sunnyvale** node.

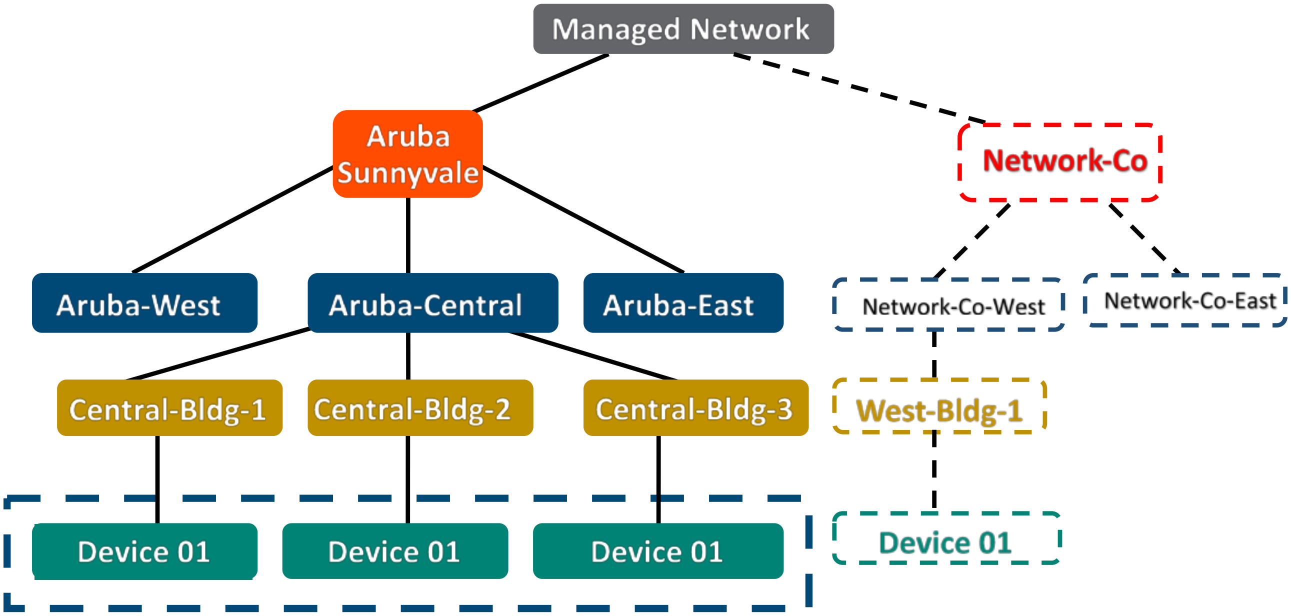

Figure 36 *Managed Network Node Hierarchy*

Configuration Notes

- When manually bringing up MCs it is important to ensure that they have been whitelisted on the MM under the appropriate configuration node
- When using ZTP to bring up MCs it is critical to ensure that the correct configuration node and MM MAC address are configured on Activate.

- Verify that the MM has learned about the MCs from Activate and whitelisted them under the configuration nodes that were specified in the Activate provisioning rule
- When specifying the MAC address of the MM for establishing an IPsec connection during initial configuration of controllers, always ensure that the **management port hardware MAC address** is used for a VMM and the **hardware MAC address** is used for an HMM
- When the MM registers with Activate, the correct MAC address is automatically populated. If controllers are using ZTP to contact Activate and register with the MM, identify the MAC address of the MM and select it from the dropdown list when configuring the provisioning rule on Activate

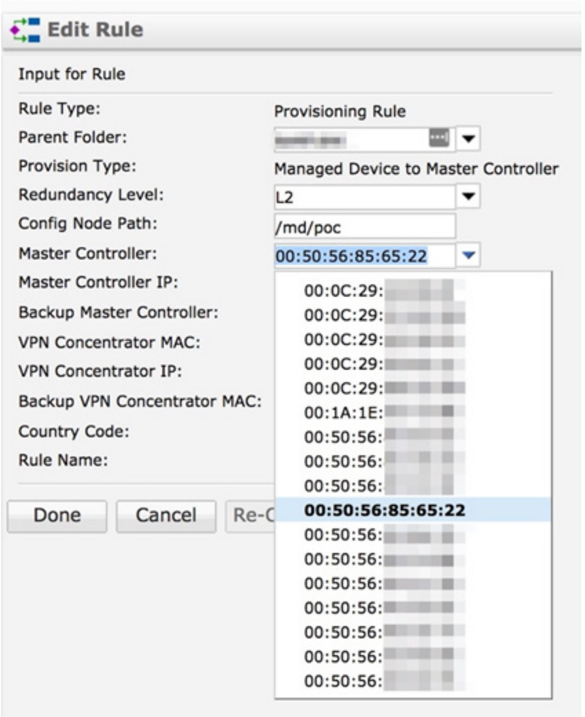

Edit Rule

Input for Rule

Rule Type:	Provisioning Rule
Parent Folder:	[redacted]
Provision Type:	Managed Device to Master Controller
Redundancy Level:	L2
Config Node Path:	/md/poc
Master Controller:	00:50:56:85:65:22
Master Controller IP:	
Backup Master Controller:	
VPN Concentrator MAC:	
VPN Concentrator IP:	
Backup VPN Concentrator MAC:	
Country Code:	
Rule Name:	

00:0C:29:[redacted]
00:0C:29:[redacted]
00:0C:29:[redacted]
00:0C:29:[redacted]
00:0C:29:[redacted]
00:1A:1E:[redacted]
00:50:56:[redacted]
00:50:56:[redacted]
00:50:56:[redacted]
00:50:56:85:65:22
00:50:56:[redacted]
00:50:56:[redacted]
00:50:56:[redacted]
00:50:56:[redacted]
00:50:56:[redacted]
00:50:56:[redacted]
00:50:56:[redacted]

Figure 37 *Selecting the MM MAC Address*

Loadable Services Module

Loadable Service Modules (LSMs) are a feature of ArubaOS 8 which allows administrators to dynamically upgrade or downgrade service modules on a running system without requiring a controller firmware upgrade or total system reboot. Each application has its own compressed image and upgrades are performed in real time without requiring a controller reboot.

Unified Communication and Collaboration

Unified Communication and Collaboration (UCC) is term that Aruba uses to describe the integration of real-time enterprise communication services such as instant messaging, voice, video conferencing, desktop-sharing, application sharing etc.

In the context of UCC as a feature on Aruba controllers, switches, and APs, it represents unification of various aspects of enterprise communication and collaboration applications. These aspects can be loosely categorized as media detection, media and traffic prioritization, monitoring and visibility, and media classification.

Aruba Controllers support the following UCC applications:

- Skype for Business
- Cisco Jabber
- Session Initiation Protocol (SIP)
- Wi-Fi Calling

NOTE

The list of application above is not comprehensive. Please refer to the ArubaOS User Guides for a full list of supported UCC applications.

UCC feature capabilities have not changed from ArubaOS 6 and UCC is not a new feature to ArubaOS 8. However, what has changed in ArubaOS 8 is the architecture of the feature and how it is deployed.

Architecture Comparison

UCC consists of a deep packet inspection (DPI) engine that runs on Local controllers in ArubaOS 6 and MCs ArubaOS 8. In ArubaOS 6, both the DPI as well as the UCC processes run on the Local controller itself. What has changed in ArubaOS 8 is that a portion of the UCC processes which Aruba refers to as the UCC service or UCC application has been moved to the MMM. The DPI functionality itself remains on the MCs.

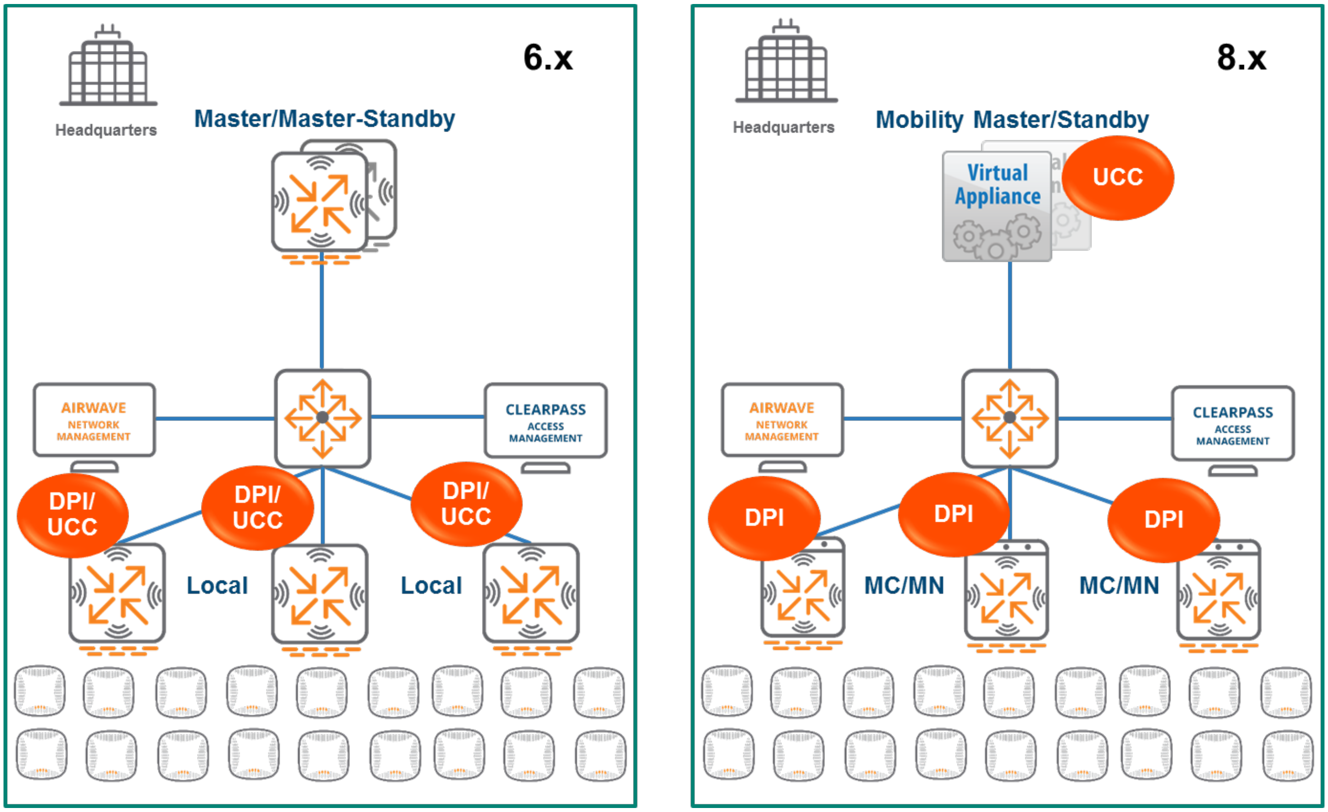

Figure 38 *UCC Architectural Comparison between ArubaOS 6 and ArubaOS 8*

New Features of UCC in ArubaOS 8

While UCC in ArubaOS 6 performs well there are a few drawbacks in its design that were improved for ArubaOS 8 which provide advantages for administrators considering a migration to ArubaOS 8:

- **Lack of Visibility** – In ArubaOS 6 UCC visibility is not centralized on the Master controller. Statistics and monitoring are maintained on the Local controllers. This design is not ideal as it requires users to log into each Local separately to monitor UCC data
- **Challenging Upgrades** – Adding support for new applications in ArubaOS 6 involves an entire controller upgrade which can be disruptive to the network
- **No SDN Aggregation** – Skype for Business Software Defined Network (SDN) API usage in ArubaOS 6 involves configuring the SDN Manager with the IP addresses of all subscribers in the network. Doing so has an adverse effect on network scalability

In contrast, ArubaOS 8 addresses the challenges listed above with improved functionality across the board through its superior architectural design and approach to UCC implementation. UCC now runs as an application (or a loadable service) on the MM. The DPI engine continues to run on the MCs which fulfill the same role as Local controllers in ArubaOS 6.

Classification and prioritization decision making functionality has been moved to the MM along with the VoIP application layer gateway which operates as part of the UCC application. In addition, the UCC feature can be upgraded independently without having to upgrade the all of the controllers in the network since it is one of the LSMs. This seamless upgradability allows administrators to add support for newer voice and UCC applications without experiencing any of the adverse effects associated with rebooting a controller.

The MM brings an important value proposition to enterprises using SfB as their UCC application. SfB SDN APIs can now be aggregated at the MM for all MCs. Doing so eliminates the need to configure the SfB SDN manager with thousands of IP addresses of individual subscribers. The MM keeps track of which MC where the call was initiated and matches the SDN API messages received from the SfB SDN manager to the call flows while programming the datapath on that particular MC. Employing the MM-based architecture introduced in ArubaOS 8 provides centralized visibility into UCC through the MM's GUI.

UCC Heuristics

In the context of UCC, heuristics is a method that the controller employs to detect and classify different types of media. It can be thought of as a form of advanced pattern matching for aspects such as packet size and ports used for flows. The UCC feature itself is comprised of two main processes: DPI and the Unified Communication Manager (UCM).

UCM is the name of the process that handles the UCC classification and programs prioritized flows in the datapath of the controller. The following steps outline how UCC in ArubaOS 8 handles a call flow when using heuristics:

1. A client initiates a call and the flow is analyzed by the DPI engine on the MC to detect the presence of media flows
2. The detection of media flows is passed from the MC to the UCC application on the MM
3. The UCC application on the MM classifies these media flows into categories such as voice, and video
4. The UCC application on the MM triggers an action on the MC to program the datapath to prioritize the flows
5. The prioritized flows are installed by the MC from client to server and server to the call recipient. The result is end-to-end prioritization for both upstream and downstream traffic

NOTE

The steps listed above assume preset communication channels between the MM and MC which exchange information related to the call such as trigger flow programming actions

The figure below illustrates the steps described above in an ArubaOS 8 architecture:

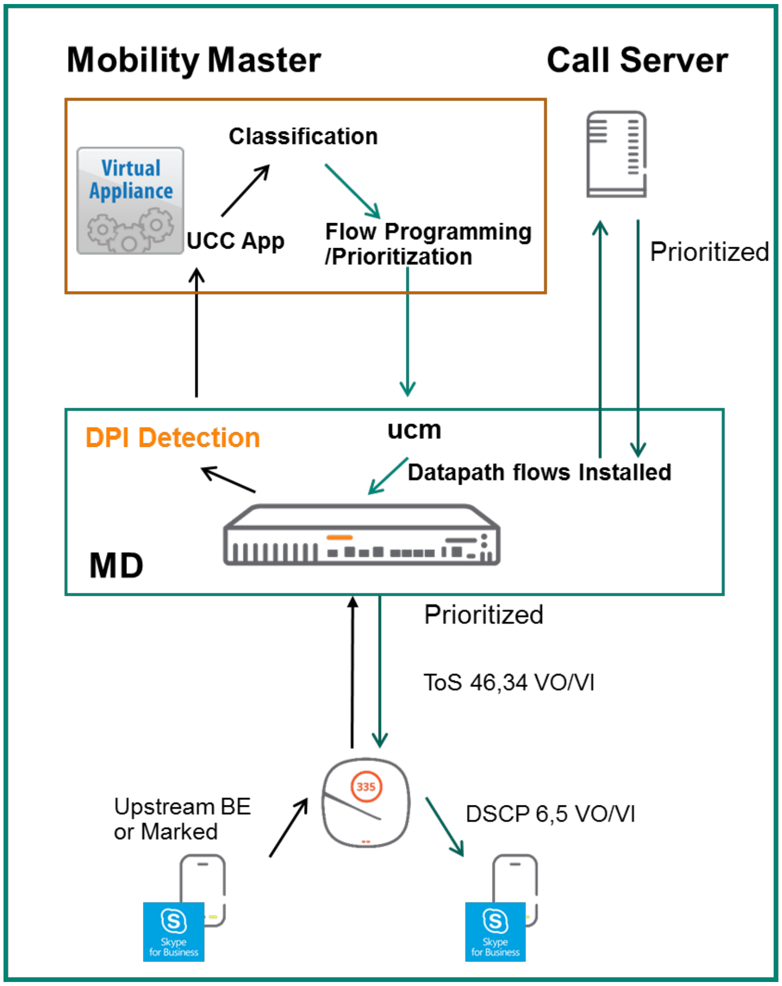

Figure 39 *ArubaOS 8 Call Flow with Heuristics*

Skype for Business

Microsoft has developed a service that provides detailed call information to switches which they call the Skype for Business Software Defined Networking Application Programming Interface. This tool was formerly known as the Lync SDN API.

The SfB SDN API consists 3 components:

- **SfB SDN Manager** - Resides next to the SfB front end server
- **SDN Dialog Listener** - Resides on the front end server
- **Subscriber** - In an Aruba architecture the subscriber would be the SfB SDN API-certified Aruba Controller (or switch)

NOTE

The SfB SDN API is not to be mistaken with the concept of SDN related to OpenFlow.

The subscriber (in this case the Aruba controller or switch) subscribes to the SDN API Manager which receives SDN API XML messages. These messages consist of XMLs containing detailed call information such as caller, recipient, port numbers, type of call, and endpoint client information.

SDN API Extensible Markup Language (XML) messages dramatically improve visibility into aspects of call quality as well as information that a controller is able to leverage for datapath flow prioritization. If SDN API XMLs used there is no need to engage heuristics for media classification as these flows provide such detailed information to the controller.

ArubaOS 6 SfB SDN API

In ArubaOS 6, the SfB SDN API is used to enable controllers to classify and prioritize media sessions by listening on a certain port over Hypertext Transfer Protocol (HTTP) or HTTP Secure (HTTPS). The SfB SDN Manager is in turn configured to send messages to the controller over the same port. The steps below outline the process for a call flow utilizing the SfB SDN API:

1. The client initiates a call. The SfB front end server triggers a session info XML and sends it to the SfB SDN Manager
2. The SfB SDN Manager sends the SDN API XML to the controller that has subscribed to it
3. DPI on the controller detects the presence of a media flows
4. The controller receives the session info XML which includes details such as caller, recipient, ports, MAC addresses, and media type
5. The UCM process matches the SDN API XML it receives with the DPI result of the sessions which were identified as media
6. The UCM programs flows in the datapath to prioritize traffic from client to call server and from server to call recipient
7. Upon call termination, the SDN manager sends a call end message to the controller which includes detailed statistics about call quality

The process of how the SfB SDN API work is represented in the figure below:

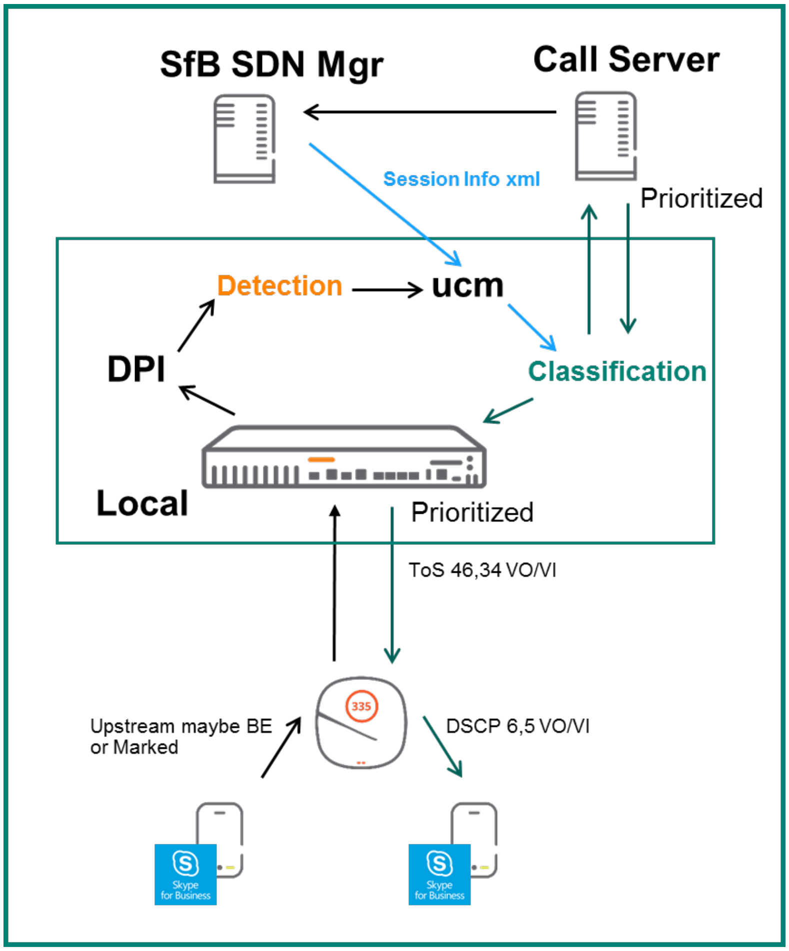

Figure 40 *ArubaOS 6 SfB SDN API Call Flow*

ArubaOS 8 SfB SDN API

SfB SDN API functionality in ArubaOS 8 is similar to ArubaOS 6 with the key difference that the SfB SDN Manager only needs to be configured with the subscriber IP address of the MM. The MM is configured to listen to SfB SDN API messages on port 32000 by default. In turn, the SfB SDN Manager sends messages over http or https to the MM using the same port.

The flow for a SfB call when using SfB SDN API with Aruba OS 8 is as follows:

1. The client initiates a call. Once the MC's DPI engine detects the presence of media flows it triggers a notification to the UCC application on the MM
2. In parallel, the SfB front end server will send a call session info XML message to the SfB SDN Manager which in turn forwards the message to the MM
3. The MM consumes the XML and correlates it with the MC that originally sent the DPI metadata
4. The UCC application on the MM uses SDN session info to classify and program datapath flows
5. The MM sends a flow programming action to the MC for the client which initiated the call

6. The MC programs the datapath accordingly and installs priority flows both upstream and downstream

With an ArubaOS 6-based architecture all UCC functionality resides on the Local controller. ArubaOS 8 differs in the sense that only the DPI resides on the MC while classification and prioritization decision making are moved to the MM. One of the key distinctions of UCC in ArubaOS 8 is being able to aggregate the SfB SDN API messages at the MM. The figure below illustrates the SfB SDN API process in ArubaOS 8:

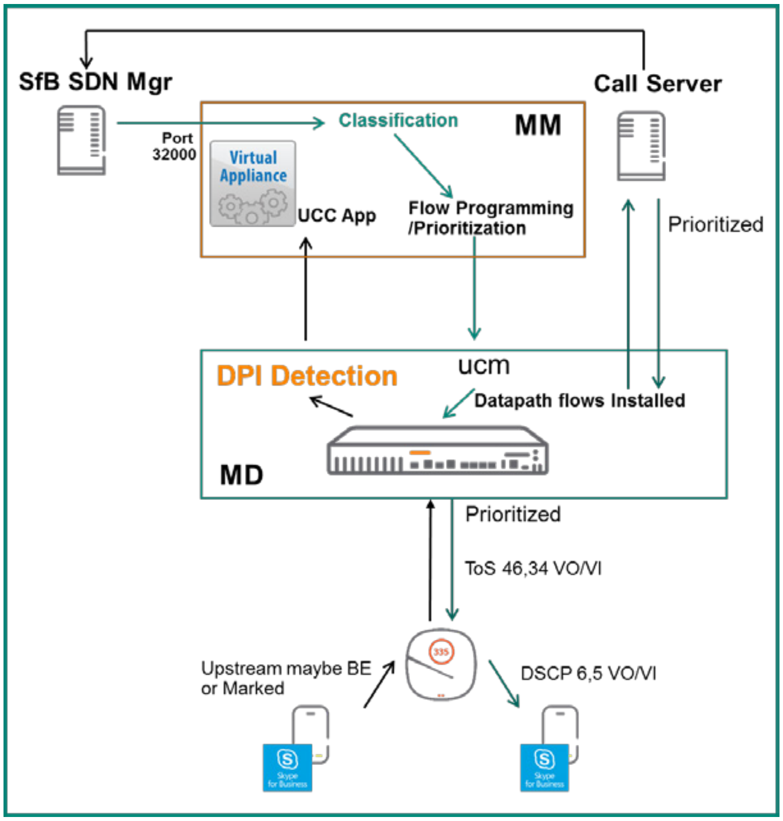

Figure 41 *ArubaOS 8 SfB SDN API Call Flow*

AirMatch

ARM in ArubaOS 6

Adaptive Radio Management (ARM) is the primary RF optimization technique used in ArubaOS 6. While ARM was a revolutionary technology when it was introduced it did suffer from a few shortcomings. Some of these included:

- Excessively frequent channel changes that lead to client disconnection and RF network instability, channel plan coupling among proximate radio neighbors,
- Uneven use of available channels
- Asymmetric EIRP planning adversely affecting client roaming behavior
- Lack of 2.4Ghz/5Ghz distinction in EIRP planning
- Lack of automatic bandwidth planning

These drawbacks have led some customers to abandon the ARM solution by either turning off the feature or by manually setting radio parameters through a long and tedious configuration process. ARM had the following characteristics in ArubaOS 6:

1. A decentralized service; each individual radio makes its own decision
2. ARM is “reactive” in nature
3. Future spectrum enhancements
4. Asymmetric EIRP planning which may not provide optimal client roaming behavior

When ARM was conceived the size of the networks was relatively small compared to a modern enterprise network and channel structures were much more basic. While it was necessary to have automation in RF planning, it was not as critical for network stability and performance as it is today. At the time it was considered acceptable practice to design a decentralized algorithm where each individual radio makes its own decision based on local information. Perennial convergence time, cascading effects, and mutual coupling or back-off were natural outcomes and regarded as minor issues. In modern production networks such occurrences are no longer acceptable and can pose significant challenges for larger, denser, and increasingly heterogeneous networks.

AirMatch was created to address all of the aforementioned challenges ARM was incapable of addressing. AirMatch is a centralized, clean slate RF optimization service. Information collection and configuration deployment paths are newly defined. The algorithm targets long-term network stability and performance in order to model and address RF challenges for the network as a whole.

AirMatch in ArubaOS 8

AirMatch provides unprecedented quality for RF network resource allocation. It collects data from the past 24 hours of RF network statistics and proactively optimizes the network for the next day.

As a best practice, the RF plan change should be deployed at the time of lowest network utilization so that client disconnects have a minimal impact on user experience. In addition to proactive channel planning done every 24 hours, AirMatch also reacts to dynamic changes in the RF environment such as radar and high noise events. AirMatch results in a stable network experience with greatly minimized channel and EIRP changes. AirMatch is defined by the following key attributes:

1. A centralized RF optimization service
2. Newly defined information collection and configuration deployment paths
3. Models and solves the network as a whole
4. Results in optimal channel, bandwidth, and EIRP plan for the network

NOTE

AirMatch only functions if the network is managed by an MM and is incompatible with an MCM architecture. In an MCM topology all channel, bandwidth, EIRP and other RF optimization decisions will continue to be made by ARM as they would be in an ArubaOS 6 architecture.

In the event that the link between the MM and MCs goes down and MM is unreachable there will be an impact to performance however AirMatch will still continue to function. Most notably the features which require the centralized coordination of the MM will be lost such as scheduled updates for RF optimization. The current RF solution will continue to function and changes resulting from high noise events and radar will still occur.

AirMatch Workflow

AirMatch is Aruba's next generation automatic RF planning service which assigns channel, bandwidth, and power to radios in the entire network. The AirMatch service runs on the Mobility Master and generates an RF solution which specifies new channel, bandwidth, and EIRP settings for each radio. The AirMatch workflow occurs using the following steps:

1. APs send RF statistics as AMON messages to MCs
2. The MCs forward the AMON messages to their MM
3. AirMatch calculates the optimal RF solution
4. The MM pushes the solution back down to the MCs
5. MCs send dot11 radio profiles to APs

AirMatch and ARM Comparison

The following table and images provide an overview of how AirMatch in ArubaOS 8 differs and improves upon the functionality provided by ARM in ArubaOS 6:

Feature	AirMatch	ARM
ArubaOS 8 Support	Mobility Master	Standalone or MCM
Computation	Centralized	Decentralized
High Noise Avoidance	Yes	Yes
Radar Avoidance	Yes	Yes
Optimization Scope	Entire RF network	Each AP
RF information Used	Past 24 Hours	Instantaneous snapshot

Table 7 *AirMatch and ARM Feature Comparison*

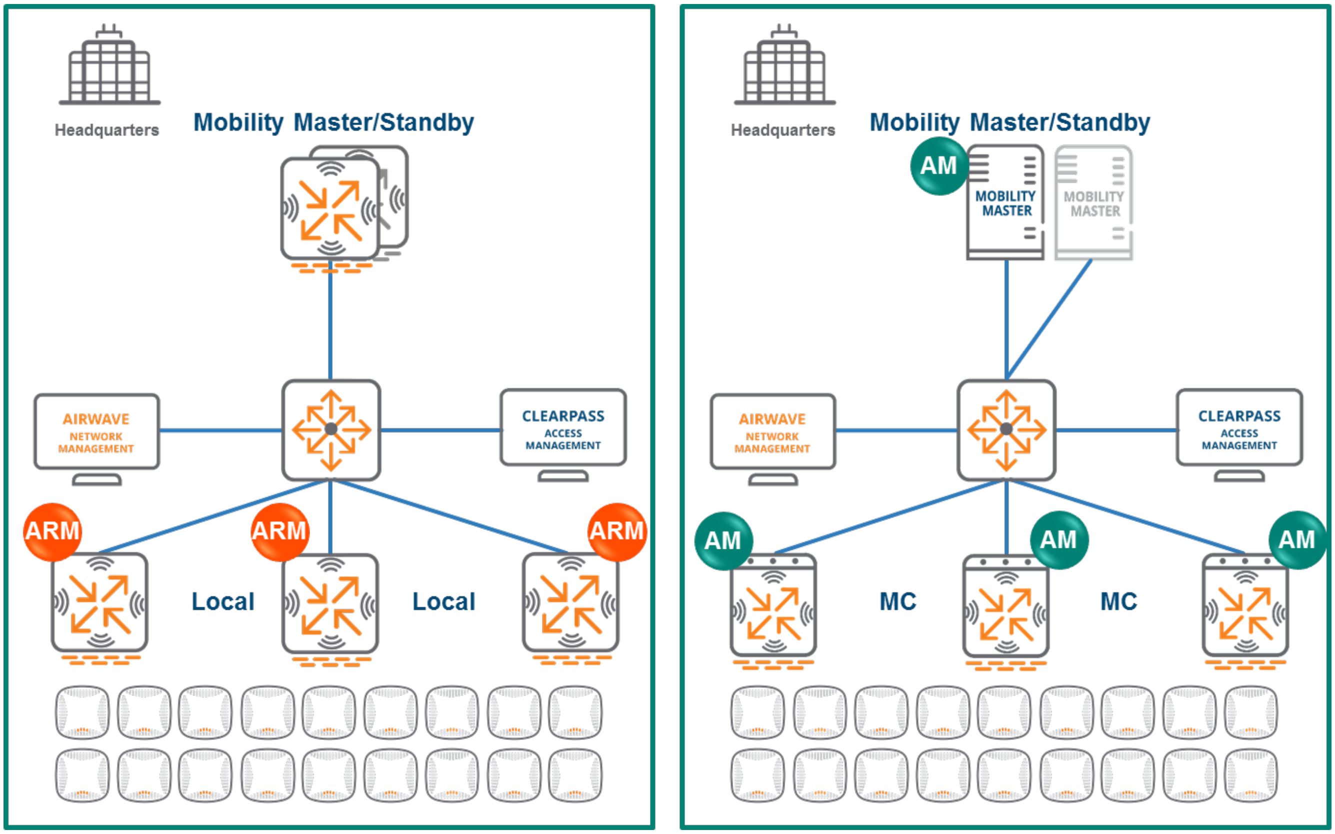

Figure 42 *ARM and AirMatch Comparison*

Web Content Classification

Web Content Classification (WebCC) is a feature on Aruba controllers and IAPs that was first introduced in ArubaOS 6. Its purpose is to classify http and https traffic according to category and reputation. Firewall rules can then be applied accordingly based on WebCC's classification. WebCC is able to prevent spyware and malware by blocking access to sites known to be dangerous as well as provide visibility into the web content categories and sites being accessed by users.

In an ArubaOS 6 deployment the WebCC process runs on the Local controllers. In ArubaOS 8 however the underlying architecture has been changed by moving WebCC process to the MM in form of an application or *loadable service*.

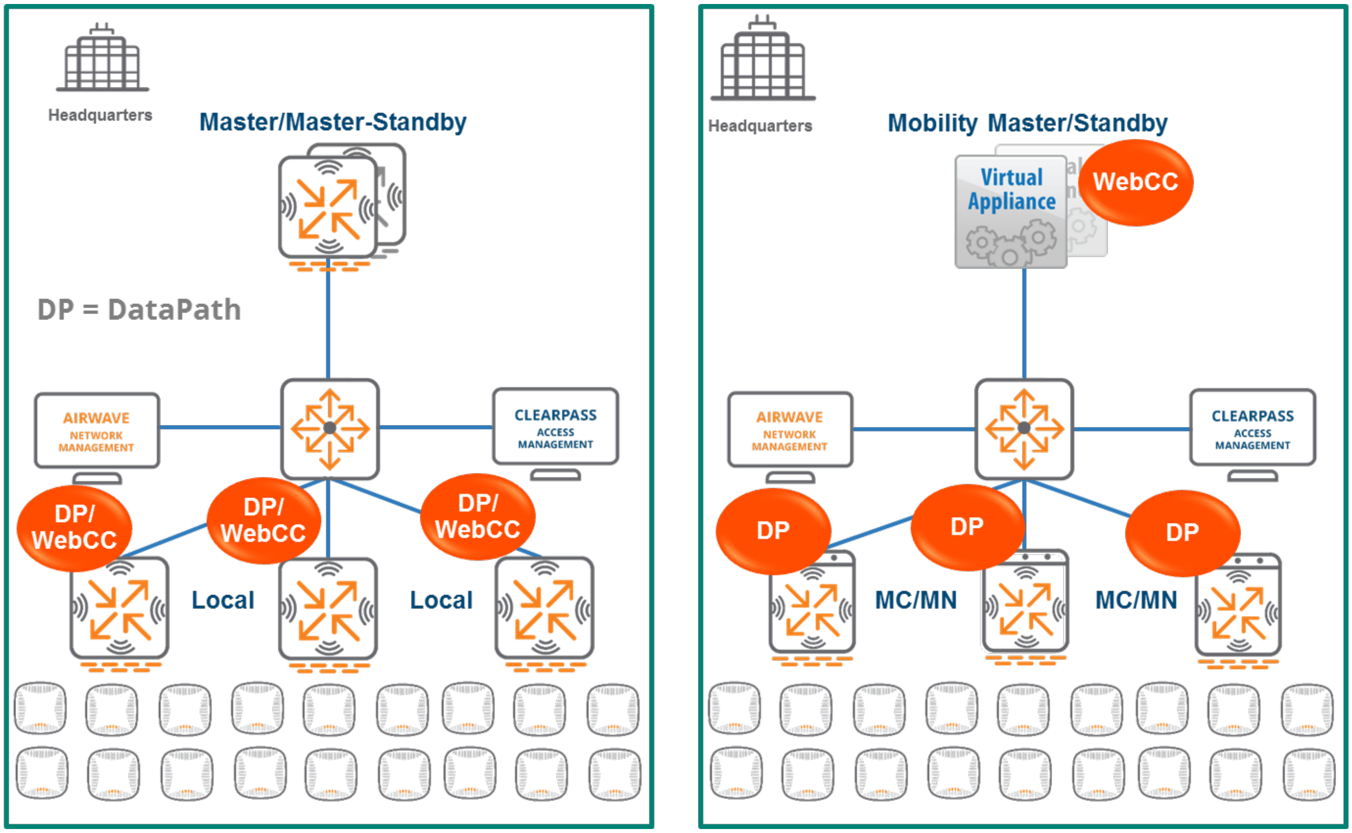

Figure 43 *WebCC Changes in ArubaOS 8*

WebCC in ArubaOS 6

WebCC runs as a process on Local controllers in ArubaOS 6 and works in conjunction with the datapath. Its primary role is to snoop http/https traffic in the datapath and inspect it to determine if further action is required. Once a client has IP connectivity and accesses a URL, the datapath intercepts http(s) traffic from client and checks against its local URL cache for a match. If datapath finds a match then classification and reputation rules are applied. The classification provided by WebCC is also used in a firewall access list which can take action to allow or deny access to the URL based on the additional information.

If the datapath cache does not find a match for the URL the client is attempting to access then a URL-miss trigger is sent to WebCC. The WebCC process then looks up the URL in the database maintained by the controller. If a match is found the URL will be classified and the information will be provided to the datapath. This information is then used in an ACL to either deny or allow access to the URL. If the WebCC process does not find a match in the URL database it performs a real-time cloud lookup from the Brightcloud repository and requests the URL's classification.

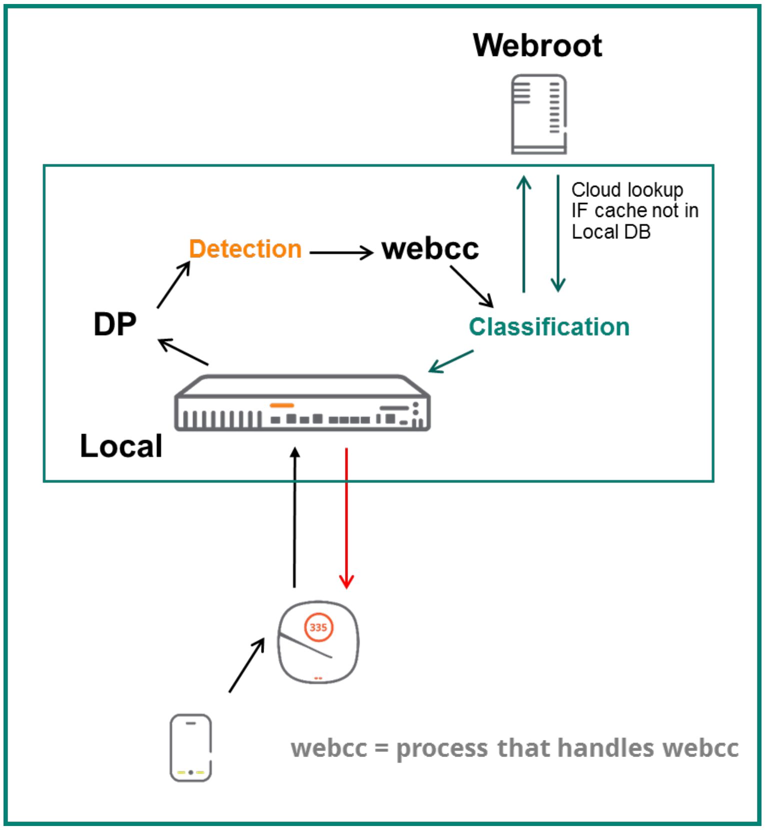

Figure 44 *WebCC Design in ArubaOS 6*

While WebCC in ArubaOS 6 offers significant advantages there are a few drawbacks with its design. Each controller maintains a URL database, however controller sizes vary as do their database and memory sizes. The web URL database that can be maintained on a controller is dependent on that controller's size. The probability of finding a URL in the local database decreases as controllers are reduced in size. This leads to increase in number of times a real-time cloud lookup is required for URL classification in the event of a URL miss. The time required to block a URL increases as well which allows users to access a URL which should have been blocked.

The other drawback of WebCC in ArubaOS 6 is each Local controller is required to individually contact Brightcloud. In addition, Local controllers have to consume memory and space to maintain their URL database.

WebCC in ArubaOS 8

The design changes to WebCC in ArubaOS 8 offer numerous advantages compared to the drawbacks inherent to the design of WebCC in ArubaOS 6:

- WebCC runs as a loadable service module MM
- MCs maintain only a shallow URL cache which saves memory

- The MM has larger memory and is capable of maintaining a URL database of up to 1 million records
- Cloud lookup for missed URL is performed solely by the MM

While the flow for WebCC in ArubaOS 8 looks similar to ArubaOS 6 there are some key distinctions to note. The most critical difference is that the WebCC process in ArubaOS 8 runs on the MM instead of the MC. The flow for the WebCC feature in ArubaOS 8 can be seen in the figure below:

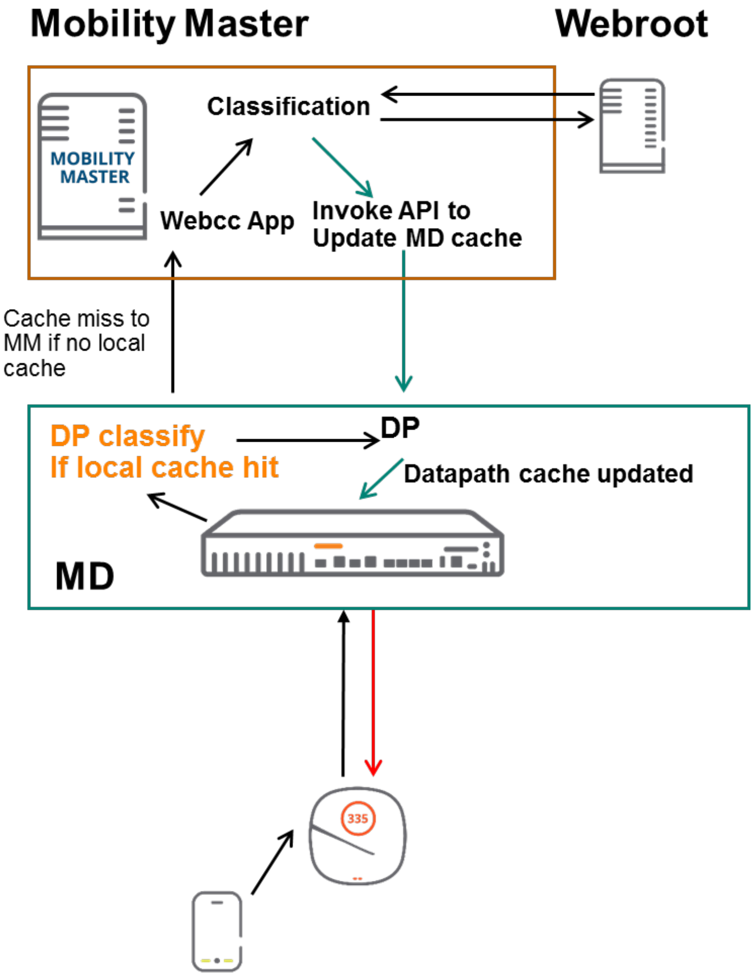
Mobility Master
Webroot
MOBILITY MASTER
Classification
Webcc App
Invoke API to Update MD cache
Cache miss to MM if no local cache
DP classify If local cache hit
DP
Datapath cache updated
MD
335

Figure 45 *WebCC in ArubaOS 8*

When a user attempts to access an http or https URL, the packet is snooped by the datapath on the MC which maintains a local URL cache. If the URL is found in the MC's local cache then classification is applied and further action can be taken to either allow or deny access based on any ACLs that have been configured.

If the MC's datapath does not locate the URL in its local cache then a URL-miss is triggered and sent to the MM. The MM looks for the URL in its database which is substantially larger than the local cache on the MC. If a match is found, the classification result is sent back down to the MC which will determine if action needs to be taken to restrict access or not based on existing ACLs. If the MM does not find a match in its local database then it will perform a cloud lookup through Webroot. The MM will update its local cache as well as the datapath of the MC which initiated the lookup request.

NOTE

The MM needs to be configured with a DNS server to be able to access Webroot over the Internet.

AirGroup

AirGroup is a component of ArubaOS that solves usability and performance issues related to the use of multicast Domain Name System (mDNS) services in enterprise and educational networks. Zero configuration networking services such as Bonjour and other mDNS services feature discovery, address assignment, and name resolution for desktop computers, mobile devices, and network services. They are designed for flat, single-subnet IP networks such as residential deployments. In large universities and enterprise networks it is common for Bonjour-capable devices to connect to the network across VLANs. As a result, user devices such as an iPad on a specific VLAN cannot discover the Apple TV that resides on another VLAN. In order to utilize the mDNS services on mobile devices in an enterprise environment, zero configuration networking multicasts need to be managed by a solution like AirGroup in order to improve network throughput, simplify connections to devices relevant to the user and location, and be properly forwarded across subnets.

AirGroup in ArubaOS 6

The mDNS protocol is designed to facilitate multicast communication and works extremely well within L2 boundaries. However, this means only mDNS-capable devices in the same VLAN are able to communicate with each other. I.e. an iPad in VLAN 10 cannot communicate with an Apple TV in VLAN 20.

Aruba created AirGroup to help facilitate communication for devices across VLANs and to provide filtering of peer-to-peer multicast traffic based on attributes including VLAN, user role, user name, user group, and location.

Each controller builds an mDNS cache table by learning and suppressing mDNS or Digital Living Network Alliance (DLNA) queries and advertisements over the air. E.g., whenever an iPad sends an AirPlay query the controller looks at its mDNS cache table and if an AirPlay service is available it responds to the iPad via unicast. Using unicast packets helps reduce channel utilization in the air.

AirGroup domains can be used for mDNS and DLNA communication between devices across different controllers. In addition, controllers are capable of integration with ClearPass to create Personal Area Networks. AirGroup servers can be defined on ClearPass and can optionally be shared along with usernames, user roles, user groups, AP group, AP name, and AP FQLN.

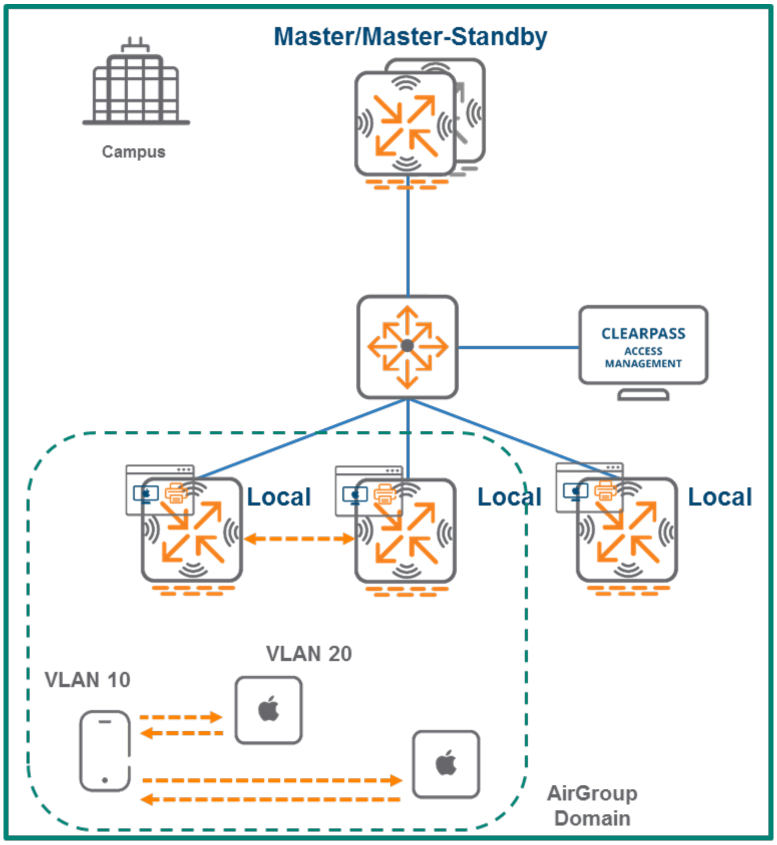

Figure 46 *AirGroup in ArubaOS 6*

AirGroup in ArubaOS 8

While AirGroup in ArubaOS 6 is capable of providing significant functionality enhancements it suffers from scalability limitations. ArubaOS 8 solves the platform scalability issues of AirGroup in ArubaOS 6 where scalability was limited by the platform capacities of controllers.

Unlike ArubaOS 6 where each controller runs AirGroup individually, in ArubaOS 8 the AirGroup functionality has been moved to the MM, i.e. the entire mDNS cache table resides on the MM. An OpenFlow controller is installed on the MM and OpenFlow agents are installed on MCs for communication of mDNS and DLNA information. Whenever the MCs intercept an mDNS or DLNA query or advertisement it will be forwarded to the MM via the OpenFlow channel.

The MM creates the appropriate mDNS/DLNA flows based on its AirGroup policies and pushes these flows to the MCs. The MCs either allow or deny mDNS and DLNA communication for the AirGroup devices on the WLAN.

AirGroup is significantly more scalable in ArubaOS 8 as the MM is equipped with appropriate resources to handle large amounts of mDNS communication on the network compared to a hardware-based controller in ArubaOS 6.

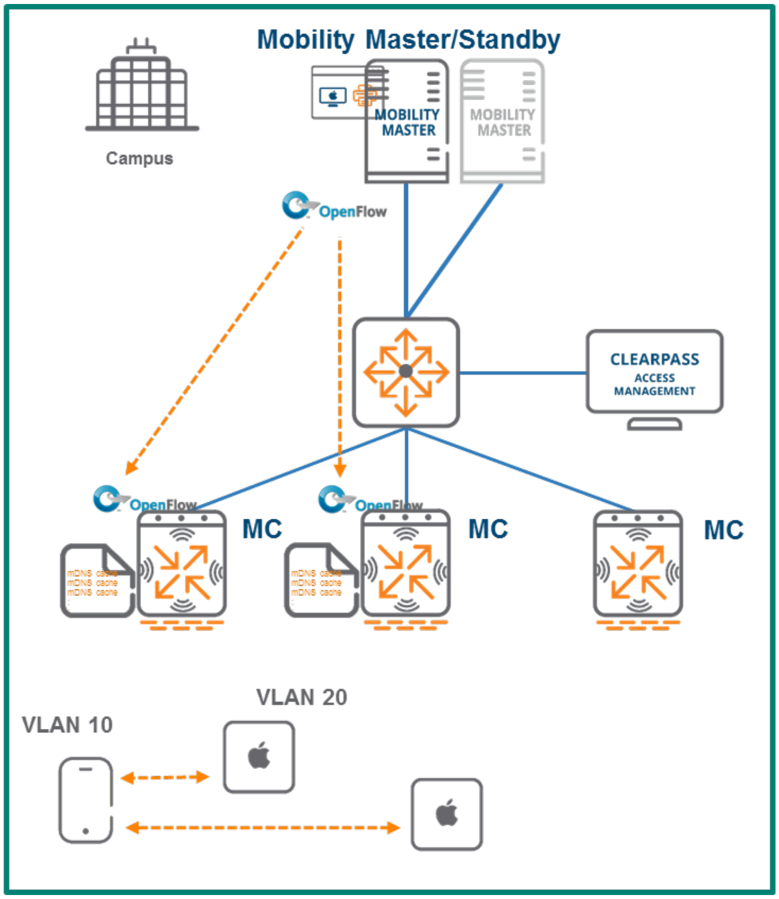

Figure 47 *AirGroup in ArubaOS 8*

AirGroup Feature Enhancements

Feature	Comments
AirGroup support for wired users	Wired users can now search for AirGroup services
AirGroup dashboard	Shows mDNS/DLNA traffic trends, server distribution, and user/server bandwidths
Ability to define number of hops	Ability share services with users up to 3 RF hops away from AP
Disallowed named VLAN	Allows to restrict AirGroup services across groups of VLAN IDs
Disallowed VLAN ID (for users)	Ability to define disallowed VLAN ID for users in addition to servers
Disallowed user-role (for servers)	Ability to define disallowed user-role for users in addition to servers

Table 8 *AirGroup Feature Enhancements*

AppRF

AppRF in ArubaOS 6

AppRF in ArubaOS 6 has the capability of identifying and applying policies to approximately 2000 applications including allowing, blocking, or rate limiting. Upgrading AppRF or adding new AppRF signatures in ArubaOS 6 requires a system-wide upgrade. E.g., even if new signatures only need to be tested on one of the Local controllers or a bug needs to be fixed in a Master-Local controller deployment, the Master Controller in the network needs to be upgraded along with all of the Locals. This limitation causes network disruption and requires scheduling a downtime for the entire network. In addition, ArubaOS 6 cannot create custom AppRF policies or custom application categories.

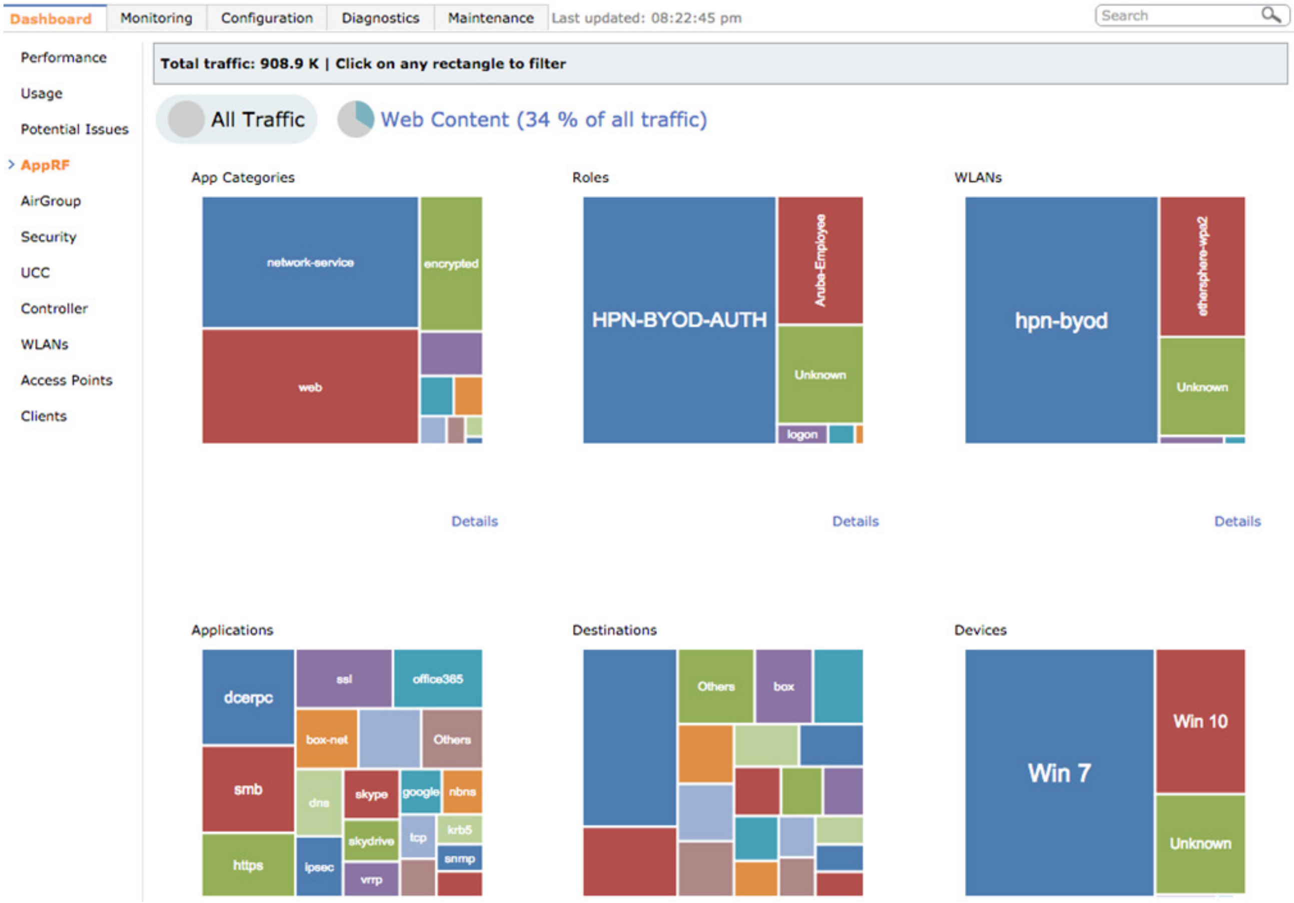

Total traffic: 908.9 K | Click on any rectangle to filter

Figure 48 *AppRF in ArubaOS 6*

AppRF in ArubaOS 8

ArubaOS 8 provides support for adding new applications to the controller without having to perform an upgrade. A proto bundle can be downloaded and activated at runtime to add support for new applications. DPI currently supports around 2,000 applications that can have rules applied to them. In ArubaOS 6, custom applications such as applications internal to the organization cannot be classified. ArubaOS 8 supports custom applications which can be pushed to MCs as desired.

New applications defined on the MM will be stored as application signatures in binary format and are delivered to MCs when configurations are pushed down. The application signature is then added to the active signature set on the MC providing the capability to support and define new applications as needed. The MM can configure up to 64 custom applications with 16 rules per application. Custom application categories can also be created and have policies applied to them. Even if an MC loses connectivity to its MM and standby MM it will not lose application classification functionality.

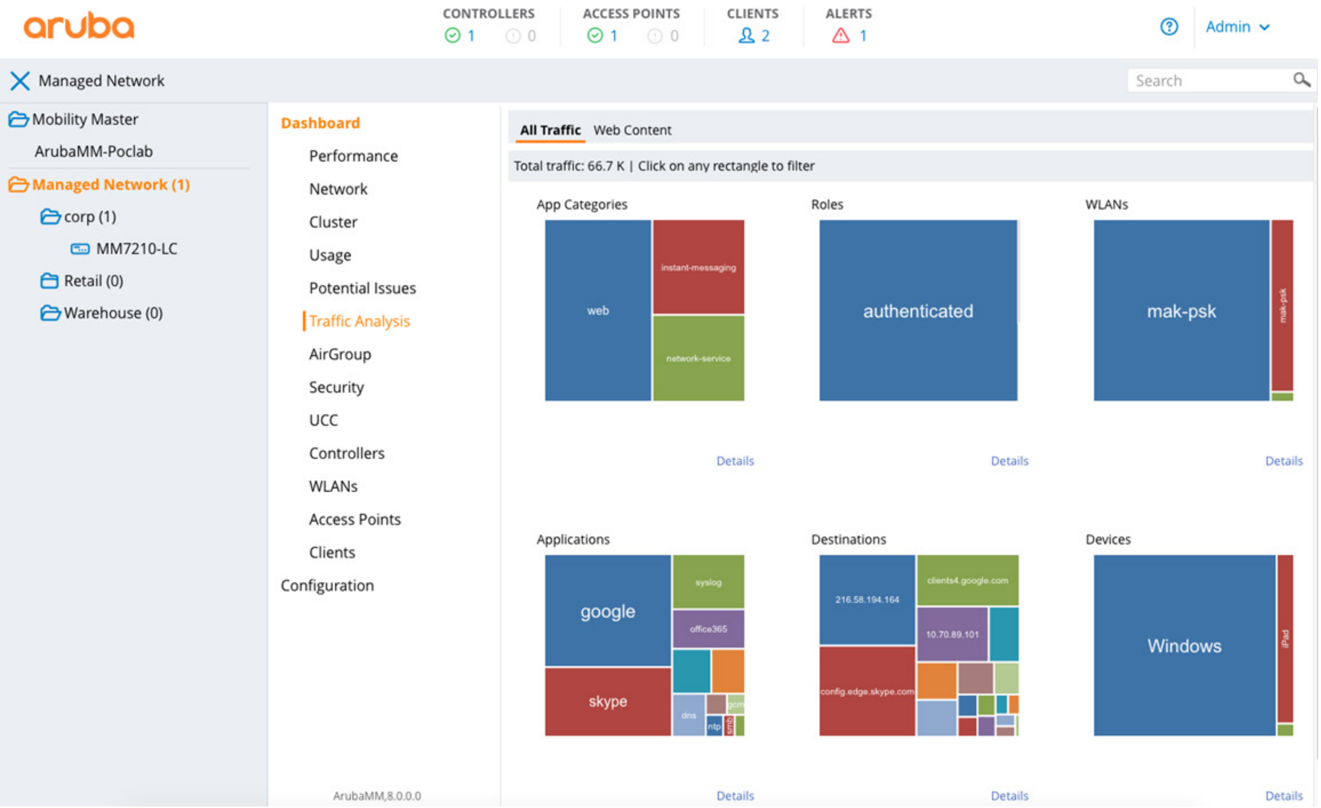

Figure 49 *AppRF in ArubaOS 8*

Application Programming Interface

In ArubaOS 8 there are three methods that can be used to automate configuration:

- Command Line Interface
- Graphical User Interface
- Application Programming Interface

ArubaOS 6 only allows configuration automation through the CLI and GUI. Unfortunately, being limited to these methods means that if the CLI output changes over time, the scripts also needed to be changed since not all outputs are generated using structured data. Modifying scripts every time there is a new code release can quickly become tedious. Similarly, the GUI is based on CLI and some GUI pages are hardcoded.

ArubaOS 8 introduces structured APIs that are based on the JavaScript Object Notation (JSON) model which uses GET and SET messages in a structured format for all configurations. Structured data means that all data is organized in a particular format where all elements that belong to a data type follow the same data model. This is achieved by separating schema from data. Schema (also called metadata) is a data model representation in JSON format which tells the user how interpret the data.

Data in the context of ArubaOS 8 is the representation of the configuration state of the MM in JSON format. The MM arranges the data in the same order as the schema so that it can be interpreted according to the schema instructions.

Configuration APIs

Configuration APIs use HTTPS to make GET and POST calls to the MM. GET APIs are used to learn about the configuration status of MCs and provide a similar output to show commands with the difference that they are in JSON format. POST APIs are used to configure MCs. E.g., the creation of new VLANs can be accomplished through POST APIs.

Configuration APIs may be appealing to service provider-oriented customers or customers who build their own large data systems using tools such as Elk Stack where they can use 3rd party APIs to interact with the entire network from a single point of configuration. A complete list of APIs is available at **https://x.x.x.x/api** where x.x.x.x represents the IP address of the controller.

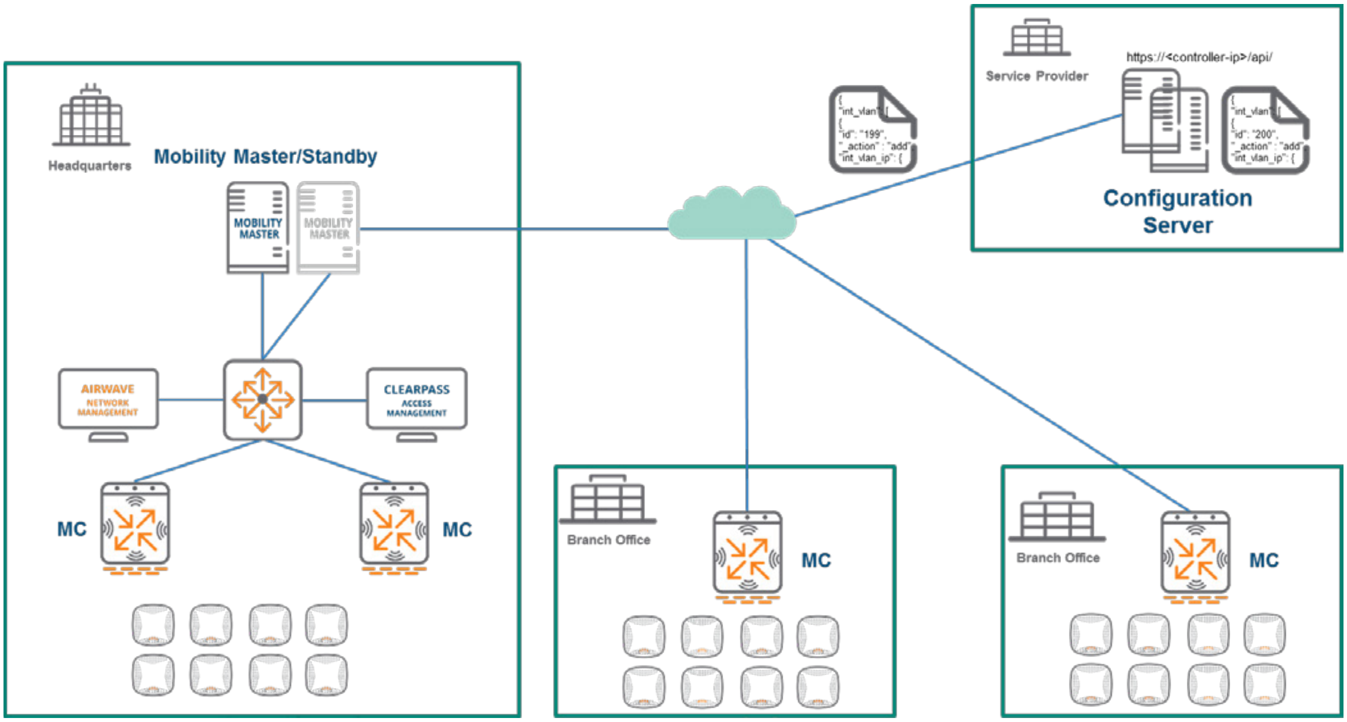
Headquarters
Mobility Master/Standby
MOBILITY MASTER
MOBILITY MASTER
AIRWAVE NETWORK MANAGEMENT
CLEARPASS ACCESS MANAGEMENT
MC
MC
{ "int_vlan": [{ "id": "199", "_action" : "add", "int_vlan_ip": {
Service Provider
https://<controller-ip>/api/
{ "int_vlan": [{ "id": "200", "_action" : "add", "int_vlan_ip": {
Configuration Server
Branch Office
MC
Branch Office
MC

Figure 50 *Configuration APIs in ArubaOS 8*

Context APIs

Context APIs are similar to the northbound APIs in Aruba's Analytics and Location Engine (ALE). Their main purpose is network analysis. The predefined APIs in ArubaOS 8 are listed below:

Context APIs

- Campus
- Building
- Floor
- Access_point
- VAP
- Station
- Radio
- Destination
- Application

From ALE

- Presence
- Location
- Geofence

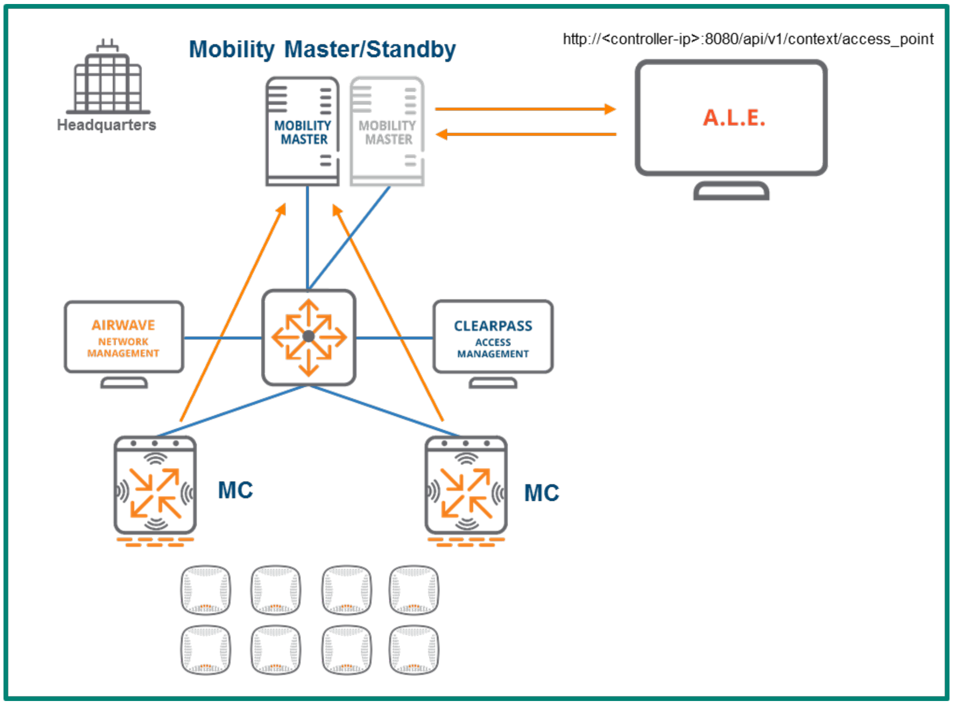
Headquarters
Mobility Master/Standby
MOBILITY MASTER
MOBILITY MASTER
http://<controller-ip>:8080/api/v1/context/access_point
A.L.E.
AIRWAVE NETWORK MANAGEMENT
CLEARPASS ACCESS MANAGEMENT
MC
MC

Figure 51 *Context APIs in ArubaOS 8*

Multizone

Multizone is a feature in ArubaOS 8 which allows IT organizations to have multiple separate and secure networks using the same AP. Historically, creating 2 secure networks in one physical location required separate APs. Multizone enables one AP to terminate two different SSIDs on two different controllers. The data is encrypted from the client all the way to the controller including when it flows through the AP. Multizone in ArubaOS 8 allows for completely secure network segregation and security even though the same AP is servicing traffic from multiple networks.

A *zone* is a collection of MCs under a single administration domain. The zone could consist of a Standalone controller or an MM and its associated MCs. A *Multizone AP* is an Aruba AP that is capable of terminating its tunnels on MCs residing in different zones. An ArubaOS 6 deployment where an AP terminates its tunnels on a single controller would be considered a single zone deployment.

Architecture

In a Multizone deployment a primary tunnel exists between the AP the primary zone. In the case of the below example the primary zone is an MM with a 3 node cluster of MCs. The Multizone profile downloaded to the AP from the primary zone. With the Multizone profile acquired, the AP learns the IP address of the data zone controller (Standalone) and establishes a data tunnel.

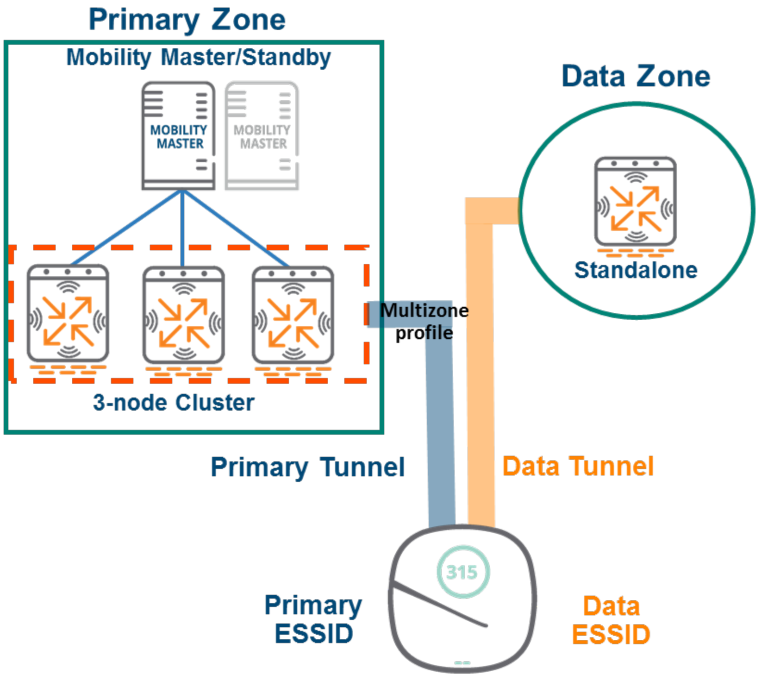

Figure 52 *Multizone Architecture*

Zone Roles

Multizone has several key objectives and abilities:

- Ability to leverage an existing AP deployment to broadcast SSIDs from different controller domains or zones
- Creating secure containers for different Basic Service Sets (BSSs) belonging to different organizations
- A wall is erected between zones where each administrative domain can only view and manage its own SSIDs

Primary Zone

- Zone that the AP connects to when booting up
- Retains full control of the AP management and configuration (AP, WLAN, and RF profiles)
- Zone where the Multizone profile is configured to enable the feature

Data Zone

- Secondary zone that an AP connects to after receiving the Multizone configuration from the primary zone
- Cannot reboot, upgrade, or provision the Multizone AP
- The only configuration allowed is the virtual AP configuration in Tunnel Mode

NOTE

The RFP license must be enable on the Primary Zone order to enable Multizone feature. No licenses are required on datazone controllers.

Key Considerations

1. MCs in all zones need to run the same ArubaOS version
2. The data zone should be using the same AP group and AP name used by the primary zone
3. The primary and data zone MCs cannot be managed from the same MM
4. There can only be a maximum of 5 zones: 1 primary zone and 4 data zones
5. There can only be a combined maximum of 12 controllers for all zones
6. The limit of 16 VAPs per radio still applies for all zones
7. Remote APs are not supported
8. All AP types are supported except for AP-9x

MM Redundancy

When designing mission critical networks, it is important to provide redundancy not only for the data plane but also for the management and control planes. In addition to the losing ability to push configurations, there are services which may be adversely affected in the absence of the MM. Having redundant MMs ensures that configuration and service-related tasks are protected and will continue to perform as needed at all times.

The following is a list of services which are impacted when the MM becomes unreachable:

- AirGroup operations (centralized mode only) and dashboard visibility
- UCC dashboard visibility
- Uncached WebCC lookups
- AirMatch recalibrations
- ClientMatch
- Configuration APIs
- Wireless intrusion detection and prevention

In ArubaOS 8 there are two types of redundancy which can be configured for an MM. Layer 2 redundancy addresses redundancy within a data center (DC). The active MM is responsible for managing all of the MCs in the network as well as any associated configuration and service-related tasks. The active MM is backed up by a standby MM using VRRP. If the active MM fails, the associated controllers will failover to the standby MM immediately upon detection of the failure. The standby MM then assumes the role of active MM.

Layer 3 redundancy addresses disaster recovery across Layer 3 separated networks and typically applies to DCs. In an ArubaOS 8 architecture this involves either one or a pair of MMs in each DC along with Layer 2 redundancy within the DC if using a pair. Within the context of Layer 3 redundancy one DC is referred to as the *primary DC* and the other as the *secondary DC*.

Regardless of which type of redundancy is being used licenses are configured on the active MM. These licenses will be automatically synced to the L2 and L3 redundant MMs.

Layer 2 Redundancy

Aruba MMs rely on VRRP as their layer 2 redundancy mechanism. The entire configuration hierarchy is automatically synced from the active MM to the standby MM with the exception of any configurations under the device configuration node of the active MM.

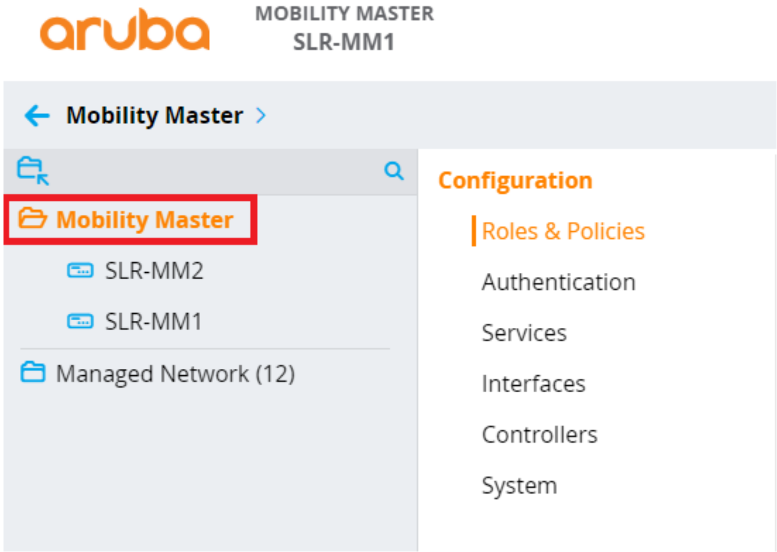

Figure 53 *Configuration for All MMs*

Configurations that are common to both the active and standby MMs are placed under the **MM** node so that they will be synced. Configurations specific to the Active MM (e.g. IP addresses and VRRP) must be placed individually on its own specific device node. MM services and managed devices cannot be configured from the Standby MM.

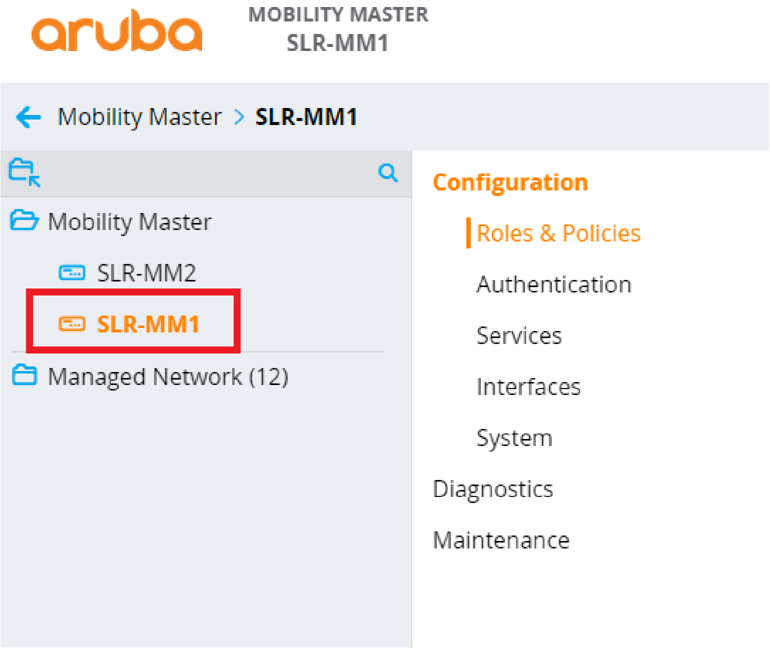

Figure 54 *Configuration Specific to the Active MM*

Topology

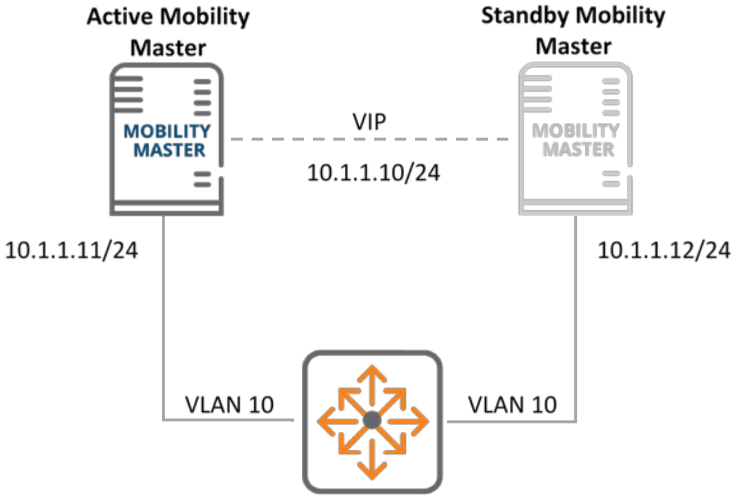

Figure 55 *Layer 2 MM Redundancy*

Synchronization

As soon as VRRP and master redundancy are configured between the active and standby MMs the entire configuration hierarchy is synced. Some of the databases that are synced to the standby MM at periodic intervals when database synchronization is configured on the active MM are listed below:

- WMS Database
- Local User Database
- Global AP Database
- AirGroup Database
- License Database
- CPSec Database

The synchronization interval is specified as part of the database synchronization configuration. The database synchronization interval is configurable. As a best practice Aruba recommends configuring the interval for no less than 20 minutes. Configuring a more frequent interval can add a substantial amount of network overhead.

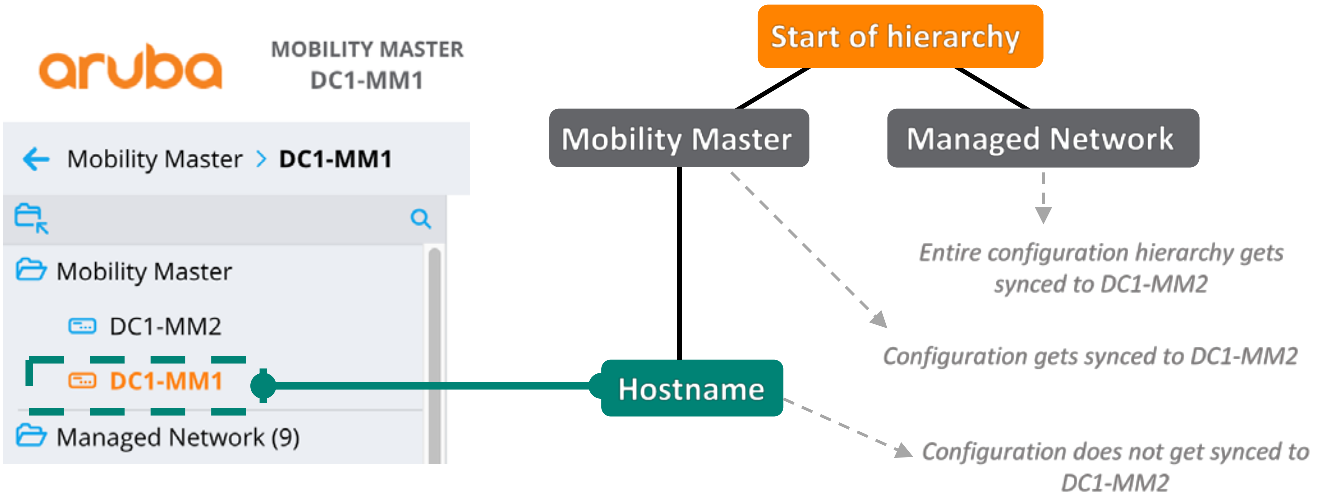
aruba
MOBILITY MASTER
DC1-MM1
Mobility Master > DC1-MM1
Mobility Master
DC1-MM2
DC1-MM1
Managed Network (9)
Start of hierarchy
Mobility Master
Managed Network
Hostname
Entire configuration hierarchy gets synced to DC1-MM2
Configuration gets synced to DC1-MM2
Configuration does not get synced to DC1-MM2

Figure 56 *Database Synchronization*

Once the active and standby MMs have performed their initial synchronization and reached a stable state, any incremental configuration change that is committed and saved on the active MM results in a configuration sync with the standby MM.

The exception to this behavior is any change made on the device configuration node of the active MM (i.e. /mynode). These changes are not synced to the standby MM. The standby MM contains its own version of the device configuration so any desired changes must be made directly on its corresponding device configuration node (i.e. /mynode on the standby MM). Configuration changes for other nodes in the hierarchy are not permitted on the standby MM.

MC Failover

MCs communicate with their active MM using the Virtual Internet Protocol (VIP) address of the VRRP instance shared by the MM pair. The active MM is master of the VRRP instance and sends out VRRP advertisements every second by default. The standby MM monitors these advertisements to ensure that the VRRP instance master is still functioning properly. If the standby MM fails to receive VRRP advertisements from the master such as in the event of a controller failure or reboot the standby MM will wait until three consecutive advertisements are missed, after which it promotes itself to be the master of that VRRP instance. The MCs continue communicating with the Virtual IP of their MM. The only impact that will affect the MCs is the time it will take them to establish IPsec sessions with the Standby MM which has become active due to the change in VIP ownership. The MCs continue to correspond with the owner of the VIP of the VRRP instance between the MMs. The actual device that owns that VIP at any particular moment and therefore is acting as the active MM is irrelevant to the MCs.

Layer 3 Redundancy

Layer 2 MM redundancy works well for single data center topologies where the active MM is supported by a standby. However, in the event of a data center power or network outage the MCs could potentially lose connectivity to both MMs which would result in a loss of functionality. Layer 3 MM redundancy was introduced to prevent such a scenario from occurring. It involves a primary MM or MM pair backed up by a secondary MM or MM pair over a Layer 3 connection to provide service continuity for controllers if the primary DC goes down. While Layer 2 redundancy can be thought of as redundancy between MMs within a data center, layer 3 redundancy can be thought of as redundancy between data centers.

Topologies

Below are two common examples of topologies that have been configured with layer 3 redundancy between MMs.

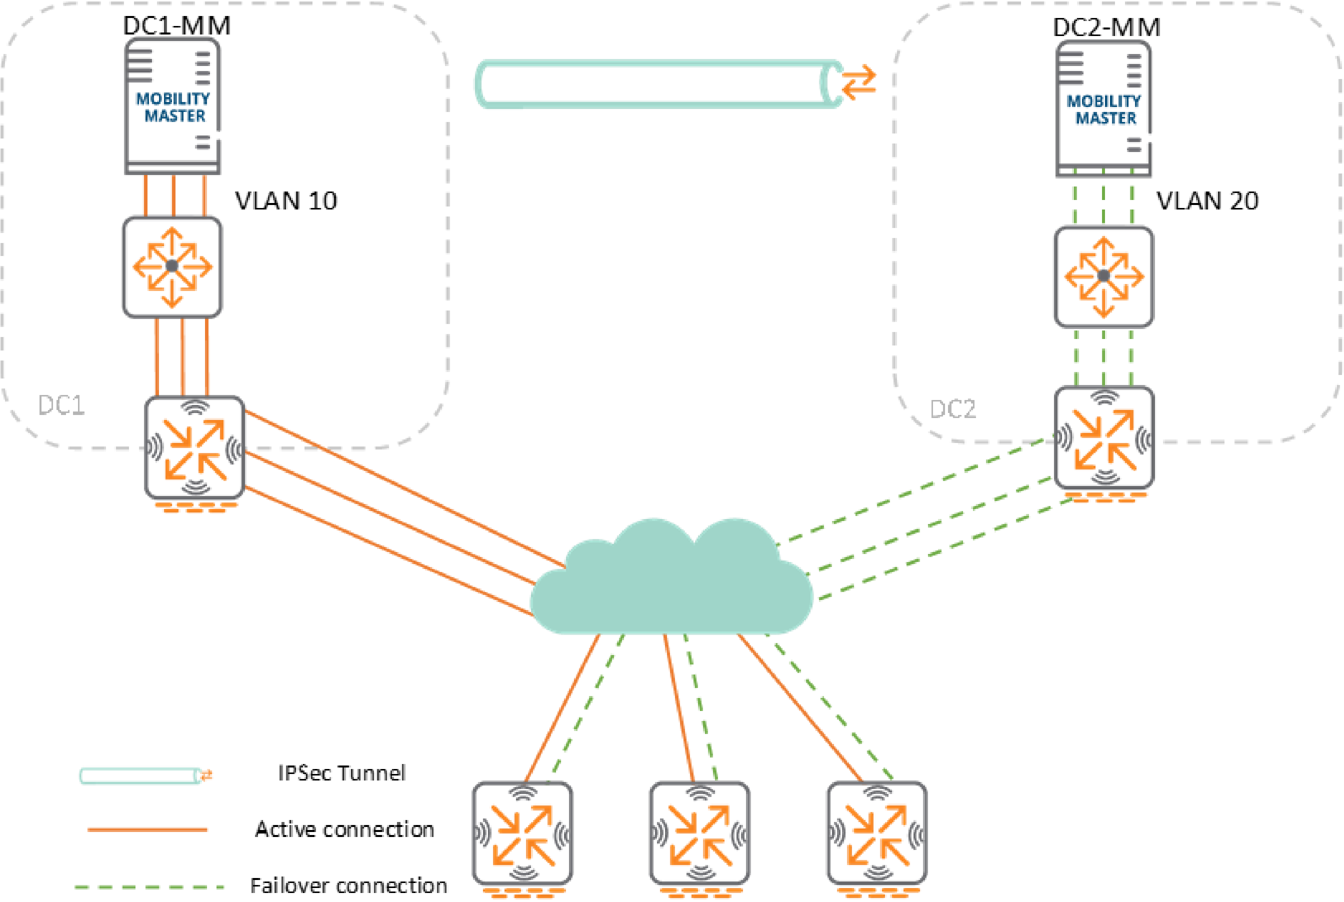

Figure 57 *Layer 3 Redundancy between Two MMs*

In the topology depicted above DC1 is acting as the primary with DC1-MM while DC2 is the secondary DC with DC2-MM. Layer 3 redundancy has been configured between the primary and secondary DCs. Since MCs are able to detect and initiate the failover, they also need to be configured with the IP address of the secondary MM.

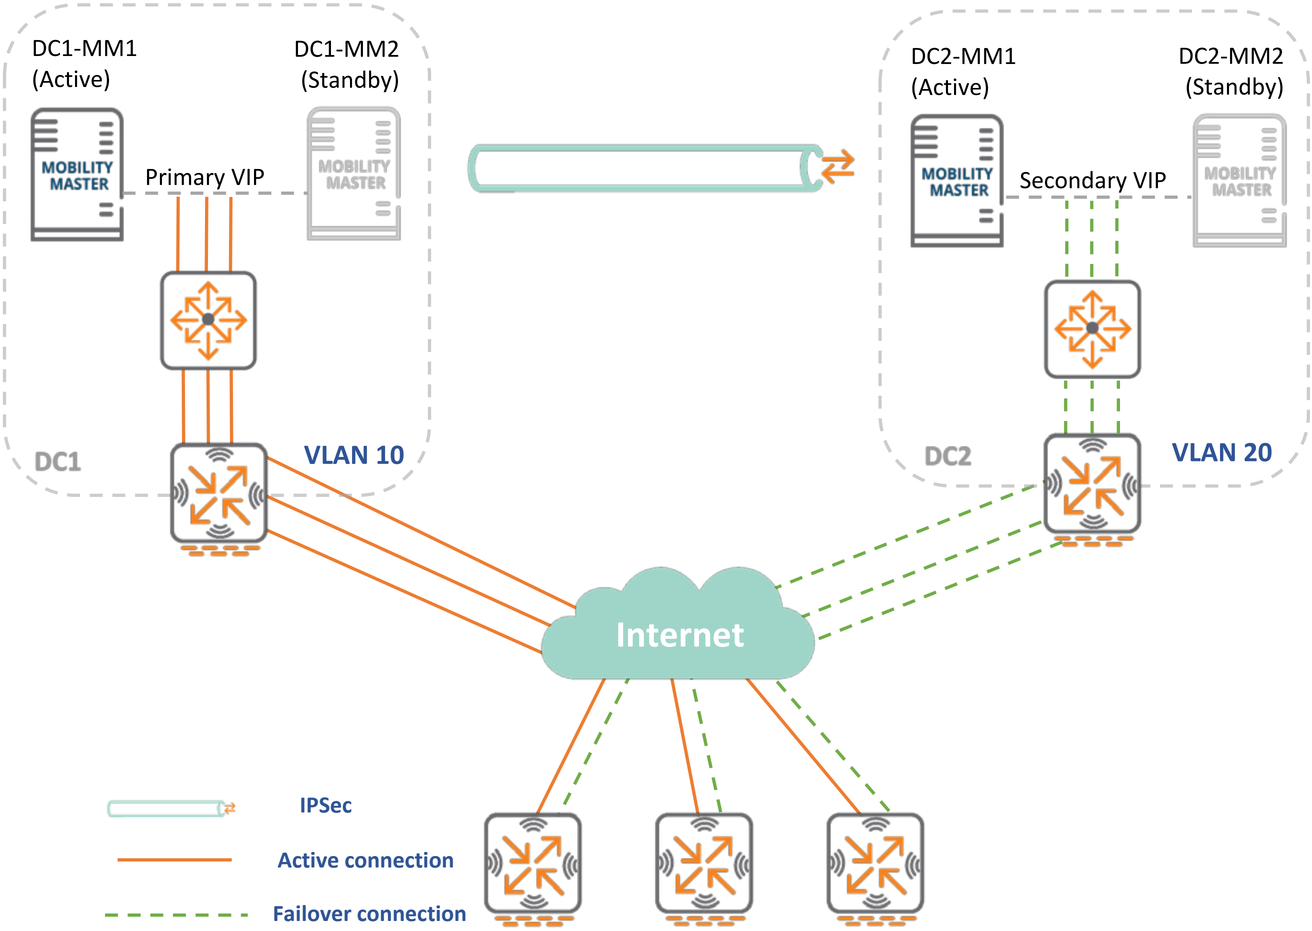

Figure 58 *Redundancy between Two MM Pairs*

In Figure 58 DC1 is the primary DC with DC1-MM1 serving as the active MM and DC1-MM2 serving as the standby. Layer 2 redundancy is configured between MMs within each DC. If DC1-MM1 fails, then DC1-MM2 takes over the VIP as the new active MM. The MCs will terminate on the VIP for the VRRP instance between DC1-MM1 and DC1-MM2.

DC2 is the secondary DC with DC2-MM1 as the active MM and DC2-MM2 as the standby. Layer 2 redundancy is configured between them in the same manner as in DC1. If DC2-MM1 fails, then DC2-MM2 takes assumes ownership of the VIP for their shared VRRP instance and becomes the new active MM in DC2.

From the perspective of the MCs, the VIP in DC1 is the primary MM IP and the VIP in DC2 is the secondary MM IP address. As with the first layer 3 topology, the MCs will need to be configured with the virtual IP addresses for both datacenters.

Synchronization

The entire configuration hierarchy, databases, and associated configurations on the active MM in DC1 are synchronized with the active MM in DC2. If Layer 2 redundancy is configured within the DC then the active MM in DC2 synchronizes the configuration hierarchy, databases, and associated configurations with its Standby as well.

NOTE

The configuration sync between MMs does not include configurations under the MM and Device nodes.

Failover

MCs actively monitor connectivity to both primary and secondary DCs using IP health checks (pings), however active connections are established only to the primary DC. The connection to the secondary DC is built only if connectivity to all available MMs in the primary DC is lost.

In the event that connectivity to the primary DC is lost, the MCs will wait for a 15-minute interval before failing over to the secondary DC. The 15-minute window gives the primary DC a chance to recover and safeguards the MMs from unwanted failover situations such as if they have been rebooted. In such instances, the MM would experience downtime while there was no intention to fail the MCs over to the secondary DC.

If connectivity to the primary MM persists after 15 minutes, the MCs will failover to the secondary MM which will only accept them if it detects its own IPsec tunnel with the primary MM is down. If the primary MM comes back up at a later time, the MCs will immediately tear the connection to the secondary MM and reconnect with the primary.

By default, the secondary MM remains in the secondary role even after failover. During this time, no configuration changes can be performed on the MCs. This is to prevent a *split-brain* state where the primary MM comes back up after the secondary MM has pushed configuration changes to the MCs resulting in different configurations in each DC. Even in the secondary role, all the MM services will continue to run as usual. The only impact of a failover event is to the configuration capabilities of the MM; operational capabilities will remain completely unaffected.

In the event that the primary MM cannot recover then the secondary MM can be converted to the primary role which will allow it to push configuration changes to the MCs. The process of promoting the secondary MM to a primary role may only be performed manually. This forces a human to verify that they are absolutely certain that the primary DC will remain down and the change is desired as well as necessary. If the secondary MM is promoted to the role of primary MM, the failed primary must be reconfigured as the new secondary so it inherits the new primary MM’s configuration and databases. Doing so prevents the failed primary from assuming its old role when it comes back and in the process creating a conflict due to having two MMs in the primary role.

Clustering

Objectives

Clustering is one of the key features introduced in ArubaOS 8 and was specifically designed to capitalize on the MM architecture and deliver maximum value for mission-critical networks. Clustering was developed to achieve the following objectives:

- **Seamless Campus Roaming** - Clients in a single large layer 2 domain will associate and stay anchored to a single MC as they roam. Users will maintain the same subnet and IP address regardless even if they roam across APs which are anchored to different controllers. This enables mobility without compromising or sacrificing performance
- **Stateful Client Failover** - User traffic will remain uninterrupted and high value sessions will be preserved in the event of a cluster member failure. Clients will not be required to re-authenticate and there will be no adverse impact to performance. The impact to performance will be mitigated to such an extent that users will not notice any degradation in their performance and they will have no knowledge that a failure has even occurred regardless of the applications they are currently utilizing
- **Access Point and Client Load Balancing** - APs and users are automatically load balanced across controllers that are members of the cluster. This process ensures an even distribution in order to deliver and maintain optimal network performance as well as to preserve capacity across all cluster members for new client associations
- **Live Upgrade** - Aruba allows customers to perform in-service cluster upgrades which allow improvements to be implemented without affecting performance while the network remains fully operational. The Live Upgrade feature allows upgrades to be completely automated. This is a key feature for customers with mission-critical networks that must remain operational 24/7

NOTE

Live upgrades can only be performed on MCs in a cluster and the APs attached to them.

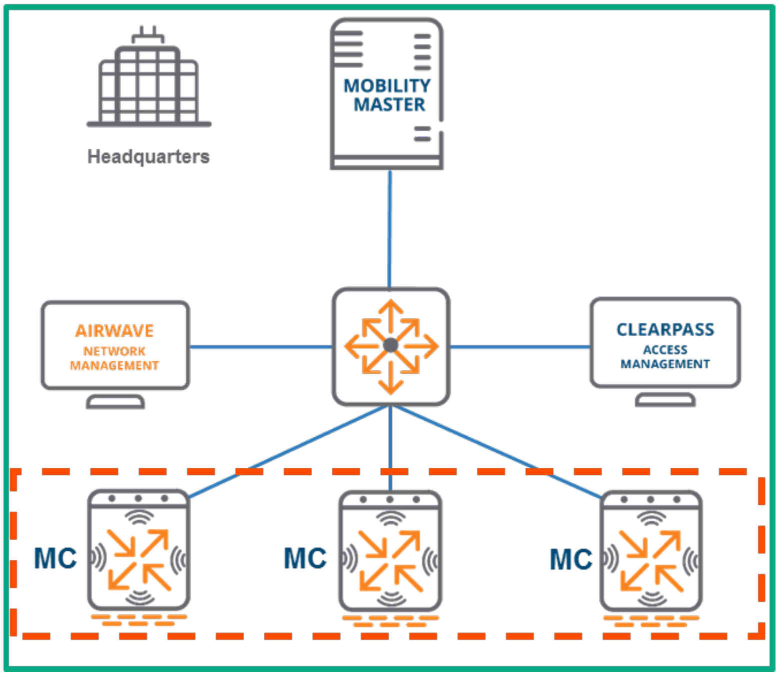

Figure 59 *Typical MC Cluster Architecture*

Highlights and Considerations

Clustering is a key feature of ArubaOS 8 however it cannot be enabled for all devices. Only MCs under management of an MM can form a cluster. MMs themselves however cannot become a member of a cluster with another MM nor with MCs. MMs strictly function as management devices for MCs in a cluster. While the redundancy options for an MM environment include both clustering as well as High Availability (HA) with AP fast failover, these are mutually exclusive features. One or the other must be chosen as they cannot both be concurrently operational.

CAUTION

All MCs in a cluster need to run the same software version so that APs that failover to a new controller will not inadvertently upgrade to a new version.

It should be noted that clustering is not supported by Standalone controllers. If Standalone controllers must be used then their primary redundancy mechanism is HA. Clustering and all of its constituent features are supported for both Campus Access Points, Remote Access Points, and meshed Access Points without requiring any additional licenses. The controller models which support clustering include the 72xx family, the 70xx family, and VMCs.

The cluster capacity for each product line is detailed in the table below:

Product Family	Devices per Cluster
72xx	12
70xx	4
Virtual	4

Table 9 *Cluster Capacity by Product Family*

While it is technically possible to combine 72xx and 70xx devices in the same cluster doing so is strongly discouraged as a long term deployment option. Such a scenario is acceptable as a temporary migration strategy however as a best practice cluster devices should always be homogeneous. If different controller models are clustered together then all controller scalability limits will be downgraded to the capabilities of the lowest controller model. E.g., if a cluster was created with two 7240 controllers and one 7210 controller then the cluster's scalability capacity will be limited to that of three 7210 controllers.

NOTE

Virtual and hardware controllers cannot be combined in a cluster under any circumstances.

If RAPs are being terminated on any of the MCs in the cluster the number of devices allowed in that cluster is limited to 4. The figure below depicts the dashboard view for a cluster:

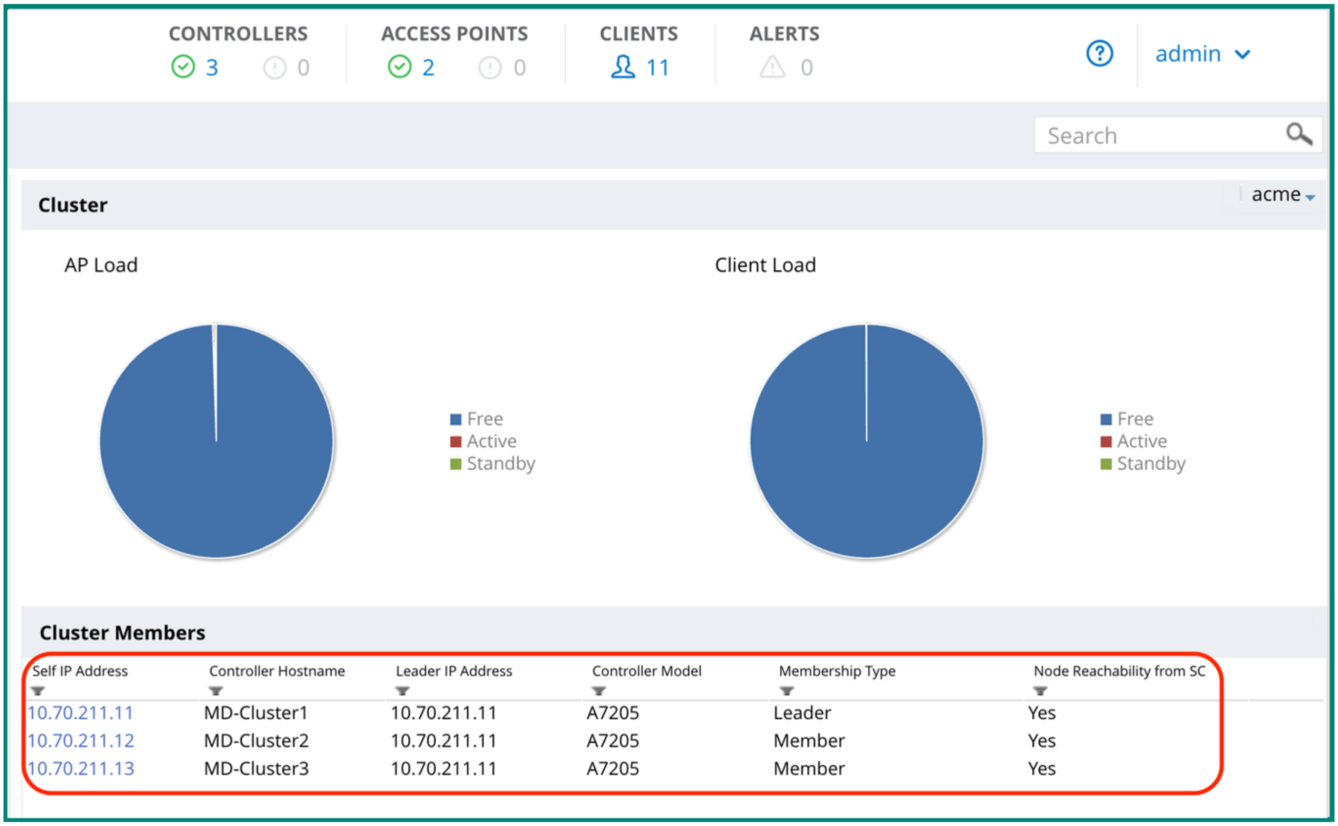

Self IP Address	Controller Hostname	Leader IP Address	Controller Model	Membership Type	Node Reachability from SC
10.70.211.11	MD-Cluster1	10.70.211.11	A7205	Leader	Yes
10.70.211.12	MD-Cluster2	10.70.211.11	A7205	Member	Yes
10.70.211.13	MD-Cluster3	10.70.211.11	A7205	Member	Yes

Figure 60 *MM Cluster Dashboard*

The view above can be accessed through the GUI by navigating to the **Cluster** tab of the main dashboard. Key statistics about a cluster can be seen under this tab including the number of controllers, APs and clients in the cluster under management by the MM as well as the current AP and client loads of the cluster members. The **Cluster Members** section at the bottom displays key statistics pertaining to the MC which are members of the cluster including their IP address, model, and which device is acting as the current cluster leader.

Cluster Formation

Handshake Process

The first step of cluster formation involves a handshake process where messages are exchanged between all potential cluster members. The handshake process occurs using hello messages which are exchanged to verify layer 3 reachability between all cluster members. Information relevant to clustering is exchanged through these messages including build, cluster name, and information about the MC sending the message. After all members have exchanged these messages they will establish layer 3 IPsec connectivity with each other in a fully-meshed configuration. The figure below depicts cluster members engaging in the hello message exchange process as part of the handshake prior to cluster formation:

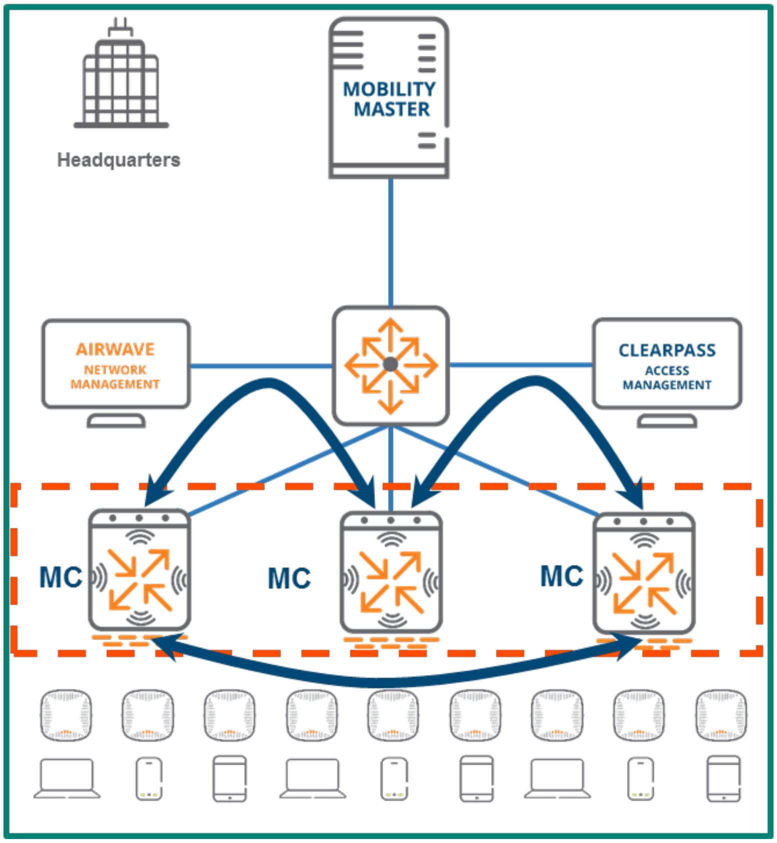

Figure 61 *Handshake Process/Hello Messages*

VLAN Probing

After the cluster has entered an L3-Connected state and the cluster members have formed IPsec connections in a full mush, each member will unicast layer 2 probes on each of its VLANs to each of the other cluster members. If the probing process is successful then the cluster will transition from an L3-Connected to an L2-Connected state meaning that all cluster members are sharing the same VLANs.

NOTE

Clusters can be formed over a layer 2 or a layer 3 network. Aruba strongly recommends configuring clusters with layer 2 connectivity to enable VRRP. Doing so provides CoA support for the cluster and facilitates controller discovery.

When discussing Aruba MC clusters, "*L2 Connected*" and "*L3 Connected*" are specific terms which refer to the state of a cluster and indicate whether or not all cluster members share the same VLANs. They are not abbreviations of the traditional networking terms *layer 2* and *layer 3*. The table below provides additional clarification on the topic:

Term	Definition
Layer 2 connectivity	MCs are connected and share the same management VLAN
Layer 3 connectivity	MCs can reach each other but do not share the same management VLAN
L2 Connected	A cluster state where all members share all of the same user VLANs
L3 Connected	A cluster state where members do not share all of the same user VLANs

Table 10 *Cluster Connectivity Distinction*

NOTE

MCs in a cluster can be configured in an L2-Connected state even if they do not share all of the same VLANs. This is done by entering a command forcing MCs to exclude certain VLANs from the probing process.

Leader Election

In every cluster one MC will be selected as the cluster leader. The cluster leader has multiple responsibilities including:

- Determining which clients are mapped to each cluster member
- Dynamically load balancing clients to ensure an even distribution of resources if a cluster member becomes overburdened. When a new member is added to the cluster the cluster leader will evenly redistribute the load across all members. This is a completely seamless process and users will not experience any performance degradation
- Identification of the standby MC for each AP and client to ensure stateful failover in the event that a controller goes down

The cluster election takes place after the initial handshake as a parallel thread to VLAN probing and the heartbeat process. The cluster leader is elected as a result of each cluster member exchanging messages which include their configured priority, platform value, and the MAC address.

Heartbeats

After the initial handshake, all cluster members will commence sending out heartbeat messages to one another at regular intervals in parallel to the leader election and VLAN probing threads. These heartbeat messages serve as the primary detection mechanism for cluster member failures. Heartbeats are integral to the process the cluster leader uses to determine the role of each cluster member.

Connectivity and Verification

A cluster's connectivity status can be viewed in the GUI of the MM by navigating to **Dashboard>Cluster>Cluster Members** and then selecting the IP address of any cluster member. This view is displayed in the figure below:

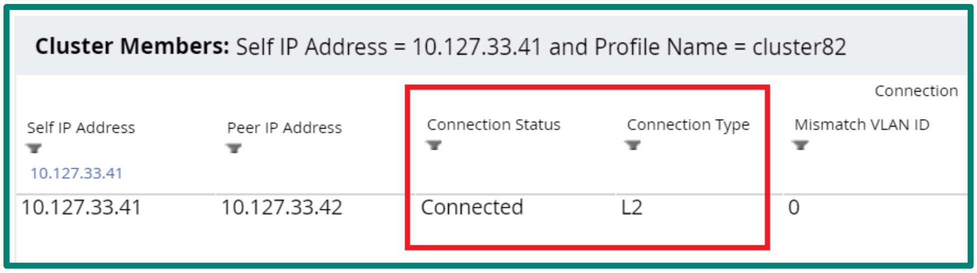

Cluster Members: Self IP Address = 10.127.33.41 and Profile Name = cluster82

Self IP Address 10.127.33.41	Peer IP Address	Connection Status	Connection Type	Connection Mismatch VLAN ID
10.127.33.41	10.127.33.42	Connected	L2	0

Figure 62 *Viewing Cluster Connectivity Status*

Cluster Roles

An MC can have any combination of the following four roles in a cluster aside from being the cluster leader:

1. AP Anchor Controller (AAC)
2. User Anchor Controller (UAC)
3. Standby AAC (S-AAC)
4. Standby UAC (S-UAC)

AP Anchor Controller

AAC Assignment

Anchoring is a concept that was introduced in ArubaOS 8 as part of the clustering feature set. Anchoring and clustering are designed to achieve the following objectives:

- Enhance user mobility through seamless campus roaming
- Ensure an even distribution of resources across the cluster to maintain the highest achievable performance level
- Enable redundancy scenarios creating fault tolerance for the cluster and minimizing the impact of an MC failure

The *AP Anchor Controller (AAC)* can be thought of as the LMS for any AP that is anchored to it. Each AP receives the IP address of the LMS and once they have been terminated they will remain anchored until the cluster leader determines that they should be moved to a different cluster member. An AP is anchored to its AAC in a three step process:

1. The AP establishes active tunnels with its AAC
2. The cluster leader dynamically assigns a standby AP Anchor Controller (S-AAC) for the AP from one of the other cluster members
3. Once designated the AP established standby tunnels to the S-AAC

The AAC and S-AAC assignment process works similarly to how HA is configured however rather than having to be manually configured the process is completely dynamic. Once the AAC is designated for an AP the subsequent steps occur automatically. A visual representation of AAC assignment is displayed in the figure below:

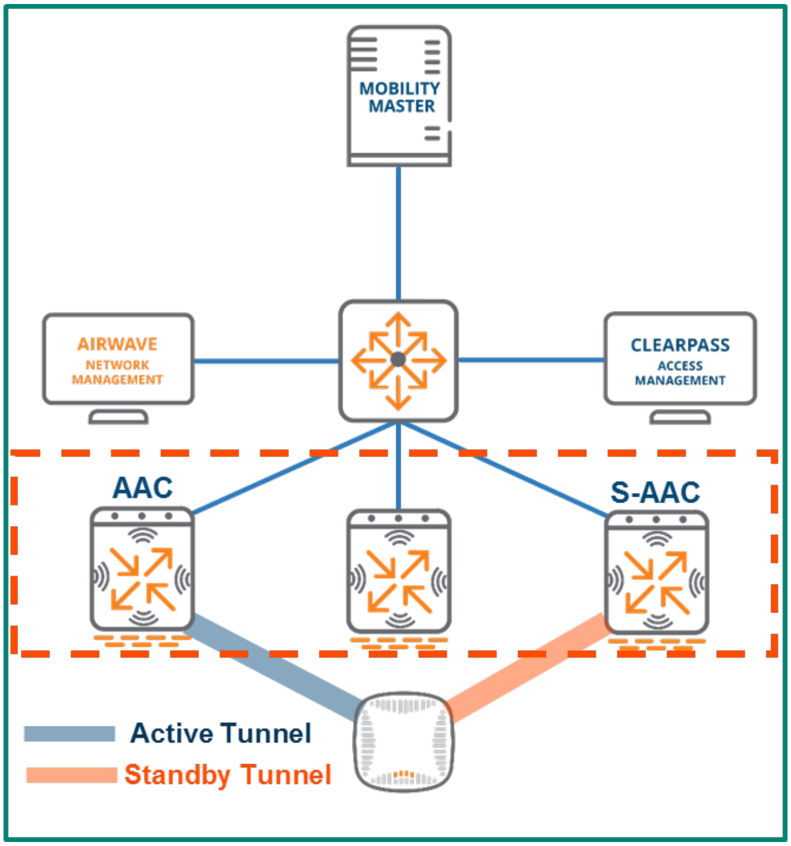

Figure 63 *AAC Assignment*

The AAC and S-AAC for an AP can be identified in the GUI of the MM by navigating to **Dashboard>Access Points**:

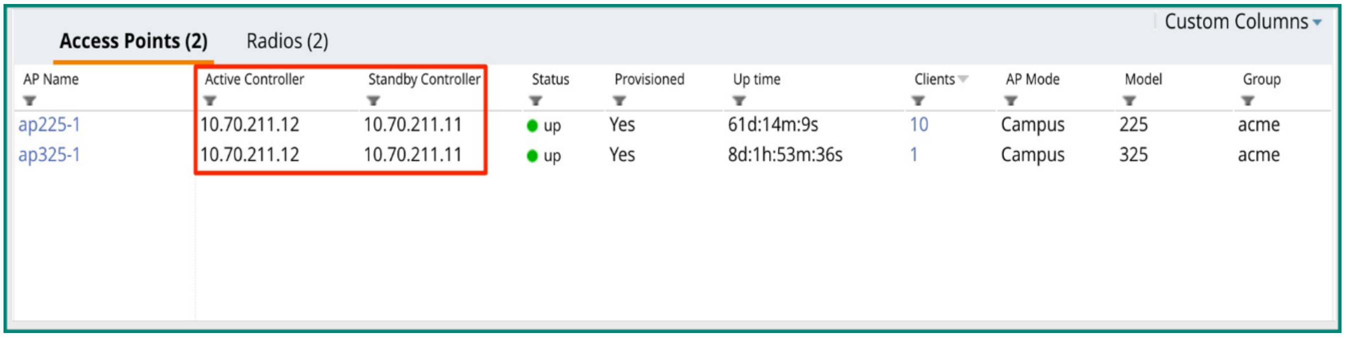

AP Name	Active Controller	Standby Controller	Status	Provisioned	Up time	Clients	AP Mode	Model	Group
ap225-1	10.70.211.12	10.70.211.11	● up	Yes	61d:14m:9s	10	Campus	225	acme
ap325-1	10.70.211.12	10.70.211.11	● up	Yes	8d:1h:53m:36s	1	Campus	325	acme

Figure 64 *AAC and S-AAC Status*

While the view above indicates that the S-AAC for both APs is the same device (10.70.211.11) it should be noted that the S-AAC is assigned by the cluster leader and not all APs terminated on an AAC will have the same S-AAC. It could just as easily be a different cluster member depending on the determination made by the cluster leader based on the conditions in the cluster environment at the time of assignment.

AAC Failover

Every AP is assigned an AAC and S-AAC from the members of the cluster. Aps will create tunnels to both MCs in advance in order to facilitate the failover process. The redundant tunnel to the S-AAC ensures that APs will transition seamlessly in the event of a cluster member failure. With clustering configured failover events will have a negligible impact to network performance and users will not have any awareness that a failure has occurred. The failover process occurs as outlined by the steps below:

1. An AAC fails. The failure is immediately detected by the S-AAC due to heartbeats
2. Upon detection of the failure the S-AAC will instruct the AP to failover
3. The AP tears down its tunnel to the AAC that has failed and fails over to the S-AAC
4. The existing AP standby tunnel becomes active with the S-AAC which assumes the role of AAC for that AP
5. A new S-AAC is dynamically assigned for the AP from the remaining cluster members
6. The AP establishes standby tunnels to the new S-AAC

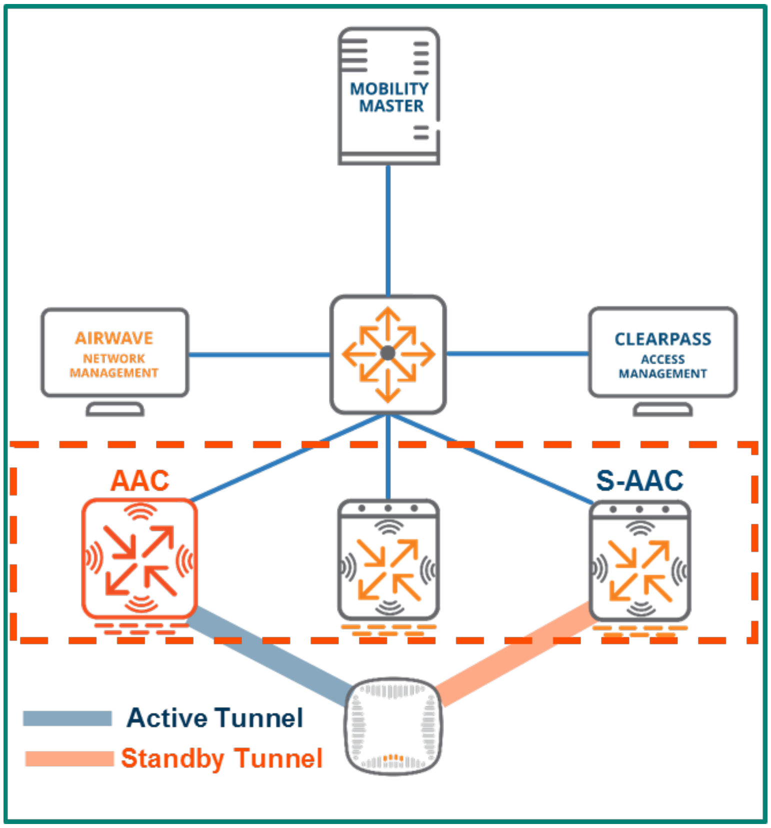

Figure 65 *AAC Fails*

The figure above shows that the AAC the AP was previously connected to has failed. At this point the S-AAC will instruct the AP to failover and tear down its active tunnel to the failed AAC.

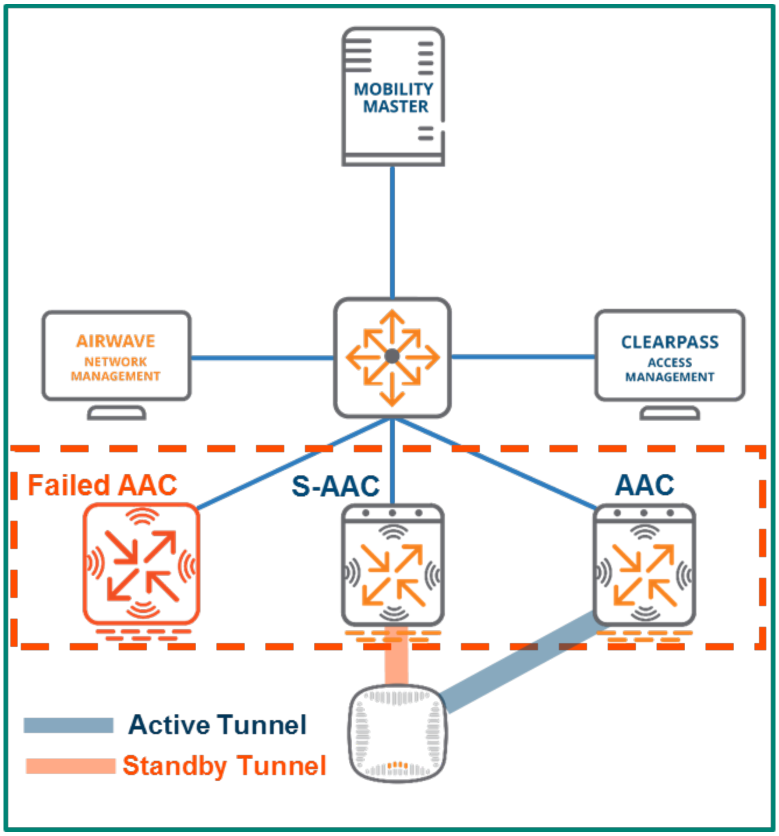

Figure 66 *AP Fails Over and New S-AAC Assigned*

The old S-AAC has now assumed the role of AAC. The old standby tunnel between the AP and the S-AAC has now become an active tunnel. Another member of the cluster was dynamically selected by the cluster leader to serve as the new S-AAC and the AP has built a standby tunnel accordingly. The failover process is now complete and all steps were completely undetectable to users.

User Anchor Controller

The concept of anchoring users to a controller using a *User Anchor Controller (UAC)* is new in ArubaOS 8 and was primarily developed to enhance the user roaming experience. When users associate to an AP, they will use the existing tunnel to their UAC if one already exists. If the AP doesn't have tunnel to their UAC established then a dynamic tunnel is created. When the client roams to a new AP, the AP they are roaming away from tears down its dynamic tunnel. User traffic is always tunneled back to their UAC regardless of which AP the client associates to as the user roams, even if that AP has a different AAC.

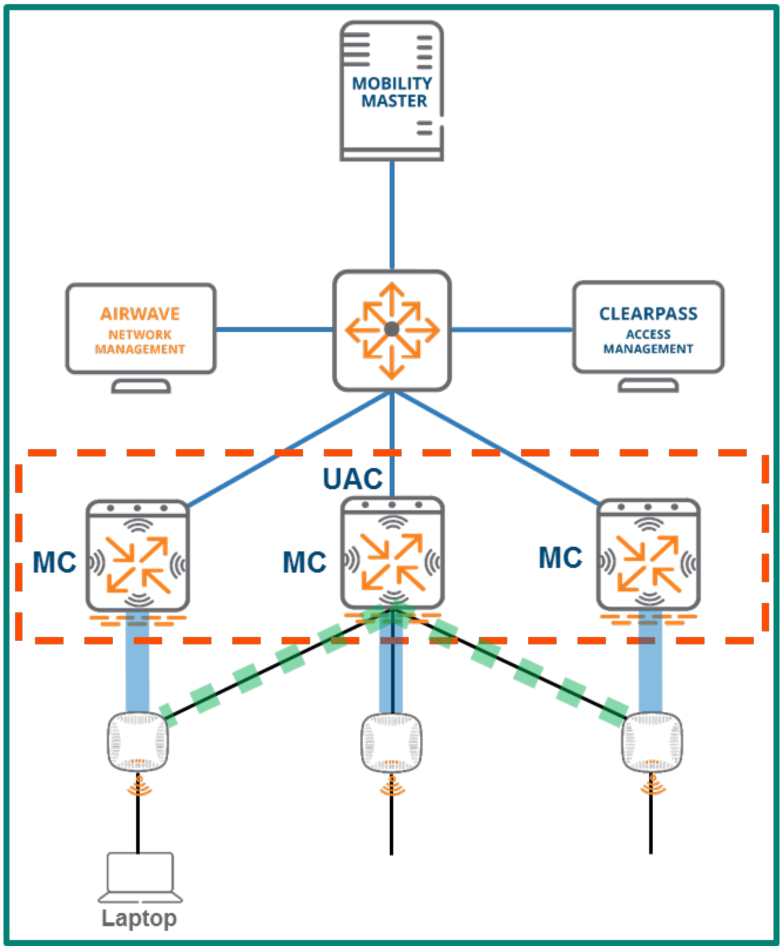

Figure 67 *Dynamic Tunnel to the Client's UAC*

In order to remain anchored, a user must first be mapped to a UAC through a hashing algorithm at the AP level. The MAC address of the client is examined and the hashing algorithm creates an index which is then compared to a mapping table. The same mapping table is pushed to all APs by the cluster leader to ensure UAC mapping consistency across the cluster. In addition, the cluster leader will dynamically select a standby UAC (S-UAC) on a per-user basis for redundancy purposes. An example of the hashing algorithm and UAC assignment process is displayed in the figure below:

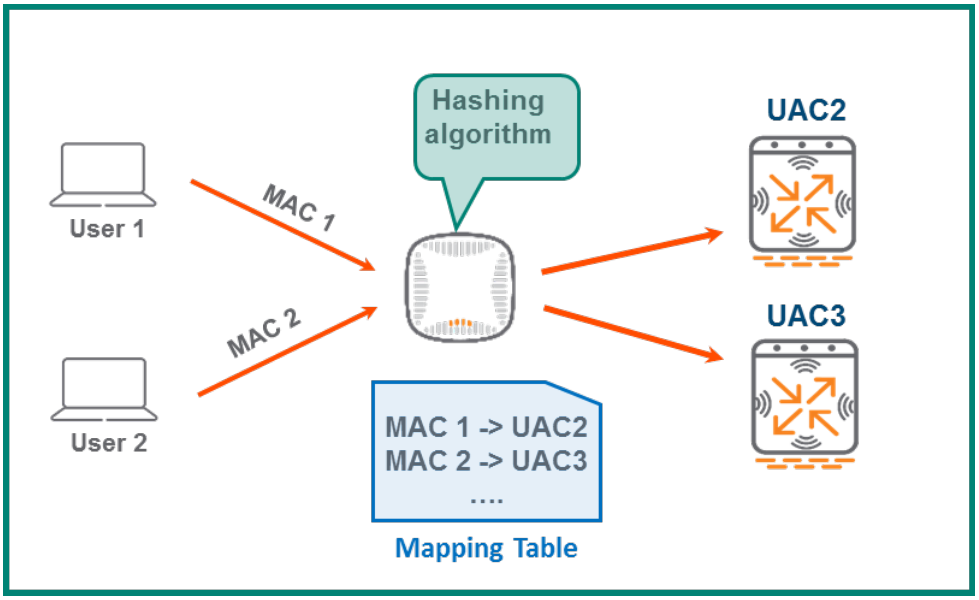

Figure 68 *UAC Assignment Process*

The UAC and S-UAC assignments for all associated clients can be identified in the GUI of the MM by navigating to **Dashboard>Clients**:

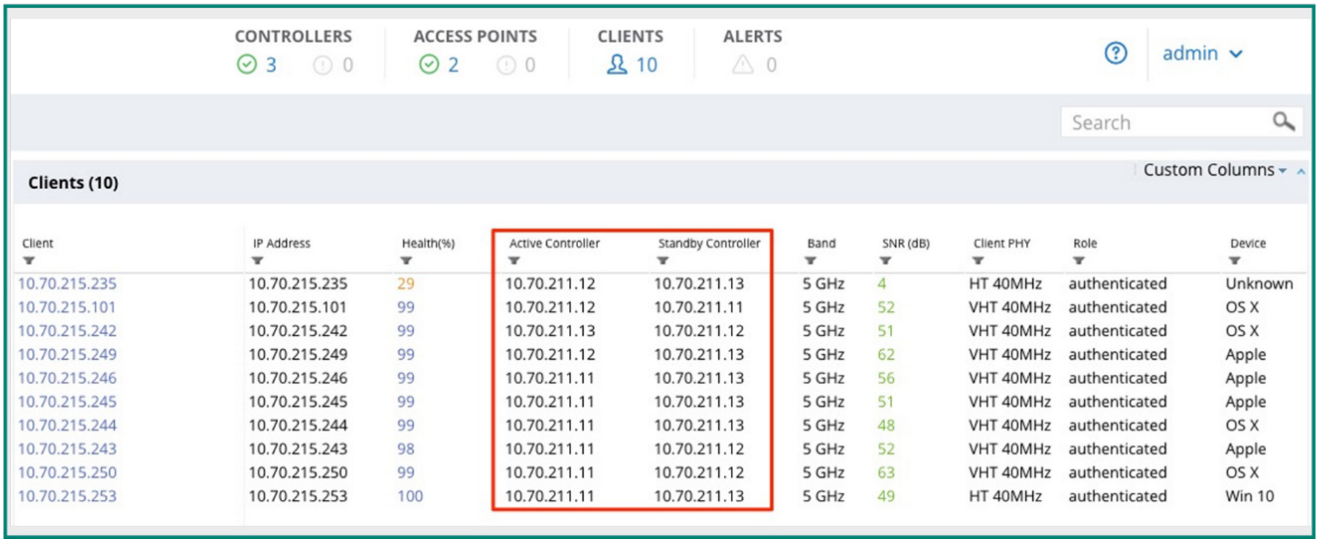

Clients (10)

Client	IP Address	Health(%)	Active Controller	Standby Controller	Band	SNR (dB)	Client PHY	Role	Device
10.70.215.235	10.70.215.235	29	10.70.211.12	10.70.211.13	5 GHz	4	HT 40MHz	authenticated	Unknown
10.70.215.101	10.70.215.101	99	10.70.211.12	10.70.211.11	5 GHz	52	VHT 40MHz	authenticated	OS X
10.70.215.242	10.70.215.242	99	10.70.211.13	10.70.211.12	5 GHz	51	VHT 40MHz	authenticated	OS X
10.70.215.249	10.70.215.249	99	10.70.211.12	10.70.211.13	5 GHz	62	VHT 40MHz	authenticated	Apple
10.70.215.246	10.70.215.246	99	10.70.211.11	10.70.211.13	5 GHz	56	VHT 40MHz	authenticated	Apple
10.70.215.245	10.70.215.245	99	10.70.211.11	10.70.211.13	5 GHz	51	VHT 40MHz	authenticated	Apple
10.70.215.244	10.70.215.244	99	10.70.211.11	10.70.211.13	5 GHz	48	VHT 40MHz	authenticated	OS X
10.70.215.243	10.70.215.243	98	10.70.211.11	10.70.211.12	5 GHz	52	VHT 40MHz	authenticated	Apple
10.70.215.250	10.70.215.250	99	10.70.211.11	10.70.211.12	5 GHz	63	VHT 40MHz	authenticated	OS X
10.70.215.253	10.70.215.253	100	10.70.211.11	10.70.211.13	5 GHz	49	HT 40MHz	authenticated	Win 10

Figure 69 *Client UAC and S-UAC Assignments*

NOTE

The Active Controller and Standby Controller columns are not included in the standard view of the Clients page in the GUI. They can be displayed by adding a customization to the page view.

Cluster Features

Seamless Roaming

The advantage to introducing the concept of the UAC is that it significantly enhances the experience for users roaming within a cluster. Once a user associates to an AP it hashes the client's MAC address and assigns it a UAC. From this point on the traffic from that user will always be tunneled to their UAC. This remains true regardless of which AP the users associate to as they roam, even if that AP happens to be terminated on a different controller. Any AP the user roams to will automatically forward the traffic to the UAC the user was assigned upon association. If an active or standby tunnel does not already exist between the AP where the user has roamed and the UAC then a dynamic tunnel will be created. A visual representation of the roaming process within a cluster is displayed in the figure below:

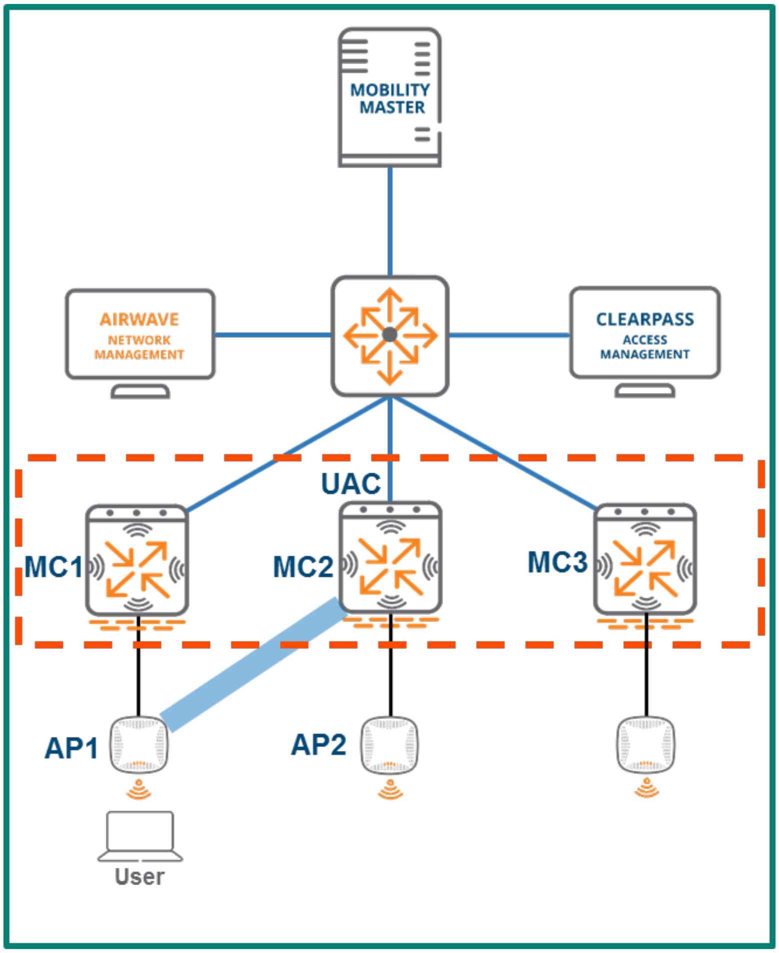

Figure 70 *Cluster Seamless Roaming*

In the figure above the user has associated to AP1 which is terminated on MC1, however the traffic is being tunneled to MC2. In this scenario, MC2 has been designated the UAC for the user so as the user roams over to AP2 or any other AP in the cluster the traffic will continue to be tunneled to MC2.

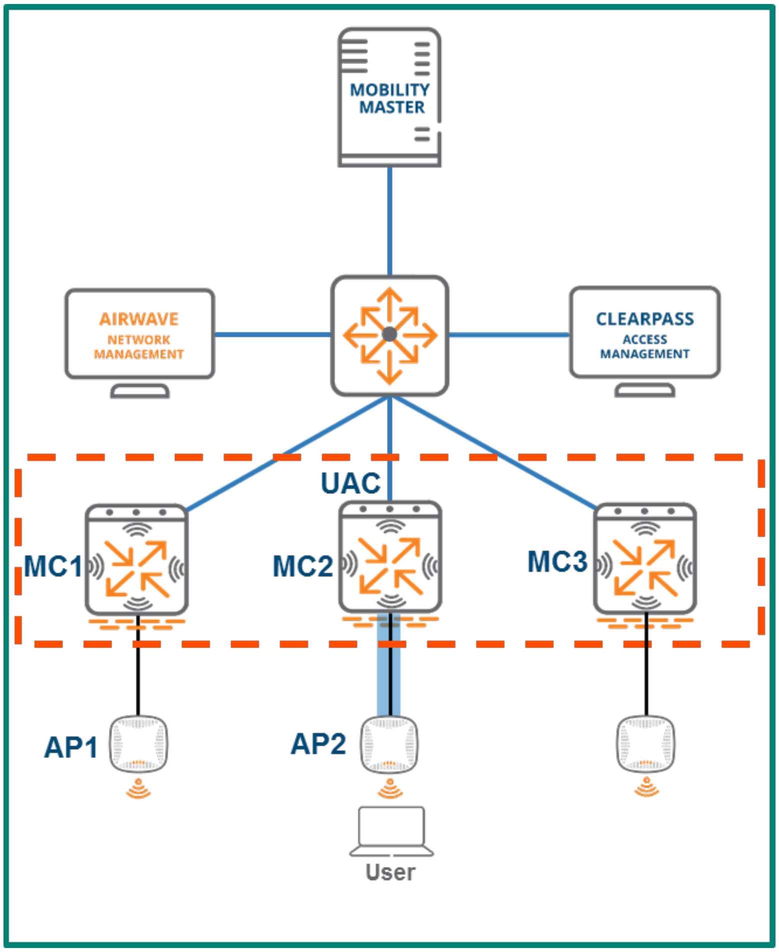

Figure 71 *Cluster Seamless Roaming continued*

Stateful Failover

Stateful failover is a critical aspect of cluster operations that safeguards users from any impacts associated with a controller failure event. There are two key conditions which must be met to enable stateful failover functionality for a cluster:

1. Redundancy mode must be enabled. It can be disabled however it is enabled by default
2. An L2-Connected state must exist between all cluster members

Provided that these two conditions have been met the client state will then be fully synchronized between the UAC and the S-UAC meaning that information such as the station table, the user table, layer 2 user state, layer 3 user state, key cache, and PMK cache will all be shared between the both devices. In addition, high value sessions such as FTP, telnet, SSH, and DPI-qualified sessions are also synced to the S-UAC. Synchronizing all of the client state and high value session information enables the S-UAC to seamlessly assume the role as the client's new UAC in the event that the client's current UAC fails. Establishing cluster redundancy in this manner guarantees stateful failover with no client deauthentication when they move from their UAC to their S-UAC.

Seamless cluster failover provides a substantial advantage over redundancy enabled with an HA configuration which would require a client to be deauthenticated in the event of a controller failure. The table below outlines the advantages of L2-Connected versus L3-Connected cluster states specifically as they pertain to redundancy, failover, and performance:

L2-Connected	L3-Connected
APs and clients are fully replicated	Only APs fully replicated
Users fully synced between nodes	Users not synced
High value sessions are synced	High value sessions not synced
Users failover with no de-auth	Users are de-authenticated upon failover
Fully redundant	Not fully redundant

Table 11 *L2-Connected vs. L3-Connected*

Client Load Balancing

Load balancing clients across MCs is another feature that helps maintain cluster performance. While the hashing algorithm applied to clients that associate to an AP for UAC assignment works well for its intended purpose, it can result in a disproportionate distribution of clients across cluster members. This can lead to inefficient usage of system resources. Load balancing enables the cluster leader to optimally distribute users across the cluster and ensure peak performance levels are maintained. The cluster leader load balances clients across the cluster by following a multi-step calculation process where it identifies the model of each controller of the cluster, counts the number of associated clients, and compares the client count against the maximum capacity for each device to calculate its load ratio.

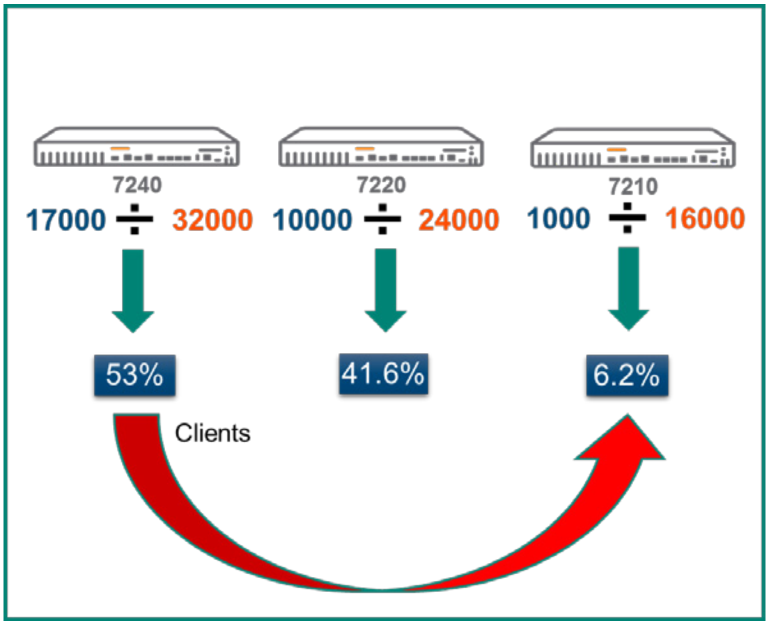

Figure 72 *Cluster Load Calculation Process*

The figure above demonstrates the load calculation process the cluster leader would perform for each member. In this scenario there are three cluster members which each have a different controller model: a 7240, a 7220, and a 7210. The blue numbers below each cluster member indicate the number of associated clients while the orange numbers indicate the maximum capacity for each particular model. The cluster leader compares the client count and checks that against the capacity for each device's model and produces a ratio for each expressed as a percentage of total capacity. The table below presents the specific triggers in place which will result in the cluster leader load balancing clients across the cluster:

Category	Threshold
Active Client Load	50%
Standby Client Load	75%
Unbalance Threshold	5%

Table 12 *Load Balancing Triggers*

The table demonstrates that a rebalancing event will be triggered when the active client load exceeds 50% or the standby client load exceeds 75% on any member of the cluster and while *unbalance threshold* also exceeds 5%. Unbalance threshold refers to the delta between the member of the cluster with the highest load percentage and the member of the cluster with the lowest load percentage and was put in place to ensure that regardless of how close the cluster members are to approaching a capacity trigger they will always remain with a roughly even distribution of clients.

AP Load Balancing

Just as the cluster leader will load balance clients to ensure an even distribution across cluster members it will also perform the same function for APs. Dynamic cluster AP load balancing is a configurable feature in ArubaOS 8. It should be enabled whenever there is a need to allow for ease of scalability when adding MCs to a cluster or a need to eliminate manual AP distribution.

The process works in the following manner:

1. An AP is dynamically assigned to an AAC upon connecting to the cluster
2. The cluster leader compares the AP count on each cluster member to the maximum capacity the controller platform
3. APs are dynamically load balanced across the cluster as needed based on platform capacity
4. The same thresholds for user load balancing are used for AP load balancing

AP Node List

The AP *node list* is a repository of IP addresses for each cluster member maintained by each AP. Upon connecting to the cluster the AP learns the IP addresses of all cluster members. These addresses are then stored as the AP's node list and saved as an environment variable. Upon booting up the AP will contact the first IP address in its node list. In the event that it does not receive a response it will try the next IP address in the list to ensure that APs are always able to find a reachable controller within the cluster.

Change of Authorization

Change of Authorization (CoA) is a feature which extends the capabilities of the Remote Authentication Dial-In User Service (RADIUS) protocol and is defined in RFC 5176. CoA request messages are usually sent by a RADIUS server to a network access server (NAS) device for dynamic modification of authorization attributes for an existing session. If the NAS device is able to successfully implement the requested authorization changes for the user session(s) then it will respond to the RADIUS server with a CoA acknowledgement also referred to as a CoA-ACK. Conversely, if the change is unsuccessful the NAS will respond with a CoA negative-acknowledgement or CoA-NAK.

In the context of an ArubaOS 8 cluster, unsolicited CoA requests for a user with an active session in progress are sent to that user's anchor controller. The UAC will then return an acknowledgement to the RADIUS server upon the successful implementation of the changes or a NAK in the event that the implementation was unsuccessful. However, a user's UAC may change in the course of normal cluster operations due to reasons such as an MC failure or user load-balancing events. Such a scenario would cause CoA requests to be dropped as the intended user would no longer be associated to the MC receiving the request from the RADIUS server. Aruba has implemented cluster redundancy features in order to prevent such a scenario from occurring.

Cluster CoA Support

The primary mechanism Aruba uses to provide CoA support for MC clusters in ArubaOS 8 is VRRP. In every cluster there are the same number of VRRP instances as there are nodes and each MC serves as the master of an instance. For example, a cluster with 5 MCs would have 5 instances of VRRP and 5 virtual IP addresses (VIPs). The master MC receives messages intended for the VIP of its instance while the remaining MCs in the cluster are backups for the all of other instances where they are not acting as the master. This configuration ensures that each cluster is protected by a fault-tolerant and fully redundant design.

NOTE

This section describes the process of Dynamic Authorization to RADIUS as described in RFC-5176 and how RADIUS communicates with Aruba controllers in a cluster. The Change of Authorization process was selected as a representation of that communication sequence.

ArubaOS reserves VRRP instance IDs in the 220-255 range. When the master of each instance sends RADIUS requests to the RADIUS server it injects the VIP of its instance into the message as the NAS-IP by default. This ensures that CoA requests from the RADIUS server will always be forwarded correctly regardless of which MC is the acting master for the instance. I.e. the RADIUS server sends CoA requests to the current master of the VRRP instance and not to an individual station. From the perspective of the server it is sending the request to the current holder of the VIP address of the instance. The figure below depicts sample architecture that will be used for the duration of the CoA section:

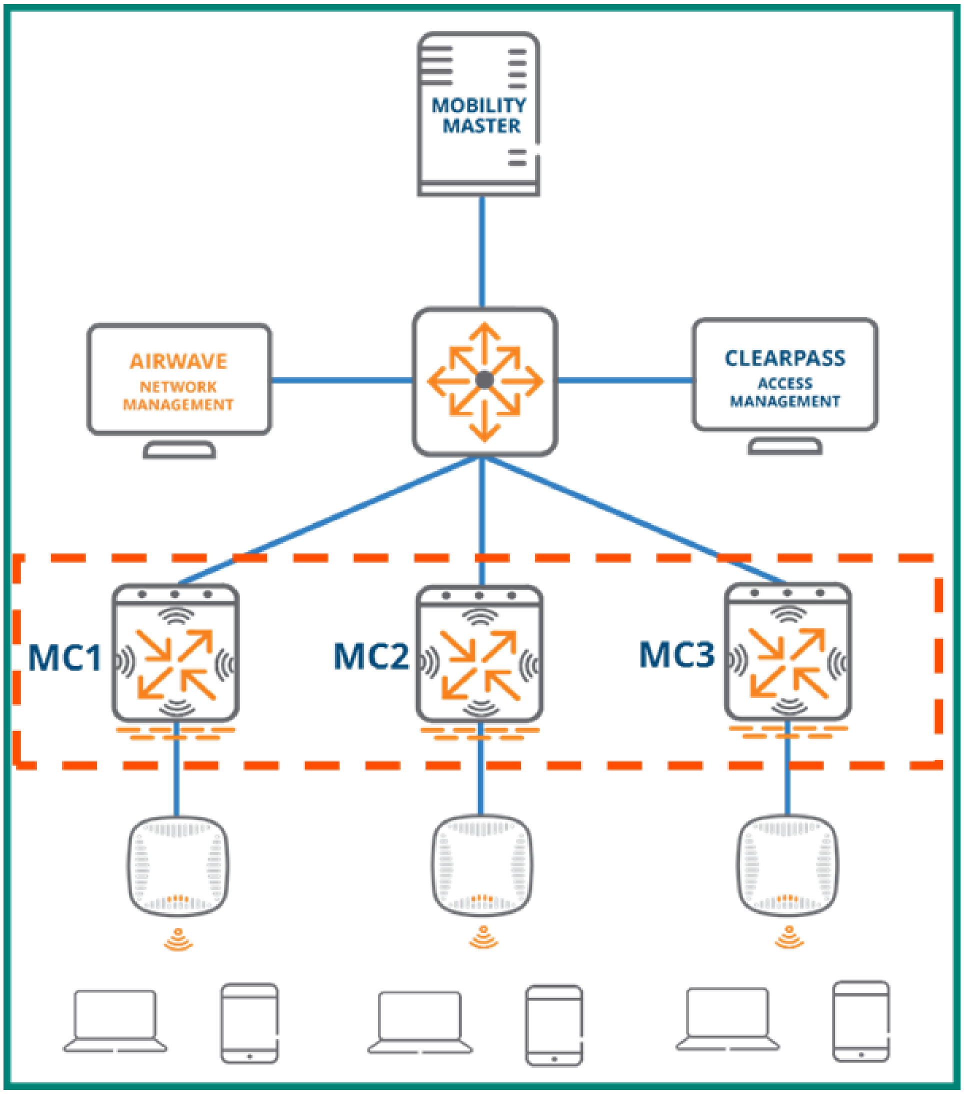

Figure 73 *Sample Architecture for CoA Demonstration*

This sample network consists of a three-node cluster with three instances of VRRP. The AOS-assigned VRRP ID range falls between 220 and 255 therefore the three instances in this cluster are assigned the VRRP IDs of 220, 221, and 222. The priorities for the MCs in each instance are dynamically assigned so that the master of the instance is assigned a priority of 255, the first backup is assigned a priority of 235, and the second backup is assigned a priority of 215. The table below outlines the priority assignments for each MC and each instance in the example network:

VRRP Instance	Virtual IP	MC1 Priority	MC2 Priority	MC3 Priority
ID 220	VIP1	**255**	235	215
ID 221	VIP2	215	**255**	235
ID 222	VIP3	235	215	**255**

Table 13 *MC Priorities and VIPs for Each VRRP Instance*

As demonstrated by the table, MC1 is the master of instance 220 with a priority of 255, MC2 is the first backup with a priority of 235, and MC3 is the second backup with a priority of 215. Similarly, MC2 is the master for instance 221 due to having the highest priority of 255, MC3 is the first backup with a priority of 235, and MC1 is the second backup with a priority of 215. Instance 222 follows the same pattern as instances 220 and 221.

CoA with MC Failure

The failure of a cluster node is an event that can adversely impact CoA operations if the network doesn't have the appropriate level of fault tolerance. If a user's anchor controller fails, the RADIUS server will push the CoA request to their UAC as usual with the assumption that it will enforce the change and respond with an ACK. However, if a redundancy mechanism such as VRRP hasn't been implemented then the request will go unanswered and will not result in a successful change. In such a scenario the users associated with the failed node will failover to their standby UAC as usual. However, the UAC will never receive the change request from the RADIUS server since the server has no awareness of cluster operations. VRRP instances must be implemented for each node to prevent such an occurrence and maintain CoA operations in the cluster.

In the figure below MC1 is the master of instance 220 with MC2 serving as the first backup and MC3 serving as the second backup. A client associated to MC1 has been fully authenticated using 802.1X with MC3 acting as the client's standby UAC. When corresponding with ClearPass, MC1 automatically inserts VIP for instance 220 as the NAS-IP. From the perspective of ClearPass it is sending CoA requests to the current master of instance 220.

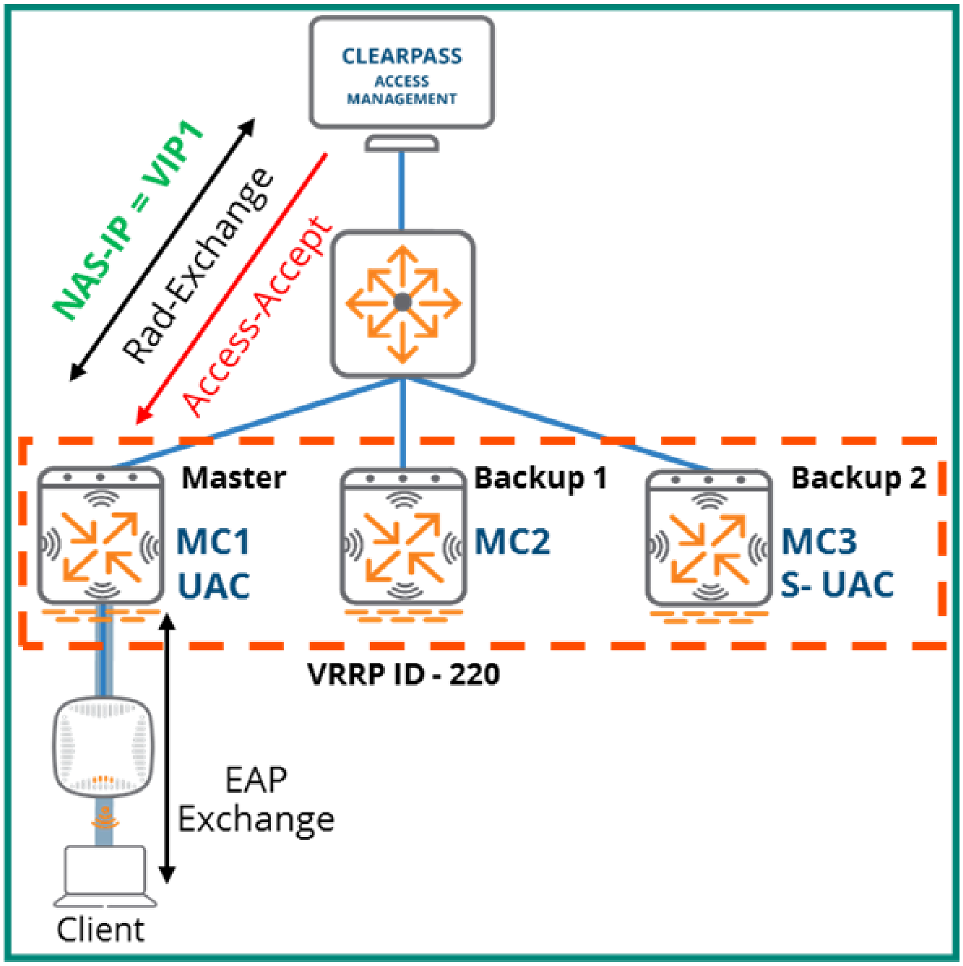

Figure 74 *User Authenticates Against ClearPass*

If MC1 fails while the client is in session the AP where the client is associated will failover to MC2. The client's session moves over to MC3 since it was the standby UAC. MC3 then assumes the role of UAC for the client. Since MC2 has a higher priority than MC3 in instance 220 it will assume the role of Master and take ownership of the VIP.

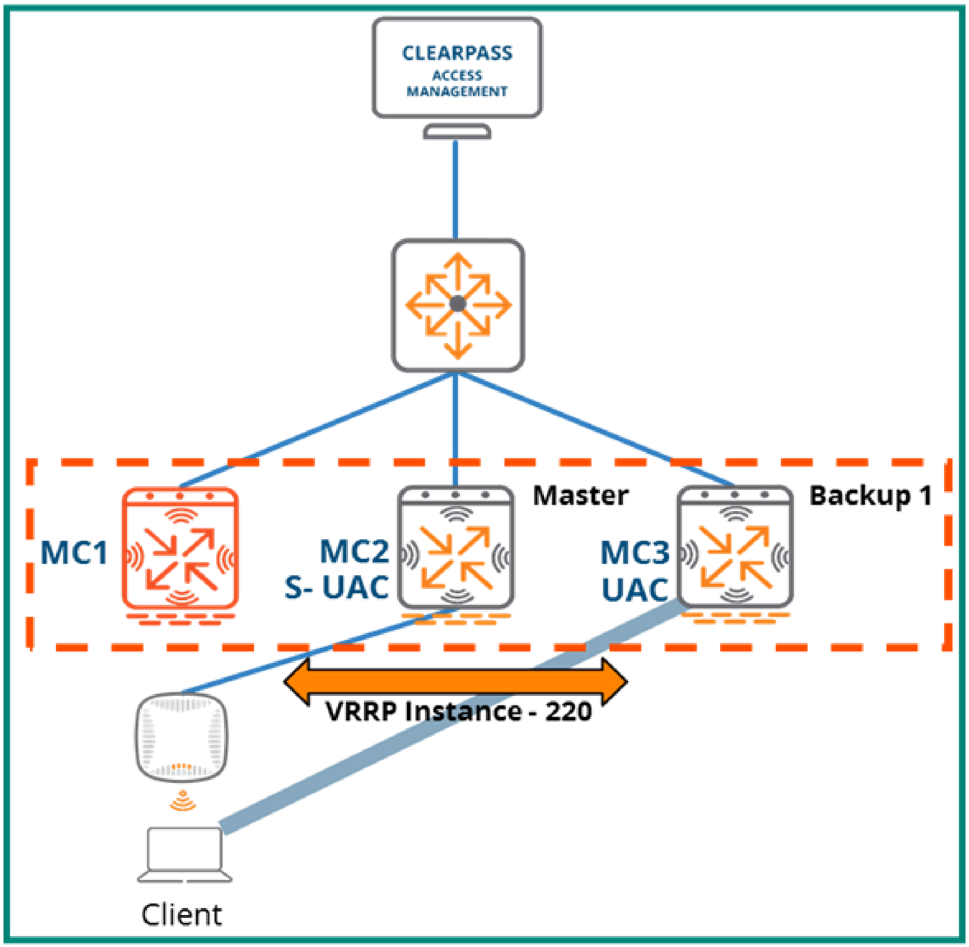

Figure 75 *MC1 Failure*

Any CoA requests sent by ClearPass for the client will be addressed to the VIP for instance 220. From the perspective of ClearPass, the VIP of instance 220 is the correct address for any CoA request intended for the client in the example. Since MC1 has failed, MC2 is now the Master of VRRP instance 200 and owns its virtual IP. When ClearPass sends a CoA request for the client, MC2 will receive it and then forward it to all nodes in the cluster. Since our cluster only has three nodes, in this case MC2 forwards the request to MC3.

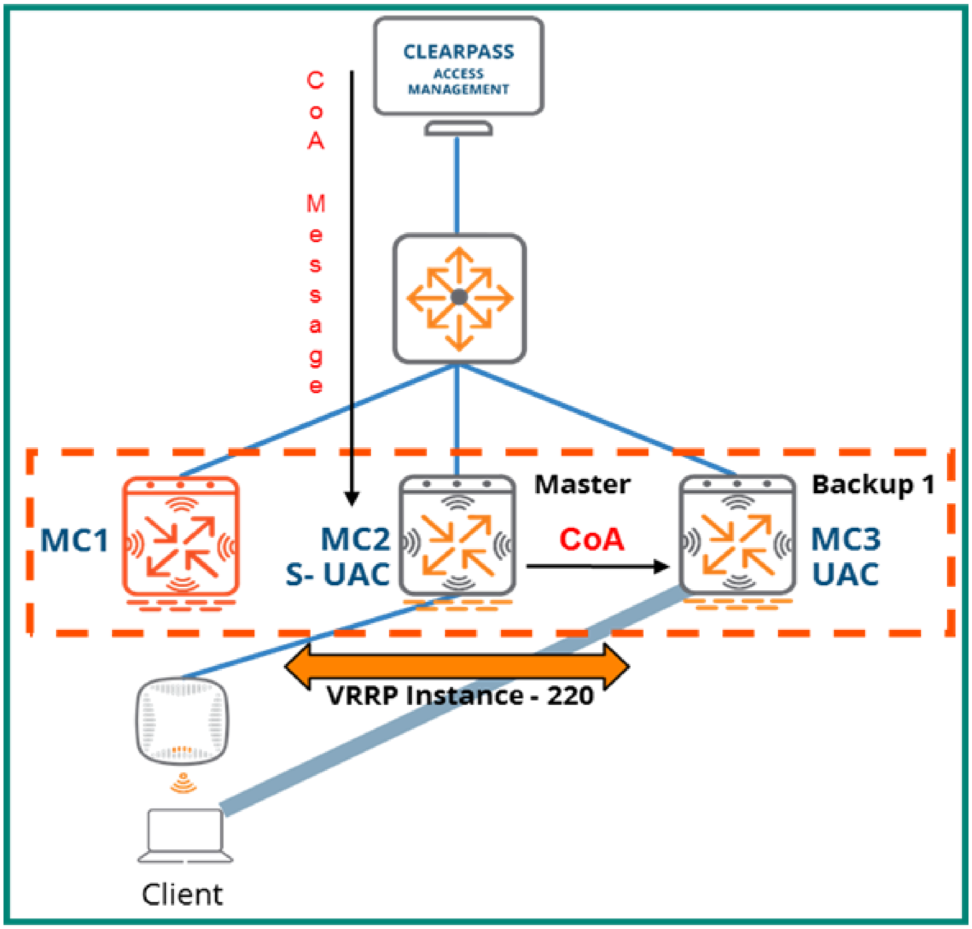

Figure 76 *CoA message forwarded to MC3*

After the change in the CoA request has been successfully implemented, MC3 will send a CoA-ACK back to ClearPass.

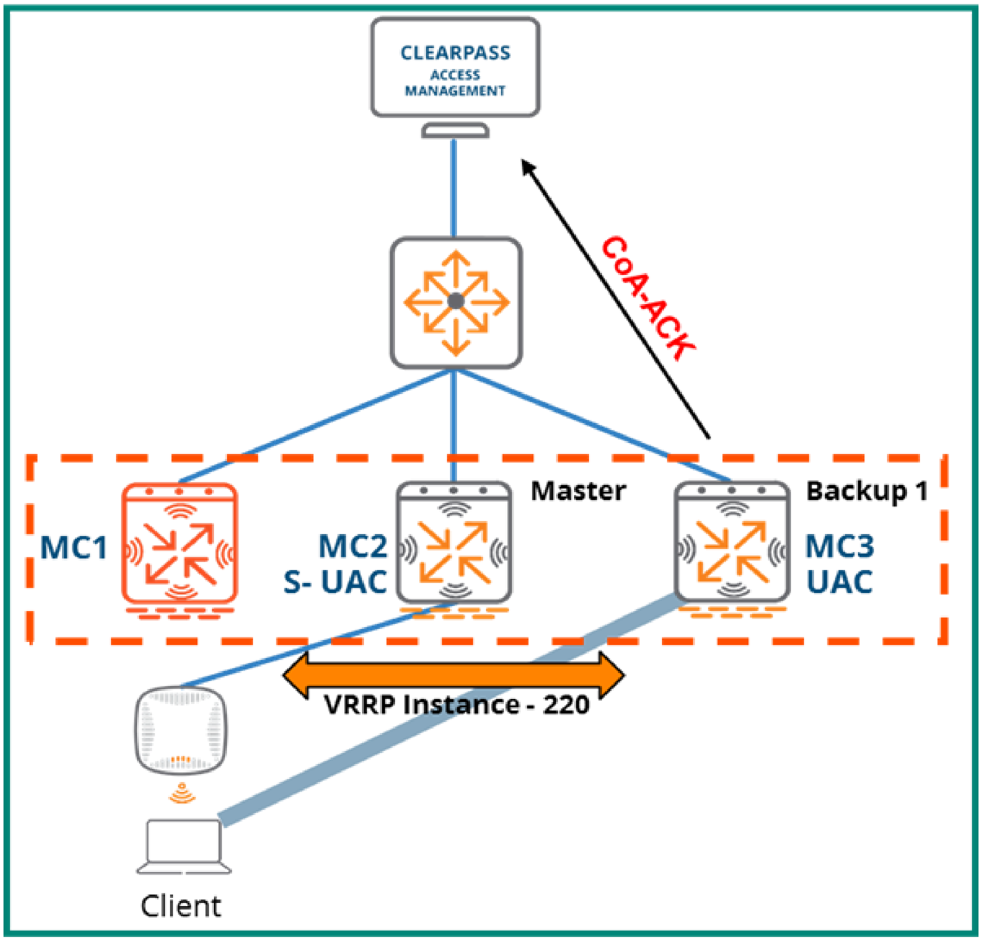

Figure 77 *CoA message forwarded to MC*

CoA with Load Balancing

While not as severe as a device failure, load balancing events within a cluster can pose a challenge for CoA operations if the proper design has not been implemented. To demonstrate Aruba's solution to prevent load-balancing events from impeding CoA functionality the same architecture will be used as the MC failure example. The example will simulate an event where a client associated to MC1 is load-balanced over to MC3 while still in session.

MC1 is operational, master of the VRRP instance 220, and owner of the VIP. As with the previous example, MC1 has inserted VIP1 as the NAS-IP in the RADIUS request which initiated the client authentication so ClearPass will send any CoA requests for the client to the VIP address of VRRP instance 220.

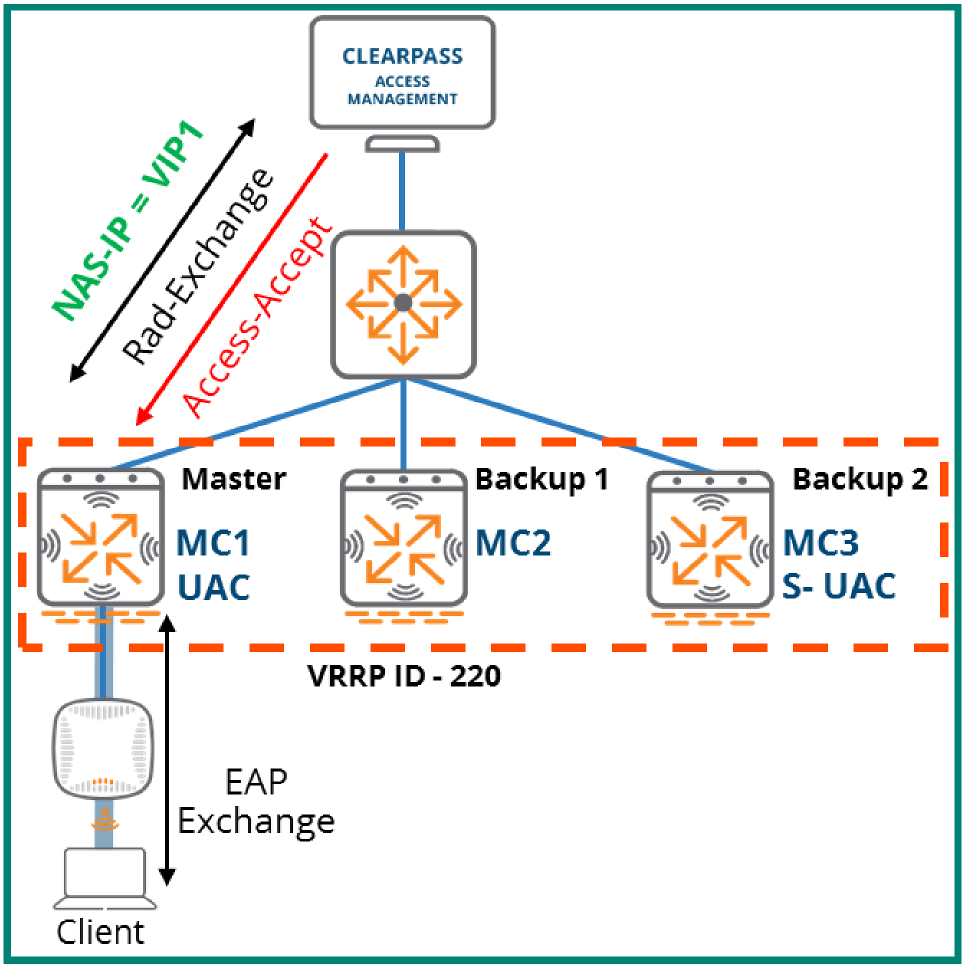

Figure 78 *Client Authenticates Against ClearPass*

Next, the client is load-balanced over to MC3 which becomes its UAC. MC1 becomes the client's S-UAC but remains Master of VRRP instance 220.

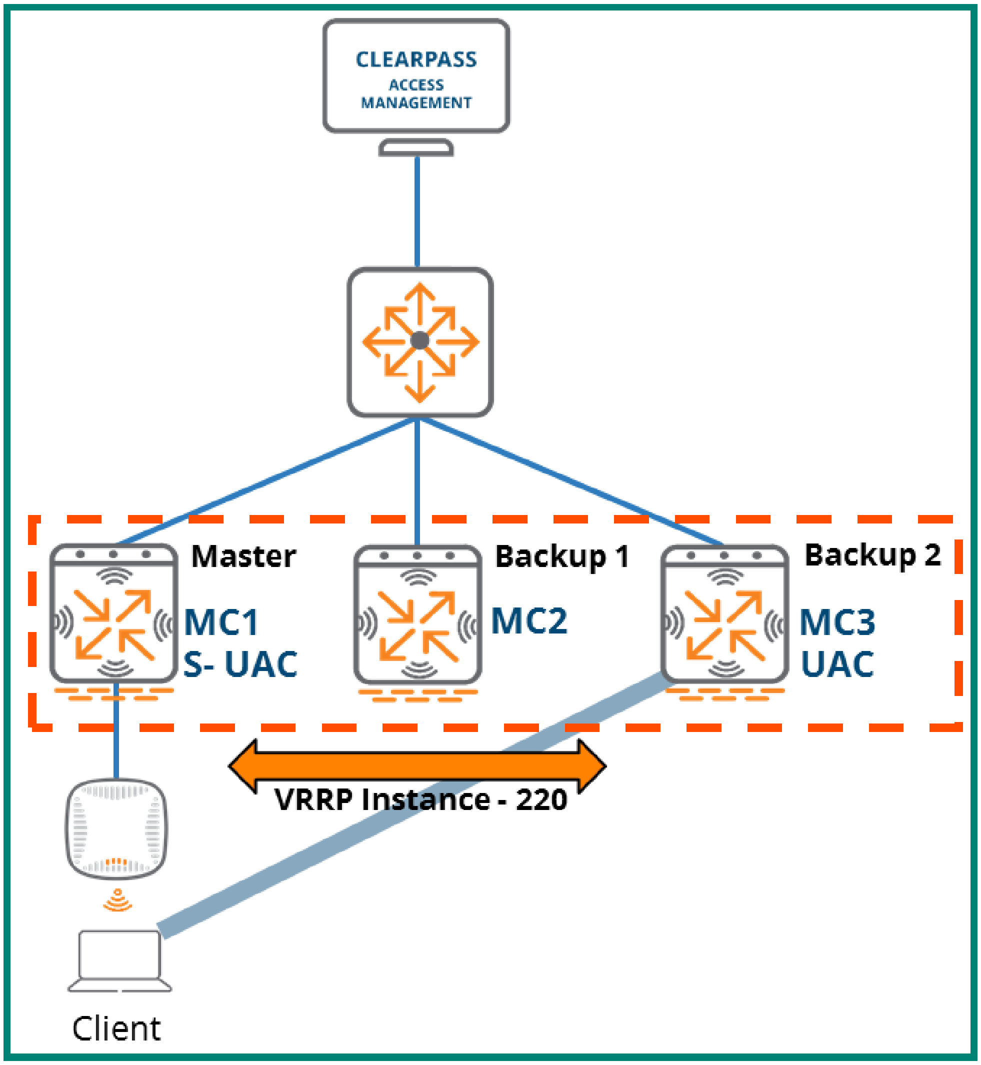

Figure 79 *Client UAC Changes*

ClearPass sends a CoA message to the owner of the VIP for VRRP ID 220:

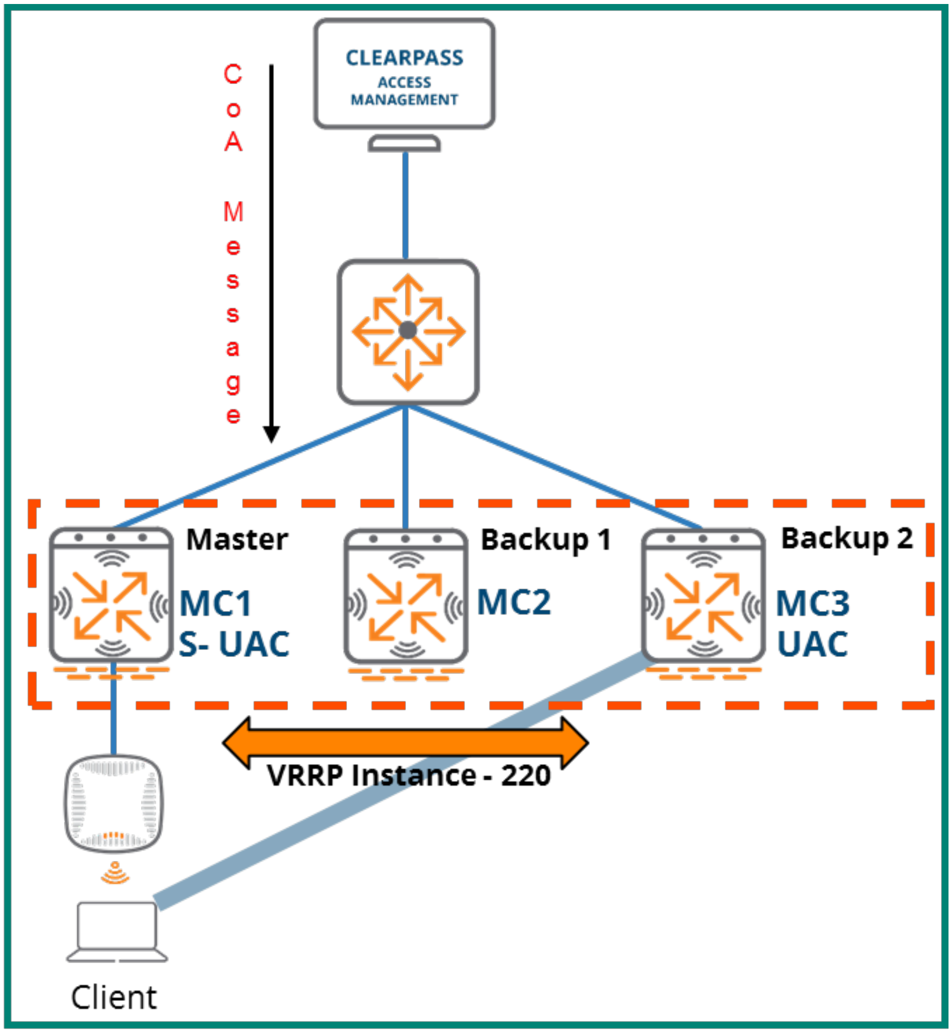

Figure 80 *CoA Message sent to VIP1*

MC1 receives the CoA request as it is the Master for the VRRP instance. It then unicasts the request to the other cluster members.

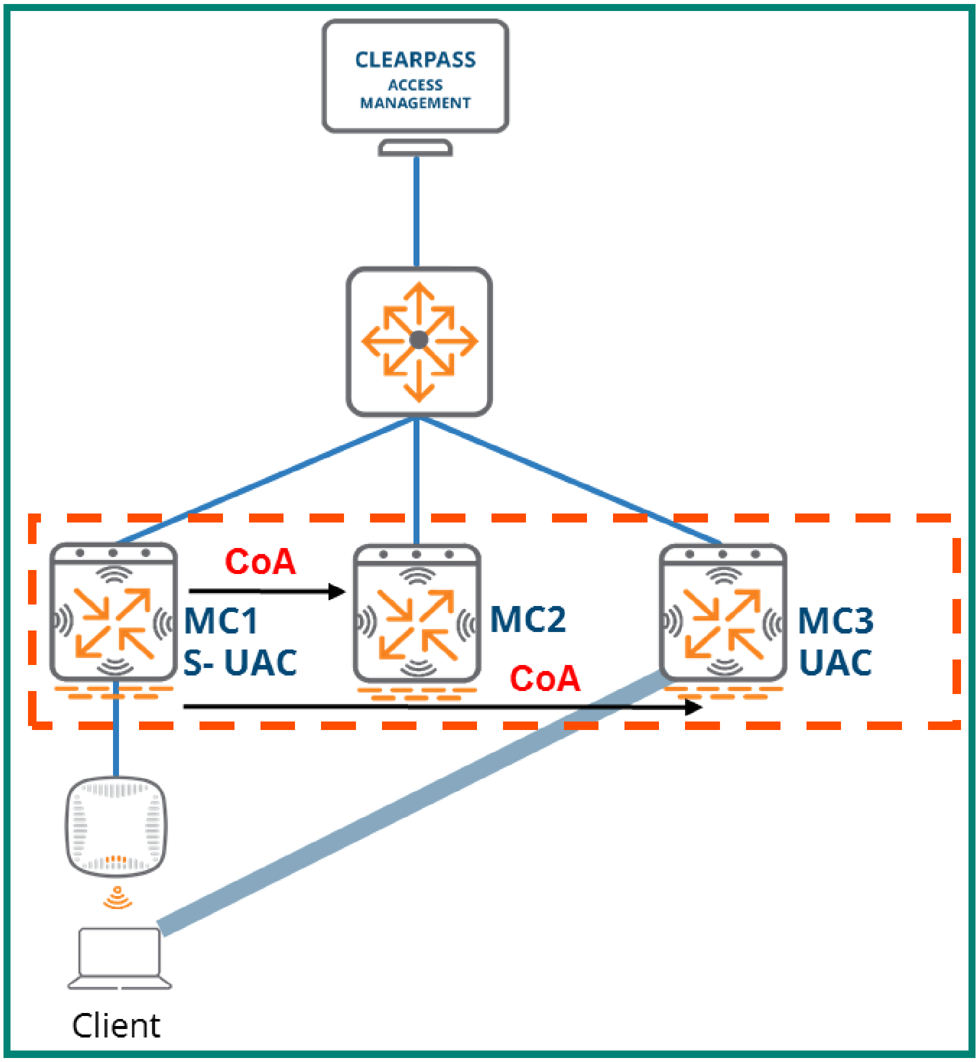

Figure 81 *MC1 Forwards the Request*

Once MC3 receives the forwarded CoA request from MC1 and successfully implements the change it will respond back to ClearPass with a CoA-ACK. The UAC implementing the change is always the MC responsible for returning the CoA-ACK or CoA-NAK to ClearPass.

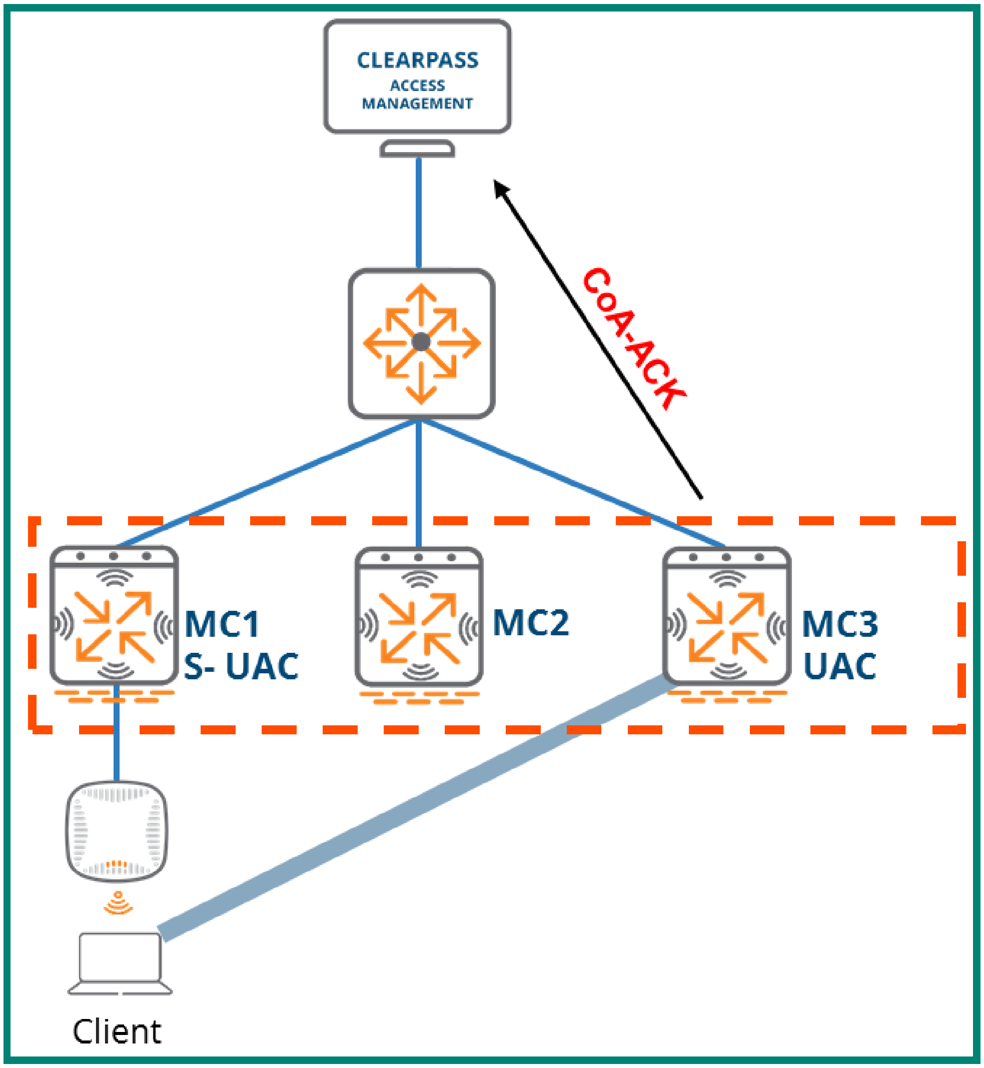

Figure 82 *MC3 Returns CoA-ACK to ClearPass*

Live Upgrade

The Live Upgrade feature allows the MCs and APs in a cluster to automatically upgrade their software to higher ArubaOS versions. MCs in a cluster can be seamlessly upgraded without any adverse impacts to client connectivity and performance. The following points outline the key details of Live Upgrade:

- Seamless in-service cluster upgrade
- No manual intervention with minimal RF impact
- Available as of ArubaOS 8.1
- Applicable to a cluster in an MM environment

Prerequisites

The following ArubaOS features are required in order to enable the Live Upgrade feature with minimal RF impact and client disruptions:

- Stateful failover through an L2-Connected cluster with redundancy enabled
- Centralized image upgrade
- AirMatch (schedule enabled)

NOTE

For additional information on how to configure the perquisites listed above please refer to the ArubaOS 8 User Guide.

Aruba best practices around AP deployment and RF coverage should always be followed to achieve positive results however when performing a Live Upgrade they are a necessity to prevent clients from losing connectivity. APs should be deployed in a capacity-based design so as to guarantee overlapping RF coverage. Doing so ensures that that a client will be always be able to roam to a different AP during the upgrade if the AP where they were previously associated needs to reboot. At a minimum a deployment should be designed so that clients can always see two APs wherever they roam.

CAUTION

In areas without adequate coverage, if the AP a client is connected to reboots during the live upgrade they will lose connectivity until the reboot has finished and the AP comes back up.

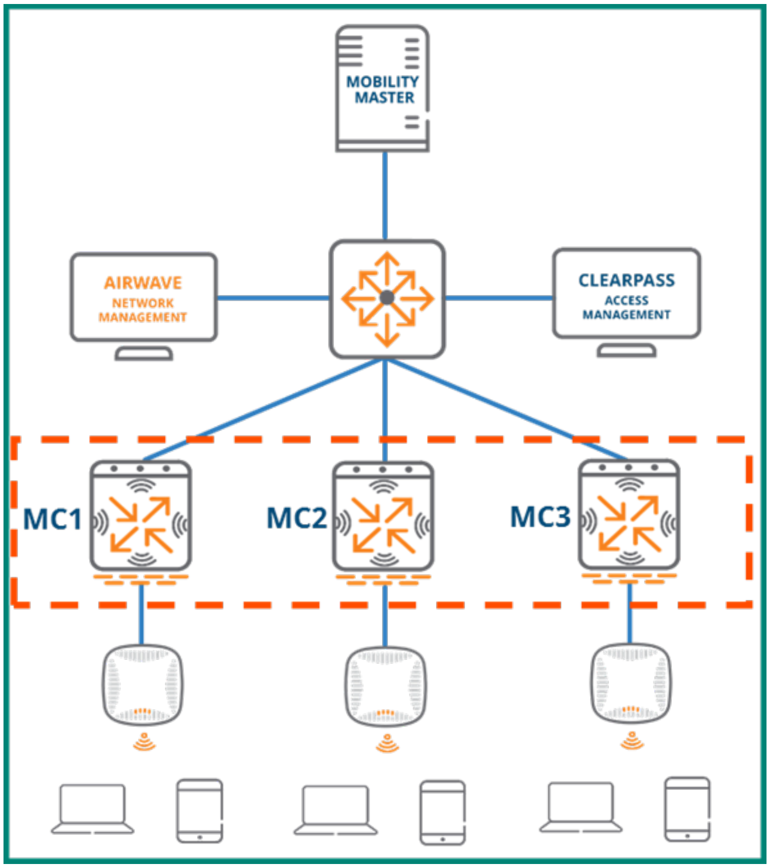

Figure 83 *Reference Live Upgrade Architecture*

Live Upgrade Flow

The Live Upgrade process involves multiple steps to ensure that the upgrade is properly applied and that clients will continue to be served throughout its duration. The high level steps required for the upgrade process are as follows and each step will be covered in detail in the subsequent sections:

1. **AP Partition** - The APs terminating on the cluster are logically grouped into partitions based on their RF channels
2. **Target Controller Assignment** - Each AP partition is assigned an MC in the cluster which serves as a post-upgrade termination target for the APs in that partition. All cluster members are used as AP partition targets with the exception of one
3. **New Firmware Pushed to MCs** - All of the cluster members download the new ArubaOS firmware through the Centralized Image Upgrade feature using a pre-configured upgrade profile
4. **Cluster Members Upgrade** - The actual upgrade process begins by rebooting one of the cluster members to the new firmware. Once the first controller receiving the upgrade comes back up loaded with its new firmware the APs in the partitions targeting that controller are pre-loaded with the new firmware. Those APs reboot one partition at a time and come up on their upgraded target controller

After all APs have rebooted for the first target controller, a second controller in the cluster is rebooted followed by the AP partitions targeting it in the same manner as described above.

After all APs are upgraded and terminated on their upgraded target controllers then the last controller that has been exempted from targeting is rebooted to come up and join the upgrade cluster.

Initial Lab AP Distribution

The image below depicts the scenario that will be used to demonstrate the key concepts of the Live Upgrade feature for the duration of this section. The selected architecture is a standard ArubaOS 8 design consisting of a fully-redundant MM pair deployed on a virtual appliance. The scenario has three MCs in a single cluster which have been configured in an L2-Connected state with full redundancy enabled between them as specified by the prerequisites. Connected to the three MCs are seven APs which are operating across a variety of different channels.

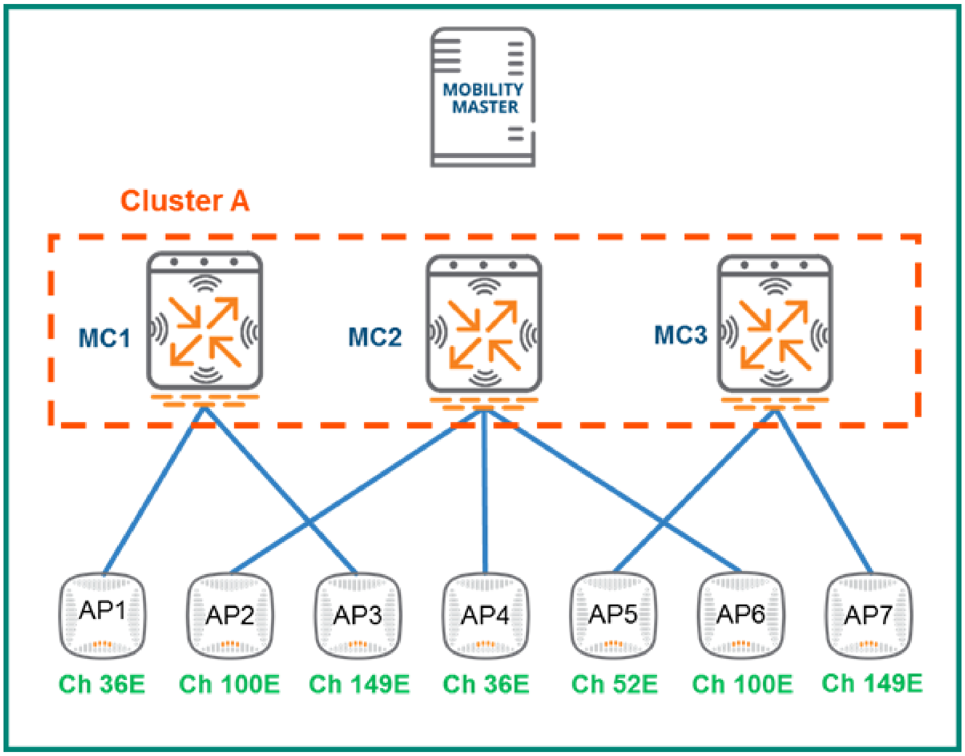

Figure 84 *Sample Network for Live Upgrade*

NOTE

The channel assignments for the APs were chosen at random strictly for purposes of demonstrating Live Upgrade operations. They do not represent a best practice recommendation for channel assignment in a production network.

AP Partition

The Live Upgrade process begins when the Upgrade Manager for the cluster sends information about all of the APs connected to the cluster to AirMatch. This is why it is a prerequisite to leave the default setting of enabling the AirMatch schedule intact to perform a Live Upgrade. Upon receipt of the AP info from the Upgrade Manager, AirMatch will segregate the APs into logical groups and update the Upgrade Manager with the partition to AP mapping assignments.

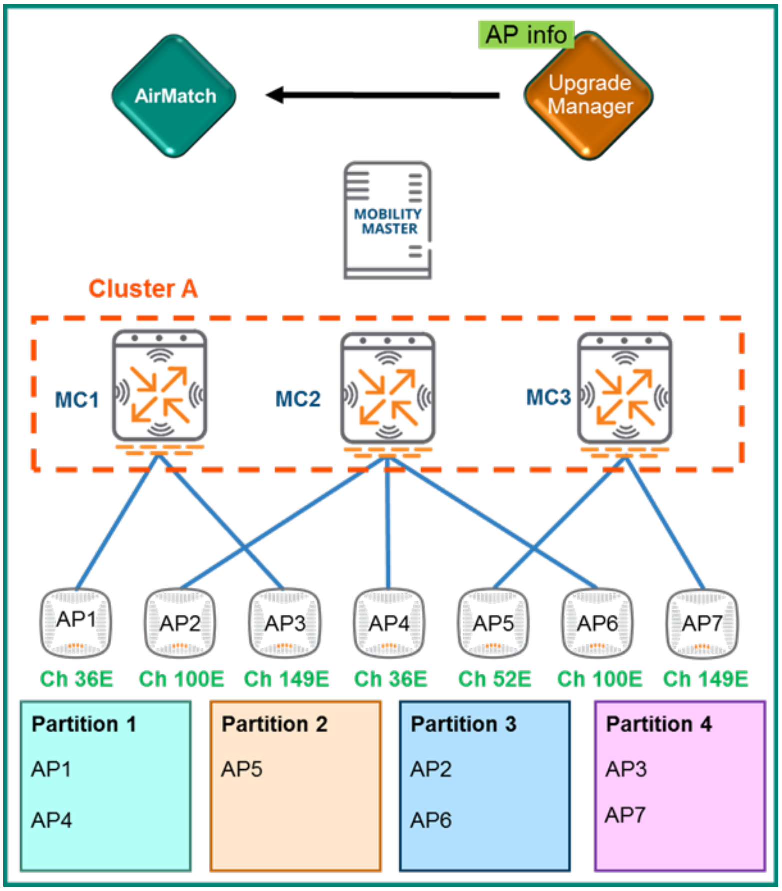

Figure 85 *AP Partitions*

Target Assignment and Firmware Download

After all APs have been logically partitioned by AirMatch based on their information and the Upgrade Manager has been updated with the assignments each partition is assigned a “target” MC. The target represents the MC managing the APs each partition after they reboot with their new firmware. As the figure below demonstrates, the APs in Partitions 1, 2, and 3 have been assigned MC3 as their target while the APs in Partition 4 have been assigned MC1:

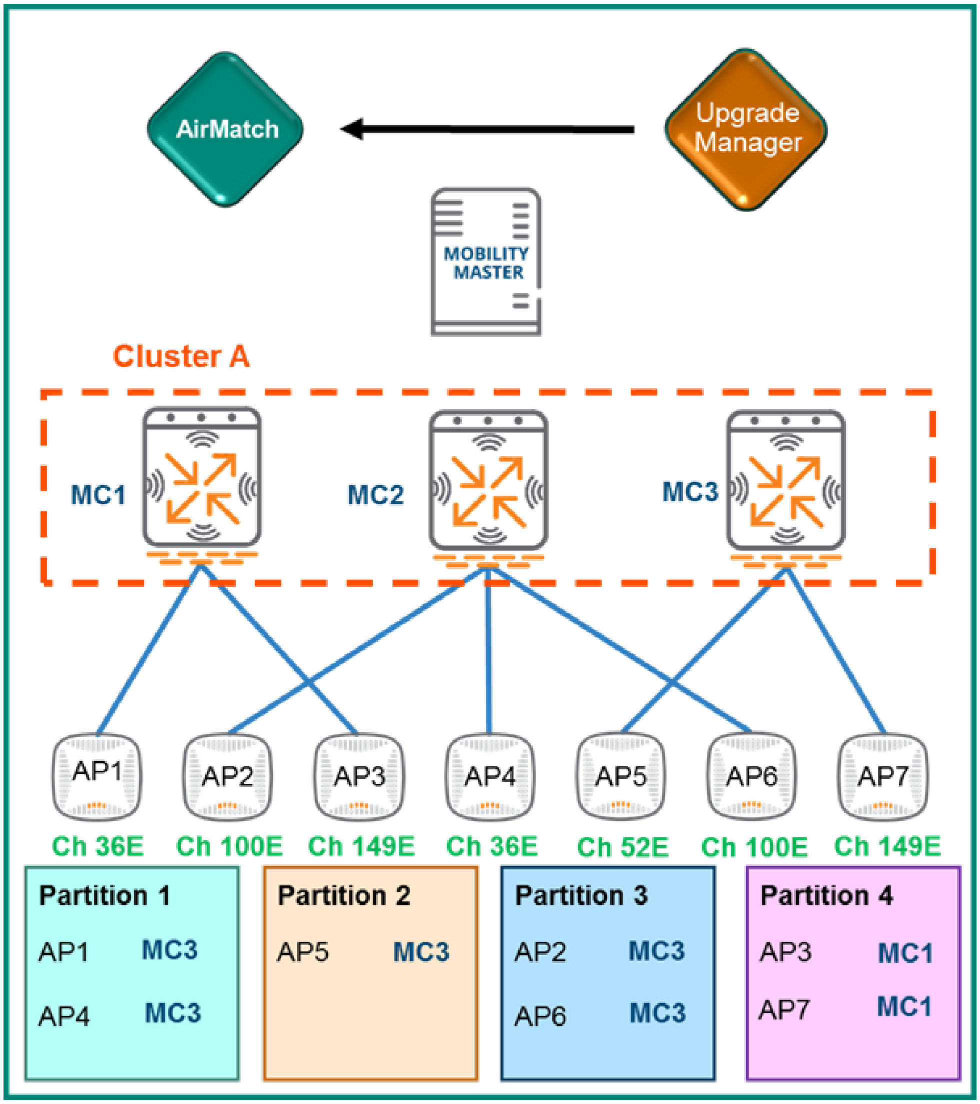

Figure 86 *Target Assignments*

MC2 has not been designated as a target for any partition or APs. It was excluded intentionally and the reasons why will be discussed in the following sections.

Once all partitions have their target assignments the MCs will download the new ArubaOS firmware one at a time. They will continue normal operation while their respective downloads take place.

Cluster Members Upgrade

Once all the APs have been partitioned, assigned a target MC, and the MCs have pre-downloaded their new firmware they are ready to commence the actual upgrade process. The sample cluster has three members, however the process will be identical regardless of how many MCs are in the cluster. Each cluster member reboots one by one and the final cluster member which was not designated as a target for any APs reboots last.

First Member Upgrade

The upgrade process starts when MC3 reboots so that it's pre-downloaded firmware upgrade can take effect. During the reboot process MC3 will go down and be unable to continue serving as the AAC for APs 5 and 7 therefore they will need to failover to MC1 and MC2, respectively. Likewise, any clients with MC3 assigned as their UAC will need to failover to MC1 and MC2.

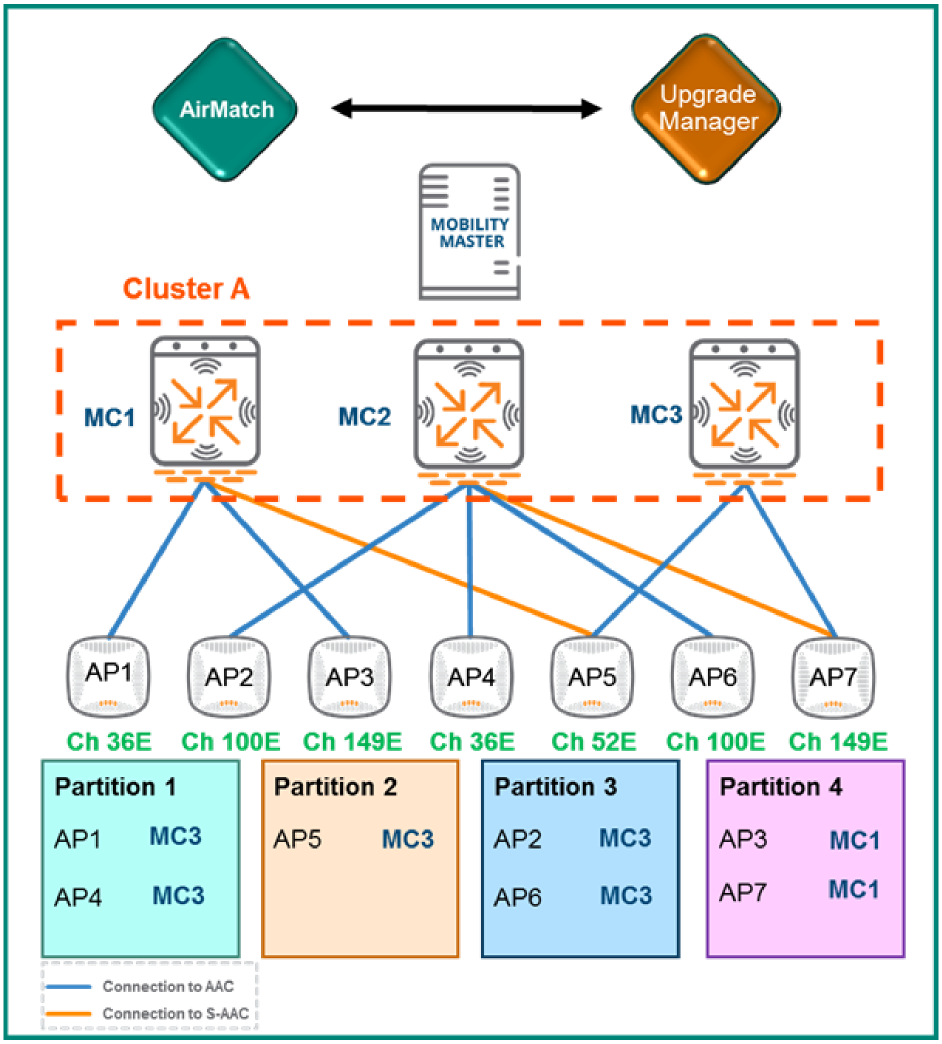

Figure 87 *MC3 begins reboot with APs and clients failing over*

After MC3 finishes rebooting and comes back online running its new firmware it will form a separate cluster since MCs in the same cluster must run the same firmware version. The new cluster MC3 has formed is represented in the figure below with a green border and will be referred to as Cluster B. At this point MC1 and MC2 along with all associated APs and clients are still in Cluster A running the previous firmware version.

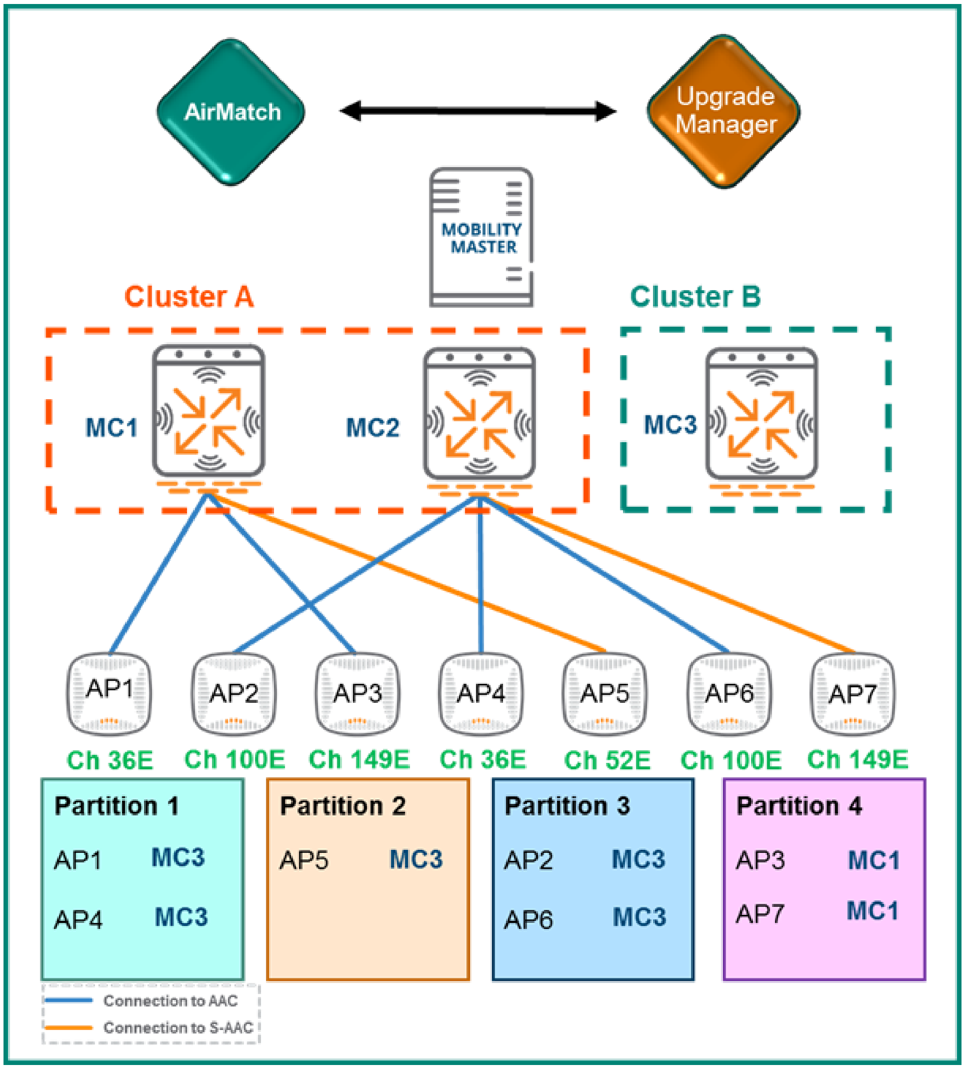

Figure 88 *MC3 Reboots*

Once MC3 is back online it will assume operation as the AAC for the APs which had it assigned as their target MC. The APs will first preload their new firmware and then reboot. The reboot will occur one partition at a time to ensure that clients will not be deprived of AP options through which to maintain connectivity. If all APs were to reboot at once ARM's coverage hole detection mechanism wouldn't be able to adequately compensate for all of the gaps. This could result in clients being forced off the network until the APs had finished rebooting. Clients associated to rebooting APs will roam to APs attached to the red Cluster A. These clients will only need to go through a 4-way 802.1X handshake rather than a full authentication process since they will retain their UAC and it is already a member of the cluster.

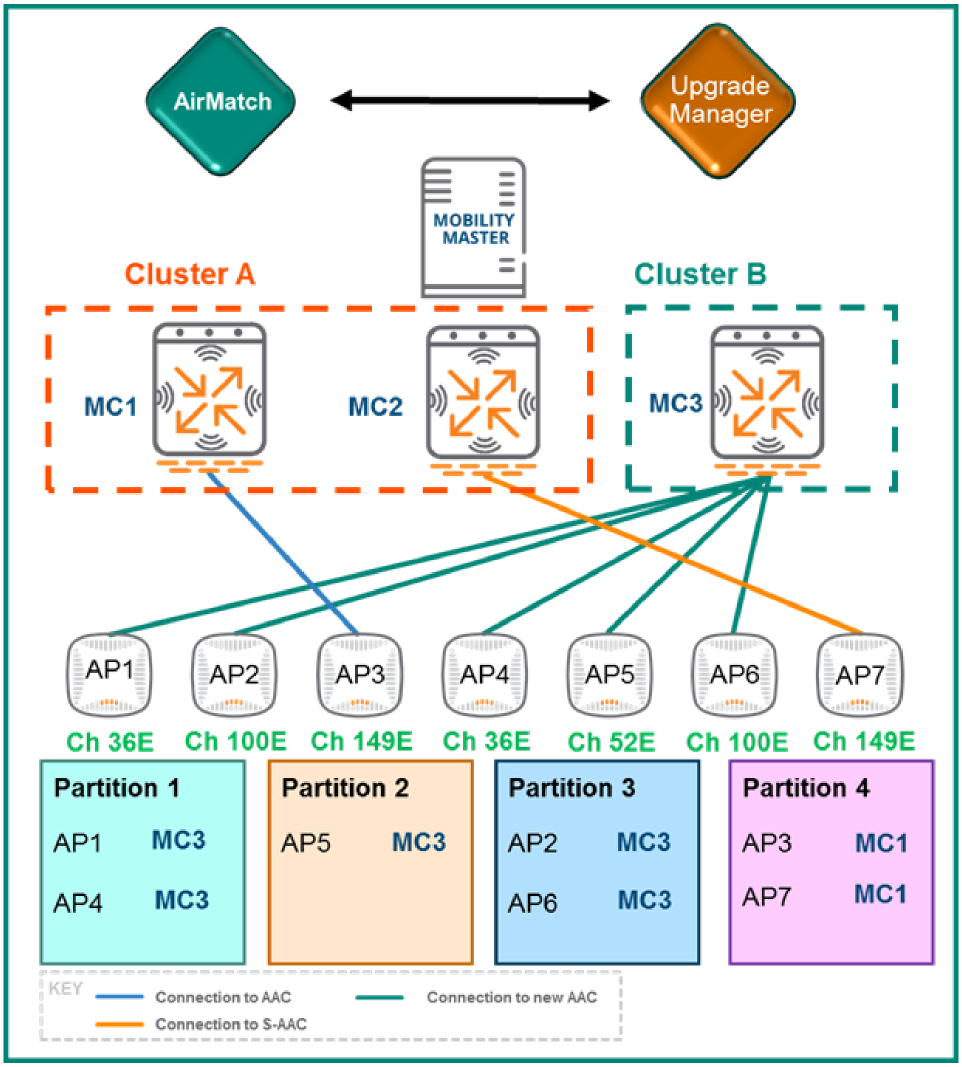

Partition 1		Partition 2		Partition 3		Partition 4	
AP1	MC3	AP5	MC3	AP2	MC3	AP3	MC1
AP4	MC3			AP6	MC3	AP7	MC1

Figure 89 *APs 1-2 and 4-6 connect to Cluster B*

Second Member Upgrade

Now that MC3 is operational with its new firmware and its designated APs have rebooted MC1 is able begin its upgrade using an identical process. MC1 will reboot and leave Cluster A making MC2 the last member of the original cluster. After MC1 reboots, any APs which it was serving as AAC will failover to their S-AAC (MC2) and any clients will failover to their S-UAC (MC2 as well). In our example this means that AP3 will failover to MC2 along with any clients whose UAC was MC1.

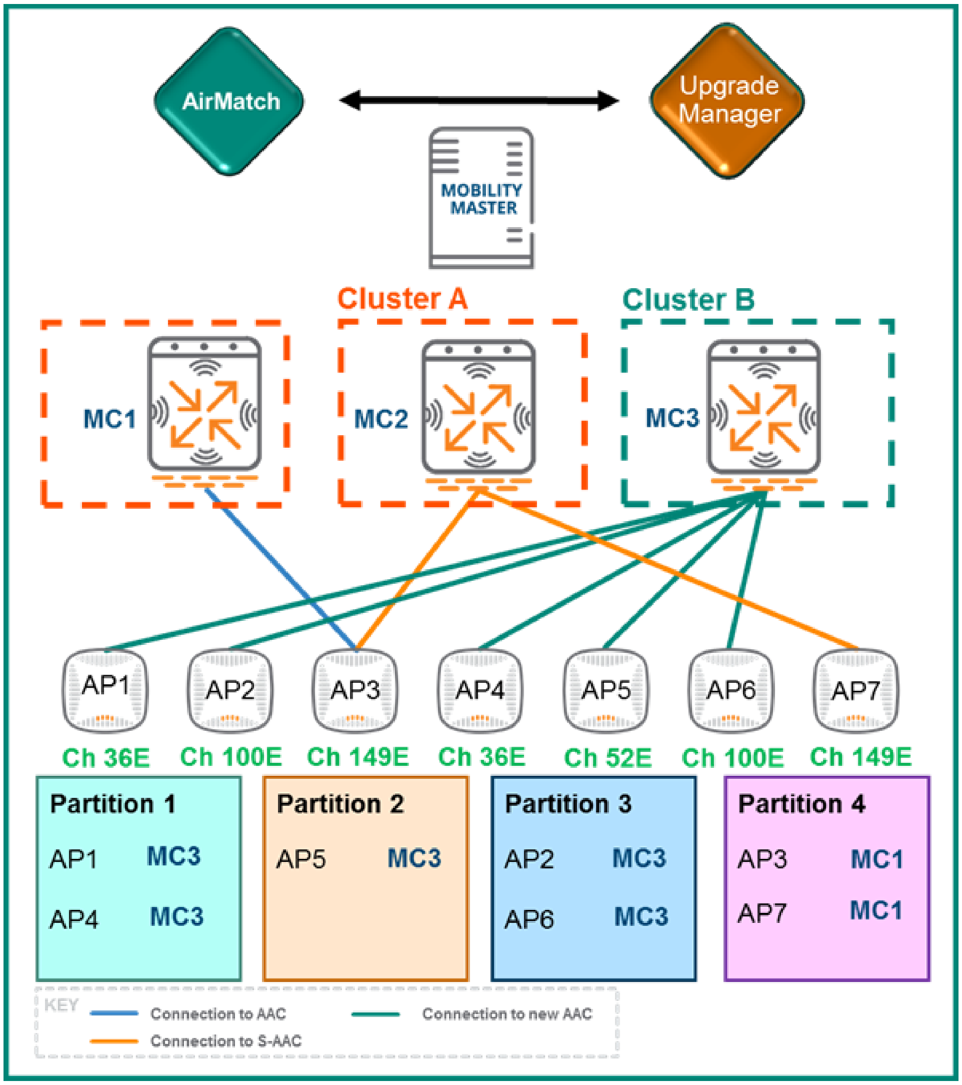

Figure 90 *APs and Clients failover to MC2 while MC1 reboots*

Once MC1 comes back online it will immediately join MC3 in Cluster B. AP3 and AP7 in Partition 4 were assigned MC1 as their target and will pre-download the new firmware so that they can reboot and join the new cluster. Once the firmware has been pre-downloaded the APs reboot causing their associated clients to immediately roam to APs attached to the green Cluster B. Even though the clients were previously connected to MC1 and MC3 the controllers are considered new devices from the client perspective since they upgraded their firmware and formed a new cluster. This will require any 802.1X client to go through the full authorization process just as they would when associating to a new device. Once AP3 and AP7 come back up with their new firmware they will connect to MC1 as their AAC.

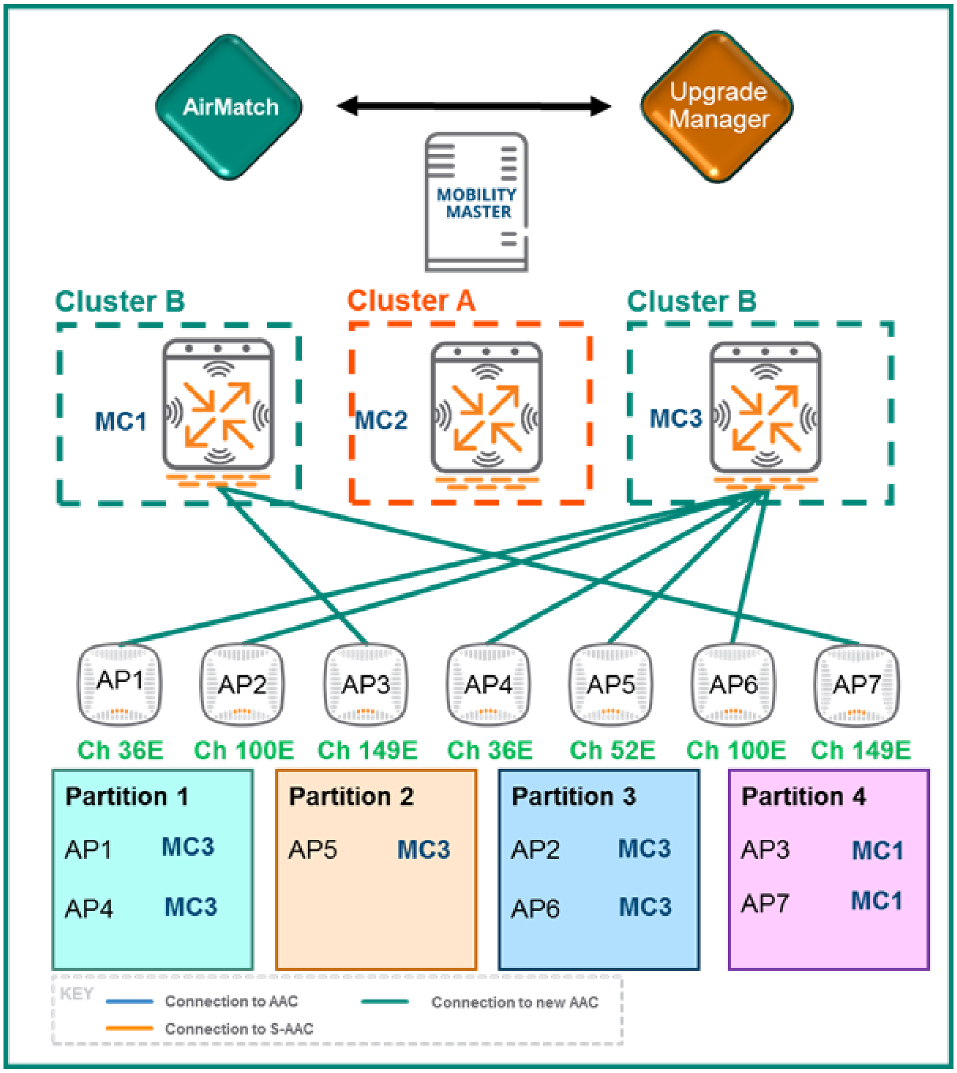

Partition 1		Partition 2		Partition 3		Partition 4	
AP1	MC3	AP5	MC3	AP2	MC3	AP3	MC1
AP4	MC3			AP6	MC3	AP7	MC1

Figure 91 *APs 3 and 7 connect to MC1*

Final Member Upgrade

After all cluster members have received their upgrade and rebooted the final member (MC2) is able to begin its upgrade as well. MC2 doesn’t have any associated clients or APs since it was deliberately excluded as a target controller for the partitions therefore no APs or clients will need to be redistributed. This means that MC2 can simply reboot, come back online with its new firmware, and rejoin the other two MCs in Cluster B. After it rejoins the cluster, MC2 will become available for AP and client load balancing as determined by the cluster leader.

NOTE

For additional information on cluster load balancing please refer to the Client Load Balancing and AP Load Balancing sections.

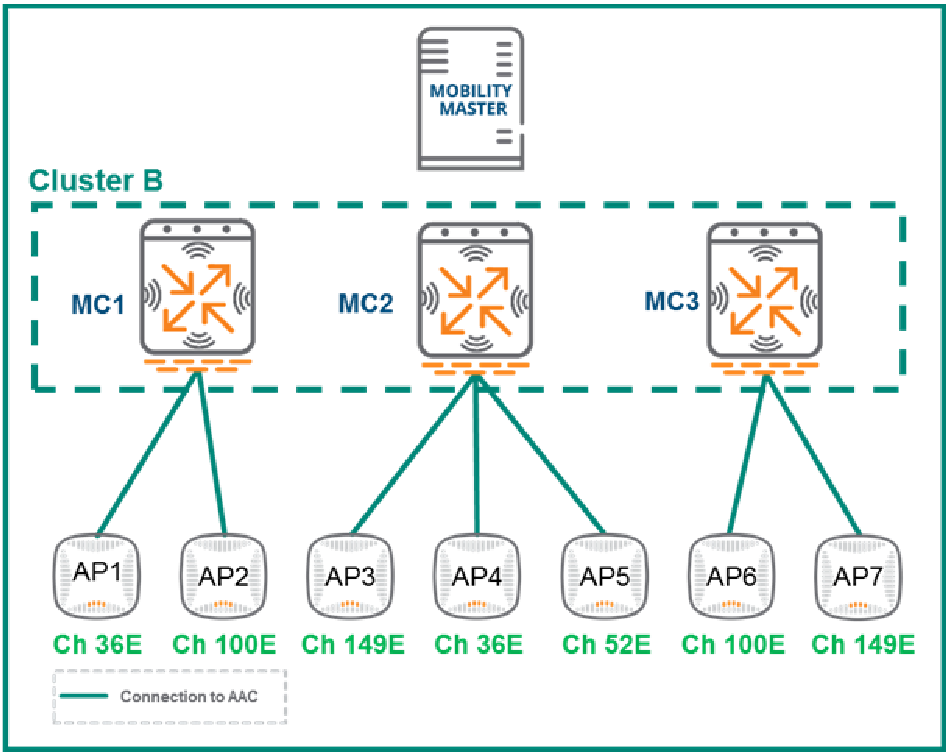

Figure 92 *Live Upgrade complete*

Centralized Licensing

ArubaOS has supported a centralized licensing model since the introduction of ArubaOS 6. The way this model worked historically was licenses could be installed on one controller and other controllers would “subscribe” to withdraw licenses from a global license pool as needed. However, a significant limitation for this model was that in some customer deployments there was no way to control how many licenses one controller could withdraw from the pool. This limitation led to some situations where license pools would be depleted when one site deployed more APs than there were licenses available in the pool.

With the introduction of ArubaOS 8, the MM now supports the creation of smaller licensing pools within the global pool. This method of segmentation allows for the limitation or reservation of licenses that a specific controller or group of controllers are allowed to withdraw from the global pool.

Licensing Concepts

License Types

Many key abilities and features of the ArubaOS 8 platform are enabled with the use of licenses. These licenses are generally installed on the MM but can also be installed on an MCM or a Standalone controller if required. Additionally, centralized licensing allows MCs under management of an MM or MCM to “subscribe” and draw the number of licenses they need on demand.

ArubaOS 8 licenses are broken up into three different classes or categories of licenses:

Device Licenses	Feature Licenses	Session Licenses
MM-VA	LIC-PEF	LIC-VIA
MM-HW*	LIC-RFP	LIC-ACR
MC-VA	LIC-PEFV	
LIC-AP	SUBX-WebCC**	
LIC-ENT		

Table 14 *Licenses in ArubaOS 8*

* MM-HW licenses are integrated within the Hardware MM.

** WebCC is a subscription-based license ‘X’ equals a 1, 3, 5, 7, or 10 year subscription.

- **Device-Based Licenses** – Licenses that enable device functionality
    - **MM-VA/MM-HW*** – Enable controllers and APs to terminate on an MM. The MM-HW licenses are pre-installed on the MM hardware appliance
    - **MC-VA** – Enable VMCs to terminate APs. These can be installed on the MM or directly on any VMC that is not terminated on a MM
    - **LIC-AP** – Licenses installed on the MM, MCM, or Standalone controller to enable AP termination
- **Feature-Based Licenses** – Licenses that enable specific software features
    - **LIC-PEF** – Enables the Policy Enforcement Firewall (PEF) feature
    - **LIC-RFP** – These enable the RF Protect (RFP) features for WIDS/WIPS and Spectrum Analysis support
    - **LIC-PEFV** – Enables PEF support on Virtual Internet Access (VIA) client roles as a platform license applied to each controller. Being phased out by LIC-VIA session licenses
    - **SUBX-WebCC** – This subscription-based license that enables web filtering support in 1, 2, 3, 5, and 7 year subscriptions
    - **LIC-ENT** – This is a bundled combination of the AP, PEF, RFP, and LIC-AW (AirWave) licenses
- **Session-Based Licenses** – Define a feature based on the number of concurrent sessions allowed across the controller
    - **LIC-VIA** – Enables Virtual Internet Access (VIA) clients to connect and establish a tunnel to a controller
    - **LIC-ACR** – Advanced Cryptography (ACR) licenses enable the use of Suite B licenses on the controller

MM Licensing

An MM in ArubaOS 8 serves as the centralized licensing server for any MCs under its management. APs and controllers will draw from the centralized MM license pool when the MM is serving as the licensing master. Table 15 below provides a description of the license consumption process on the MM.

VMM license consumption occurs using a slightly different method because the machines are virtual appliances. Only one MM-VA-XX license needs to be purchased even if multiple VMMs are deployed for redundancy. If a VMM is being used to support up to 5,000 devices then only one MM-VA-5K license is required and multiple VMMs can be provisioned to manage the WLAN. Additionally, smaller MM-VA licenses can be stacked to support larger numbers of devices on the MM. However, there is a point where stacking smaller licenses can cost more than a single larger MM-VA license.

Licenses are pre-installed on hardware MMs because they are hardware appliances and they are not capable of stacking licenses, leading to a higher licensing cost. E.g., if a deployment has to support up to 5,000 devices, then two MM-HW-5K appliances will be required, even though together they will still only support up to 5,000 devices.

For all other licenses (AP, PEF, RFP, etc.) it is only necessary to purchase the necessary license quantities and the licensing database will be shared between the MMs.

License Type	Consumption Method
MM-VA/MM-HW	Each controller or AP will consume 1 license
MC-VA	Each AP terminated on a VMC will consume 1 license
AP, PEF, RFP	Each AP terminated on a controller will consume 1 license
LIC-PEFV	Each LIC-PEFV license is applied to each controller
SUBX-WebCC	Each AP terminated on a controller will consume 1 license
LIC-VIA	Each "VIA user session" on a controller will consume 1 license
LIC-ACR	Each "SuiteB" client or tunnel will consume (1) license

Table 15 *License Consumption in ArubaOS 8*

License Model Examples

- **Sample Deployment 1** – 800 APs with AP, PEF, RFP licenses managed by hardware MCs (MCs) and Virtual MMs (VMMs). With this sample deployment, each of the 800 APs will need an AP, PEF, and RFP license. Having a VMM requires enough MM-VA licenses to cover every AP and MC under its management. The MM-VA-1K license provides a pool of 1000 licenses. 800 of the 1000 licenses in the pool are consumed by APs meaning there are 200 licenses left to cover the MCs as well as any future addition of devices.

License Type	Quantity
MM-VA-1K	1
LIC-AP	800
LIC-PEF	800
LIC-RFP	800

Table 16 *License Sample Model 1*

- **Sample Deployment 2** – 250 APs with AP and PEF licenses using VMCs and VMMs. In this case, the MM-VA-500 license provides a pool for up to 500 devices on the MM. The MC-VA-250 license will enable up to 250 APs to terminate on any number of VMCs under the MM (could be a single VMC or multiple VMCs).

License Type	Quantity
MM-VA-500	1
MC-VA-250	1
LIC-AP	250
LIC-PEF	250

Table 17 *License Sample Model 2*

- **Sample Deployment 3** – 6,000 APs with AP, PEF, and RFP licenses using hardware MCs and Hardware MMs (HMM) with redundancy for clustering. The sizing specifications below show two 10k MM appliances, six 7240XM controllers to support all 6,000 APs within a cluster, and 6000 AP licenses for AP, PEF, and RFP respectively.

License Type	Quantity
MM-HW-10K	2
7240XM MCs	6
LIC-AP	6000
LIC-PEF	6000
LIC-RFP	6000

Table 18 *License Sample Model 3*

- **Sample Deployment 4** – 2000 APs with AP, PEF, and RFP licenses supporting 1,500 clients that require VIA and Suite B cryptography, using Hardware MCs and Hardware MMs. The MM-HW-5k license provides a pool for up to 5,000 devices on the MM.

License Type	Quantity
MM-HW-5K	2
7240XM MCs	2
LIC-AP	2000
LIC-PEF	2000
LIC-RFP	2000
LIC-VIA	1500
LIC-ACR	1500

Table 19 *License Sample Model 4*

- **Sample Deployment 5** – 2000 APs with AP and PEF licenses using Hardware MCs and HMMs supporting up to 50 clients per AP up to 10,000 total clients. The client count is much higher than average so even though the total AP count is only 2,000 the largest HMM must be able to accommodate all of the clients.

License Type	Quantity
MM-HW-10K	2
7240XM MCs	2
LIC-AP	2000
LIC-PEF	2000

Table 20 *License Sample Model 5*

MCM Licensing

The MC Master (MCM) is similar to the “Master-Local” controller architecture from Aruba OS 6 where a dedicated hardware controller serves as the central licensing server for all the managed locals as well as a central configuration point.

NOTE

MCM mode is only supported on the 7030 and 7200 series controllers. 7024 and smaller and VMCs cannot serve as an MCM device.

Licenses can be configured and installed on to the MCM controller. The device (except for the MM-VA license), feature, and session based licenses should scale and use the same considerations based on their consumption requirements (similar to the MM).

However, the MC-VA licenses for VMC controllers that will be managed by an MCM will be installed on the individual VMCs. I.e. VMCs under an MCM cannot share licenses as the could under an MM. Additionally, the MC-VA license must be installed whole and must match the platform. E.g., a single MC-VA-250 license cannot be purchased and then split up across 5 VMCs with 50 licenses each, The MC-VA-250 would need to be installed on an MC-VA-250 (or smaller) non-MM managed VMC.

Standalone Licensing

When licensing a Standalone controller, all of the same device, feature, and session based licenses can be installed on a Standalone controller and are consumed the same manner.

With MC-VA licenses on standalone controllers must be installed whole and must match the platform. For example, a single MC-VA-250 license cannot be purchased and then split up across 5 VMCs with 50 licenses each, The MC-VA-250 would need to be installed on an MC-VA-250 (or smaller) non-MM managed VMC.

License Activation and Migration

MyNetworking Portal

The MyNetworkingPortal (MNP) is HPE's licensing portal and is used for a variety of support activities including:

- Activation of new licenses from purchases
- Activation of demo and evaluation licenses
- License management (migrate, change ownership, etc.),
- Software and user guides

MNP contains a main landing page that is displayed after login directs users to the appropriate resources they are attempting to locate.

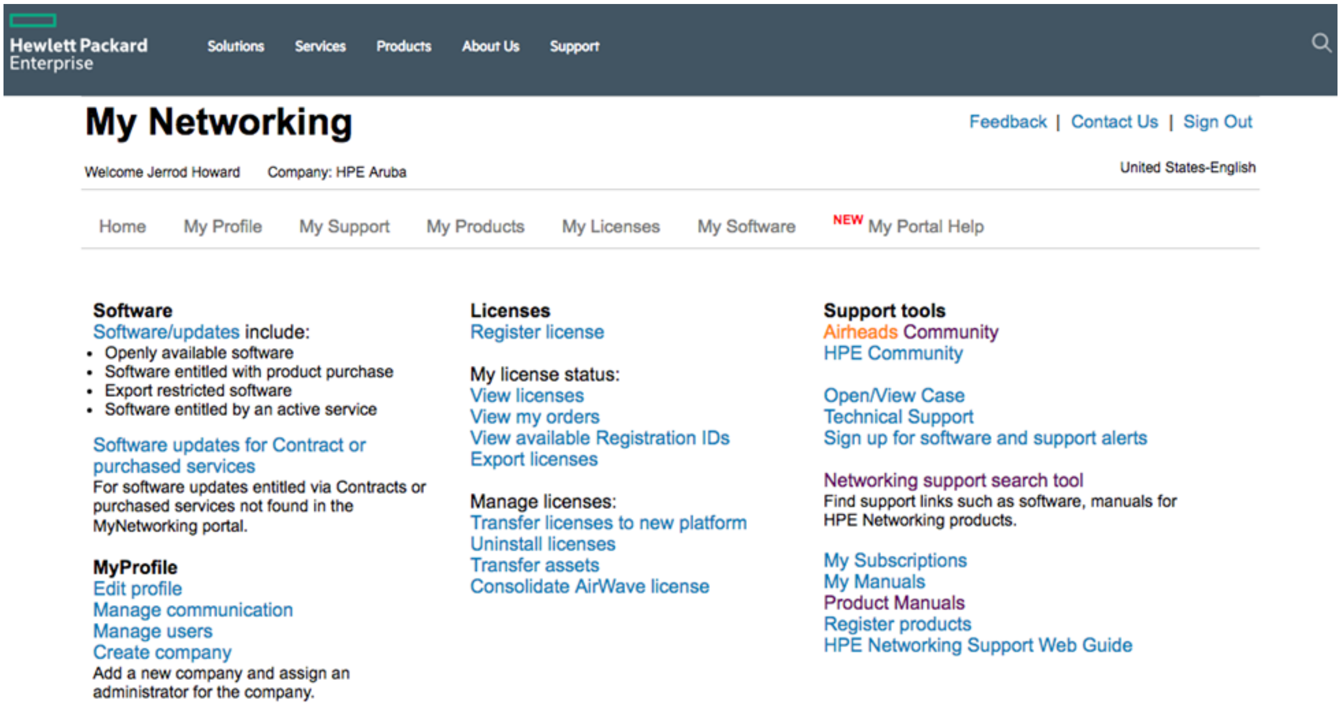

My Networking

Welcome Jerrod Howard Company: HPE Aruba

United States-English

Software
Software/updates include:
- Openly available software
- Software entitled with product purchase
- Export restricted software
- Software entitled by an active service

Software updates for Contract or purchased services
For software updates entitled via Contracts or purchased services not found in the MyNetworking portal.

MyProfile
Edit profile
Manage communication
Manage users
Create company
Add a new company and assign an administrator for the company.

Licenses
Register license

My license status:
View licenses
View my orders
View available Registration IDs
Export licenses

Manage licenses:
Transfer licenses to new platform
Uninstall licenses
Transfer assets
Consolidate AirWave license

Support tools
Airheads Community
HPE Community

Open/View Case
Technical Support
Sign up for software and support alerts

Networking support search tool
Find support links such as software, manuals for HPE Networking products.

My Subscriptions
My Manuals
Product Manuals
Register products
HPE Networking Support Web Guide

Figure 93 *MyNetworkingPortal Dashboard*

When a new order is placed, a Sales Order is generated that contains an associated *Registration ID* or *Certificate ID*. Once logged in to MNP, click on **Register License** and enter the same Registration ID or Certificate ID from the Sales Order. A prompt will then appear asking for the email address corresponding to the order. This email address may belong to an individual or an organization depending on how the order was made, however in either case the appropriate address will appear on the sale invoice. Once entered, a list of all the licenses that are ready for activation are listed. The overall quantity of licenses as well as the number of licenses that awaiting activation will be listed. As licenses are activated from the new sales order the number of available licenses will decrease. MNP can be accessed through the following URL:

https://hpe.com/networking/mynetworking/

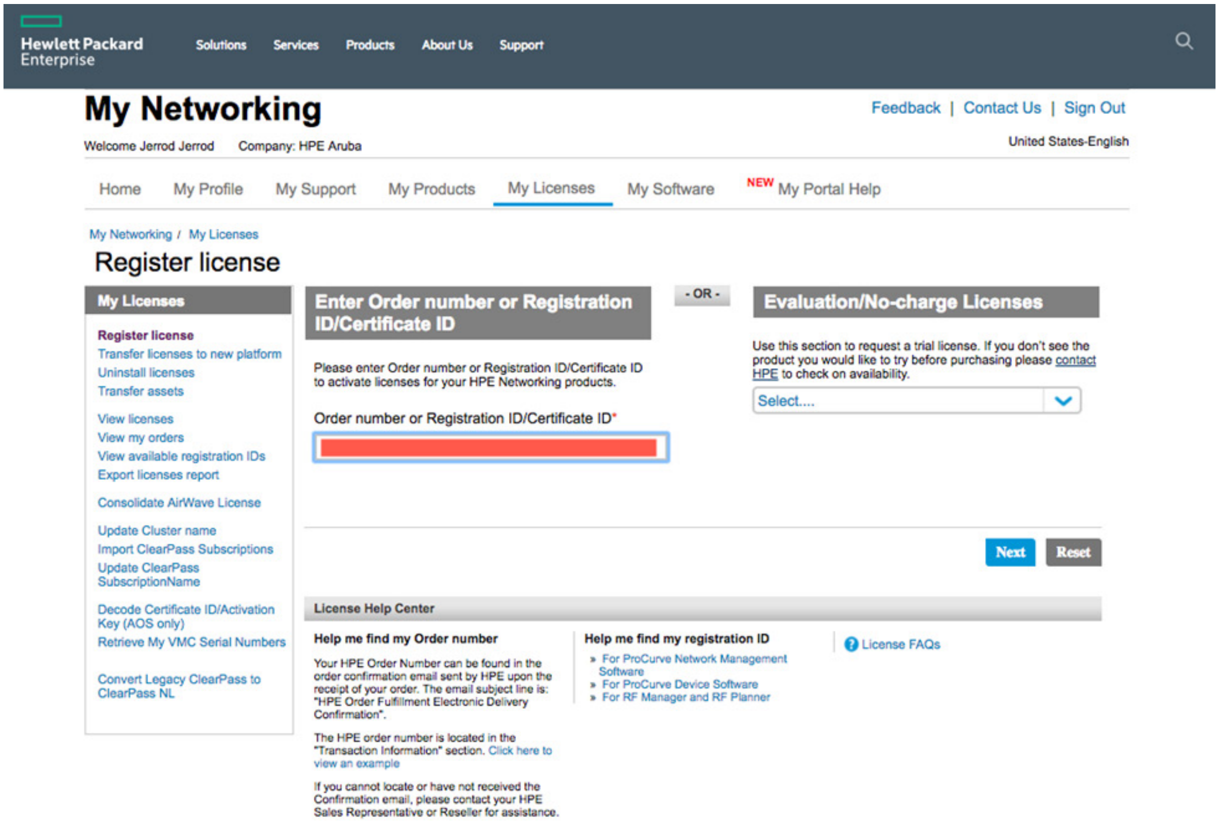

Hewlett Packard Enterprise
Solutions
Services
Products
About Us
Support
My Networking
Feedback | Contact Us | Sign Out
Welcome Jerrod Jerrod
Company: HPE Aruba
United States-English
Home
My Profile
My Support
My Products
My Licenses
My Software
NEW My Portal Help
My Networking / My Licenses
Register license
My Licenses
Register license
Transfer licenses to new platform
Uninstall licenses
Transfer assets
View licenses
View my orders
View available registration IDs
Export licenses report
Consolidate AirWave License
Update Cluster name
Import ClearPass Subscriptions
Update ClearPass SubscriptionName
Decode Certificate ID/Activation Key (AOS only)
Retrieve My VMC Serial Numbers
Convert Legacy ClearPass to ClearPass NL
Enter Order number or Registration ID/Certificate ID
- OR -
Evaluation/No-charge Licenses
Select....
Next
Reset
License Help Center
Help me find my Order number
Help me find my registration ID
» For ProCurve Network Management Software
» For ProCurve Device Software
» For RF Manager and RF Planner
License FAQs

Please enter Order number or Registration ID/Certificate ID to activate licenses for your HPE Networking products.

Order number or Registration ID/Certificate ID*

[redacted]

Use this section to request a trial license. If you don't see the product you would like to try before purchasing please contact HPE to check on availability.

Your HPE Order Number can be found in the order confirmation email sent by HPE upon the receipt of your order. The email subject line is: "HPE Order Fulfillment Electronic Delivery Confirmation".

The HPE order number is located in the "Transaction Information" section. Click here to view an example

If you cannot locate or have not received the Confirmation email, please contact your HPE Sales Representative or Reseller for assistance.

Figure 94 *License Registration in MNP*

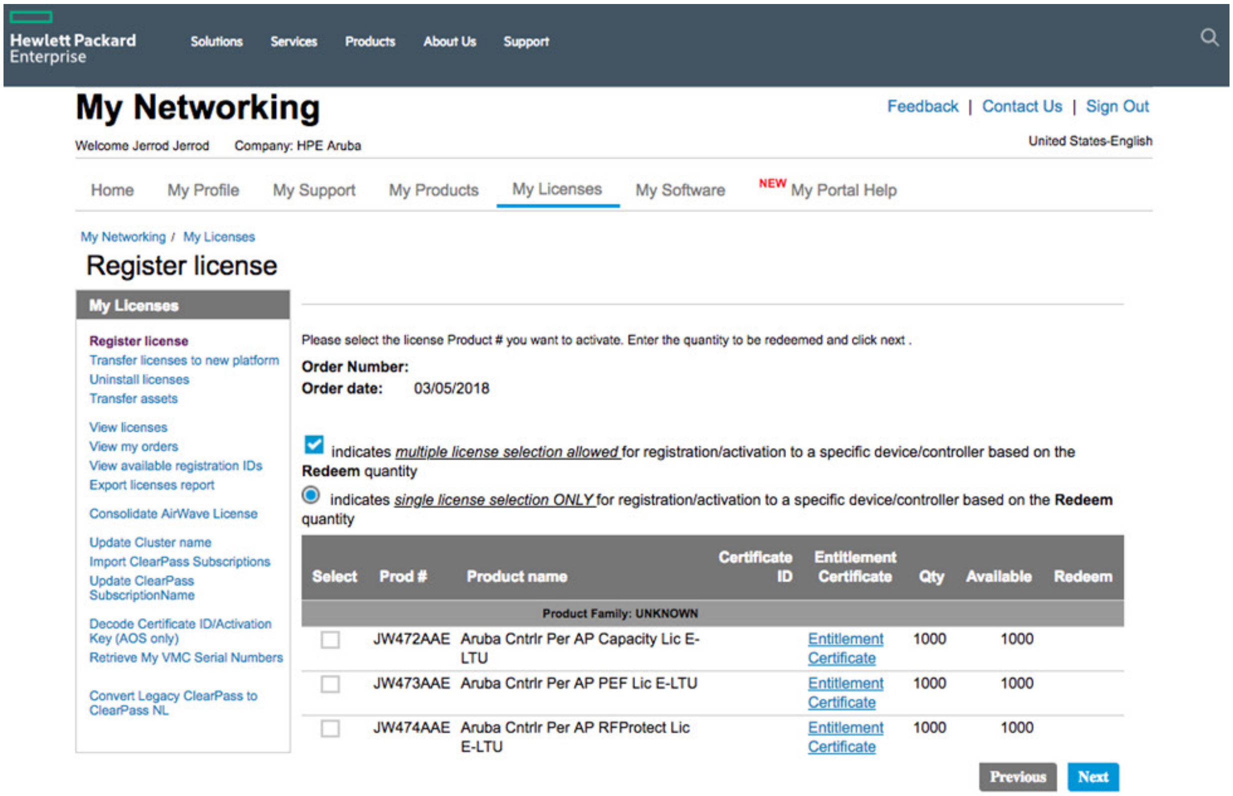

Hewlett Packard Enterprise
Solutions
Services
Products
About Us
Support
My Networking
Feedback | Contact Us | Sign Out
Welcome Jerrod Jerrod
Company: HPE Aruba
United States-English
Home
My Profile
My Support
My Products
My Licenses
My Software
NEW My Portal Help
My Networking / My Licenses
Register license
My Licenses
Register license
Transfer licenses to new platform
Uninstall licenses
Transfer assets
View licenses
View my orders
View available registration IDs
Export licenses report
Consolidate AirWave License
Update Cluster name
Import ClearPass Subscriptions
Update ClearPass SubscriptionName
Decode Certificate ID/Activation Key (AOS only)
Retrieve My VMC Serial Numbers
Convert Legacy ClearPass to ClearPass NL
Previous
Next

Please select the license Product # you want to activate. Enter the quantity to be redeemed and click next .

Order Number:
Order date: 03/05/2018

☑ indicates *multiple license selection allowed* for registration/activation to a specific device/controller based on the **Redeem** quantity

● indicates *single license selection ONLY* for registration/activation to a specific device/controller based on the **Redeem** quantity

Select	Prod #	Product name	Certificate ID	Entitlement Certificate	Qty	Available	Redeem
		Product Family: UNKNOWN					
□	JW472AAE	Aruba Cntrlr Per AP Capacity Lic E-LTU		Entitlement Certificate	1000	1000	
□	JW473AAE	Aruba Cntrlr Per AP PEF Lic E-LTU		Entitlement Certificate	1000	1000	
□	JW474AAE	Aruba Cntrlr Per AP RFProtect Lic E-LTU		Entitlement Certificate	1000	1000	

Figure 95 *License Registration in MNP*

All that is required for license activation is the controller serial number in the case of hardware MCs and MMs or the license passphrase in the case of any virtual appliance or controller. This serial number or license passphrase can be found in the GUI of the controller under **System>Licensing>MM/Controller Licensing** and clicking the blue + sign, or from the CLI with the commands `show inventory` for hardware serials or `show licenses passphrase` for virtual appliances or controllers.

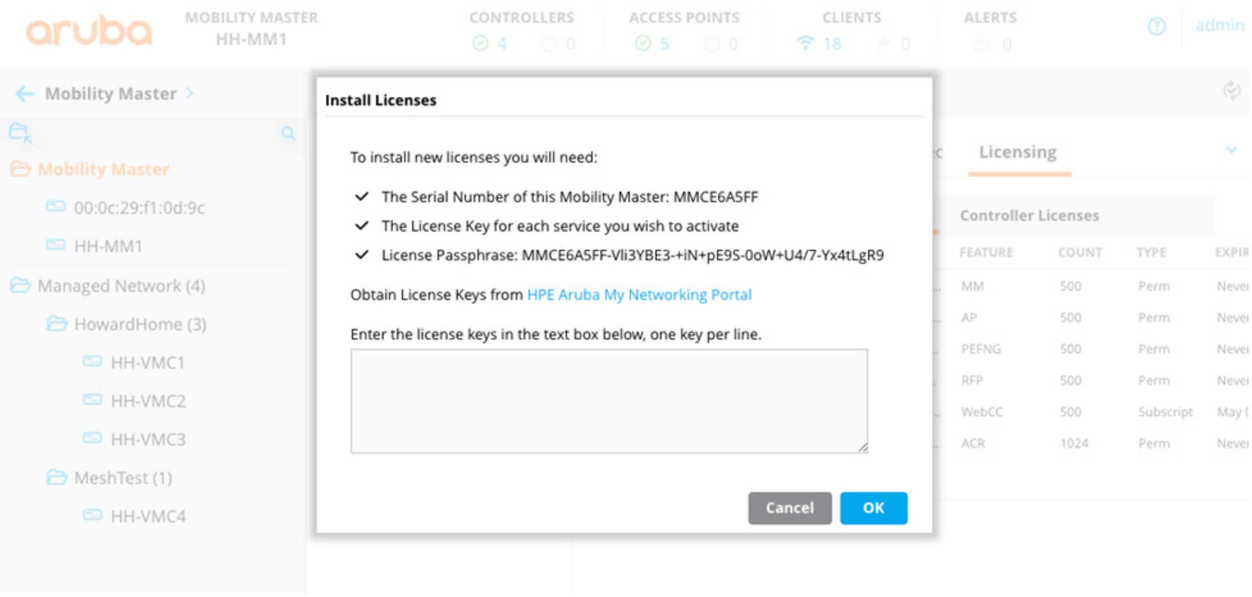

Install Licenses

To install new licenses you will need:

- ✓ The Serial Number of this Mobility Master: MMCE6A5FF
- ✓ The License Key for each service you wish to activate
- ✓ License Passphrase: MMCE6A5FF-Vli3YBE3-+iN+pE9S-0oW+U4/7-Yx4tLgR9

Obtain License Keys from HPE Aruba My Networking Portal

Enter the license keys in the text box below, one key per line.

FEATURE	COUNT	TYPE	EXPIR
MM	500	Perm	Neve
AP	500	Perm	Neve
PEFNG	500	Perm	Neve
RFP	500	Perm	Neve
WebCC	500	Subscript	May (
ACR	1024	Perm	Neve

Figure 96 *License Activation in MNP*

To migrate a license from ArubaOS 6 to ArubaOS 8 though MNP:

1. Navigate to the MNP portal page and select **Transfer Licenses to New Platform**. MNP will display all the serial numbers owned by that account
2. Locate of the serial number of the device currently holding the licenses and click the **Next**
3. Select the new controller where the licenses will migrate from the dropdown along with the new ArubaOS 8 serial or passphrase
4. Click **Transfer**

MNP will provide new license keys and update the serial number database within MNP with the new license activation keys. These new licenses keys would be pasted in to the new MM, MCM, or Standalone 8.x controller.

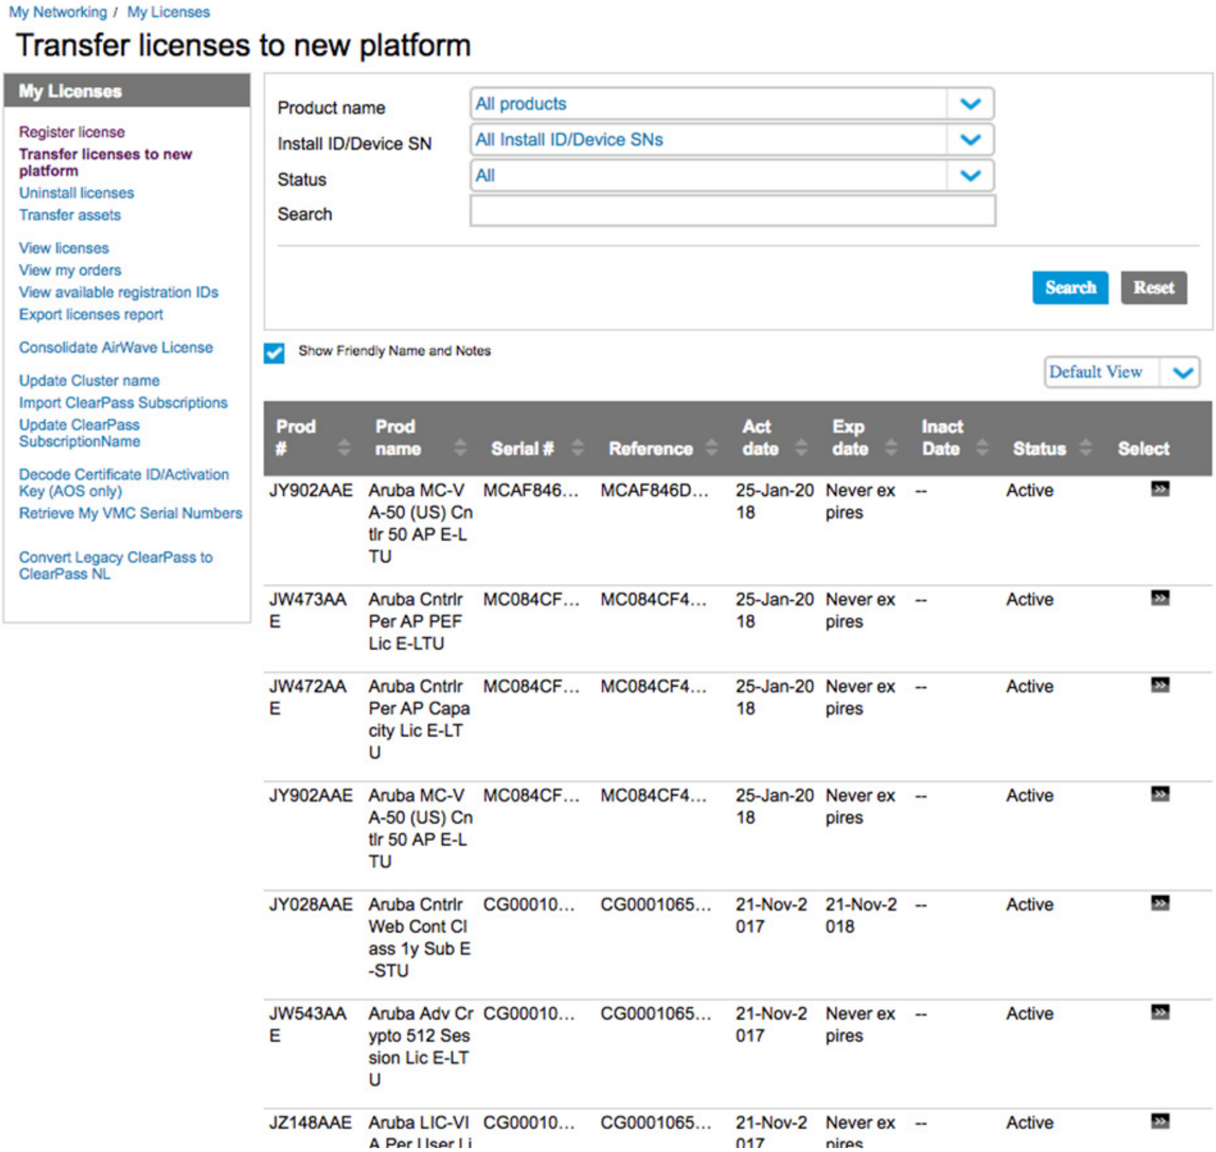

My Networking / My Licenses

Transfer licenses to new platform

My Licenses

- Register license
- Transfer licenses to new platform
- Uninstall licenses
- Transfer assets
- View licenses
- View my orders
- View available registration IDs
- Export licenses report
- Consolidate AirWave License
- Update Cluster name
- Import ClearPass Subscriptions
- Update ClearPass SubscriptionName
- Decode Certificate ID/Activation Key (AOS only)
- Retrieve My VMC Serial Numbers
- Convert Legacy ClearPass to ClearPass NL

Product name: All products

Install ID/Device SN: All Install ID/Device SNs

Status: All

Search:

☑ Show Friendly Name and Notes

Prod #	Prod name	Serial #	Reference	Act date	Exp date	Inact Date	Status	Select
JY902AAE	Aruba MC-VA-50 (US) Cntlr 50 AP E-LTU	MCAF846…	MCAF846D…	25-Jan-2018	Never expires	--	Active	»
JW473AAE	Aruba Cntrlr Per AP PEF Lic E-LTU	MC084CF…	MC084CF4…	25-Jan-2018	Never expires	--	Active	»
JW472AAE	Aruba Cntrlr Per AP Capacity Lic E-LTU	MC084CF…	MC084CF4…	25-Jan-2018	Never expires	--	Active	»
JY902AAE	Aruba MC-VA-50 (US) Cntlr 50 AP E-LTU	MC084CF…	MC084CF4…	25-Jan-2018	Never expires	--	Active	»
JY028AAE	Aruba Cntrlr Web Cont Class 1y Sub E-STU	CG00010…	CG0001065…	21-Nov-2017	21-Nov-2018	--	Active	»
JW543AAE	Aruba Adv Crypto 512 Session Lic E-LTU	CG00010…	CG0001065…	21-Nov-2017	Never expires	--	Active	»
JZ148AAE	Aruba LIC-VIA Per User Li	CG00010…	CG0001065…	21-Nov-2017	Never expires	--	Active	»

Figure 97 *Transferring Licenses to a New Platform in MNP*

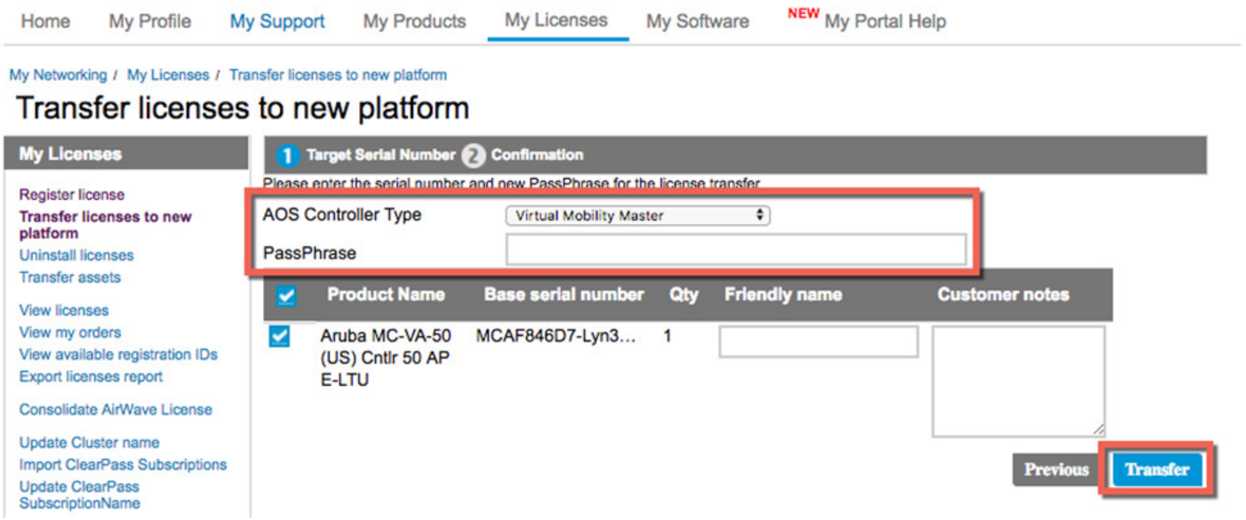

Transfer licenses to new platform

Please enter the serial number and new PassPhrase for the license transfer

AOS Controller Type	Virtual Mobility Master
PassPhrase	

☑	Product Name	Base serial number	Qty	Friendly name	Customer notes
☑	Aruba MC-VA-50 (US) Cntlr 50 AP E-LTU	MCAF846D7-Lyn3…	1		

Figure 98 *Transferring Licenses to a New Platform in MNP*

Aruba Support Portal

The Aruba Support Portal (ASP) is Aruba's new licensing support mechanism designed to activate new licenses from purchases, activate demonstration and evaluation licenses, manage licenses (migrate, change ownership, etc.), and access support documentation. ASP is a more streamlined and simplified support portal that is easier to use, more advanced in additional capabilities, and will be Aruba's main licensing support portal moving forward for all Aruba products.

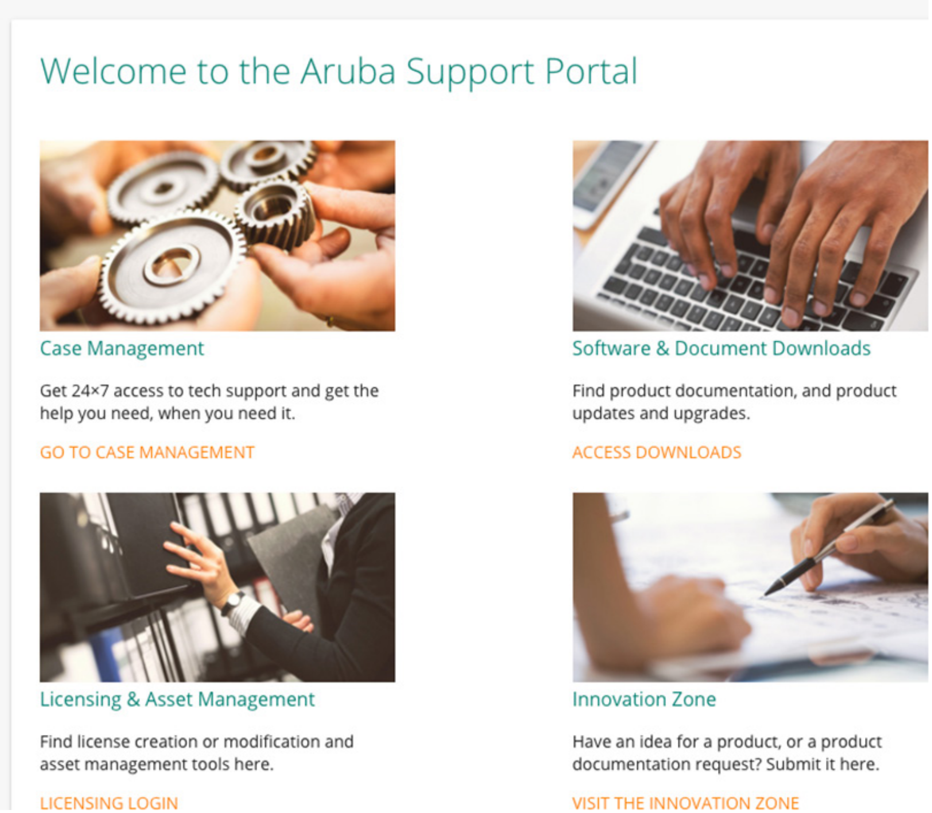

Welcome to the Aruba Support Portal

Case Management

Get 24×7 access to tech support and get the help you need, when you need it.

GO TO CASE MANAGEMENT

Software & Document Downloads

Find product documentation, and product updates and upgrades.

ACCESS DOWNLOADS

Licensing & Asset Management

Find license creation or modification and asset management tools here.

LICENSING LOGIN

Innovation Zone

Have an idea for a product, or a product documentation request? Submit it here.

VISIT THE INNOVATION ZONE

Figure 99 *Aruba Support Portal*

License Activation

ASP works similarly to MNP in the sense that new orders will have an *order number* or *certificate ID* associated with them for license activation. From the **Licensing Login** page, the order number or certificate ID can be entered in to the proper field. Next the controller serial number or passphrase will be entered along with the number of licenses requiring activation. Once all of the requisite information has been entered into the portal clicking the **Activate Certificate** button will generate the new license key. ASP can be accessed through the following URL:

https://asp.arubanetworks.com

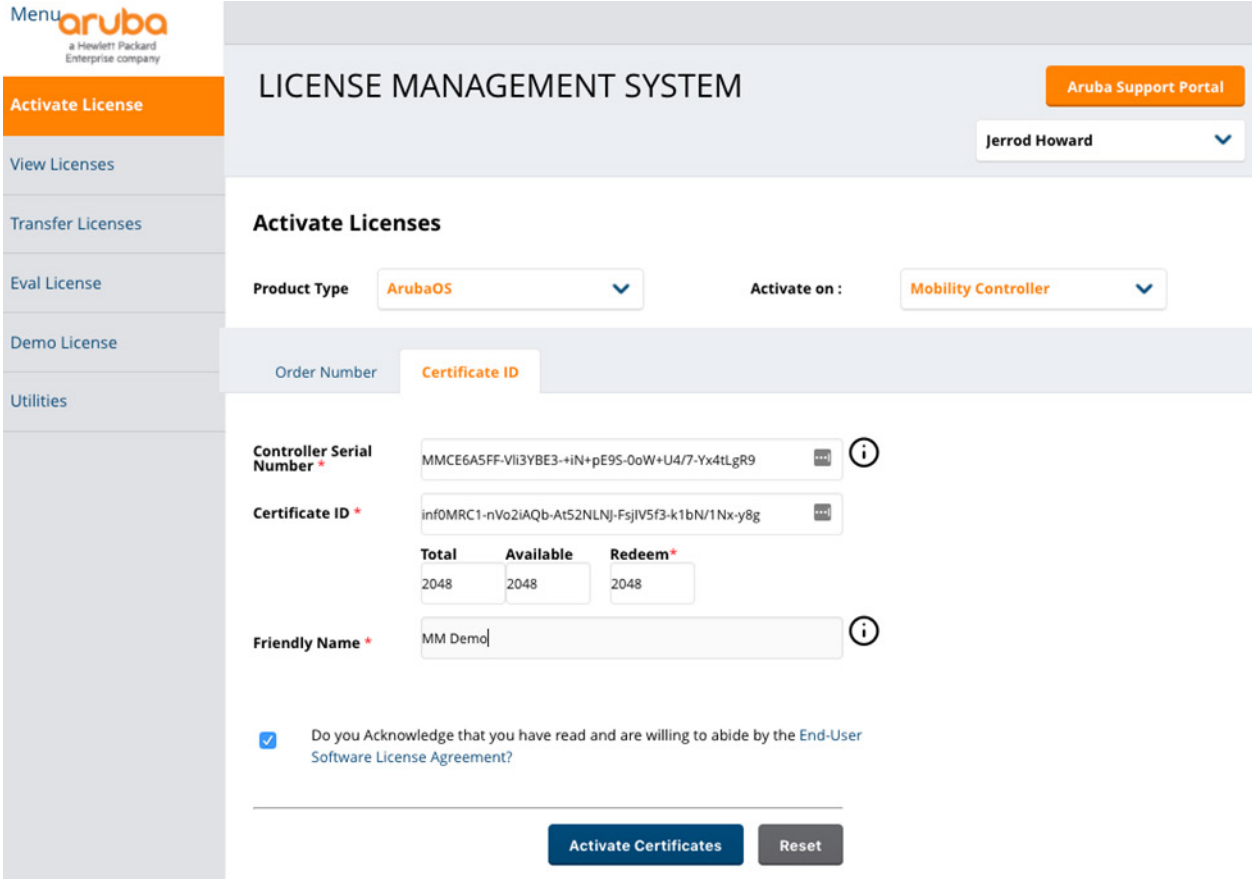

Activate Licenses

Product Type: ArubaOS

Activate on : Mobility Controller

Controller Serial Number *: MMCE6A5FF-Vli3YBE3-+iN+pE9S-0oW+U4/7-Yx4tLgR9

Certificate ID *: inf0MRC1-nVo2iAQb-At52NLNJ-FsjIV5f3-k1bN/1Nx-y8g

Total	Available	Redeem*
2048	2048	2048

Friendly Name *: MM Demo

☑ Do you Acknowledge that you have read and are willing to abide by the End-User Software License Agreement?

Figure 100 *License Activation in ASP*

License Migration

ASP can also be used to migrate licenses between devices. The portal contains a section dedicated to license migration called **Transfer Licenses** which will reveal the different types of devices which currently hold licenses such as MCs, Hardware MMs, Virtual MMs, etc. After navigating to this section, a page will be displayed for entering the serial of the device where the licenses that will be migrated currently reside. After the serial has been entered another window will appear with a prompt asking for the destination for the newly migrated licenses. Once all requisite information has been entered into the tool, ASP will update its database accordingly to display the new ArubaOS 8 license keys associated with the serial number or passphrase of the new device.

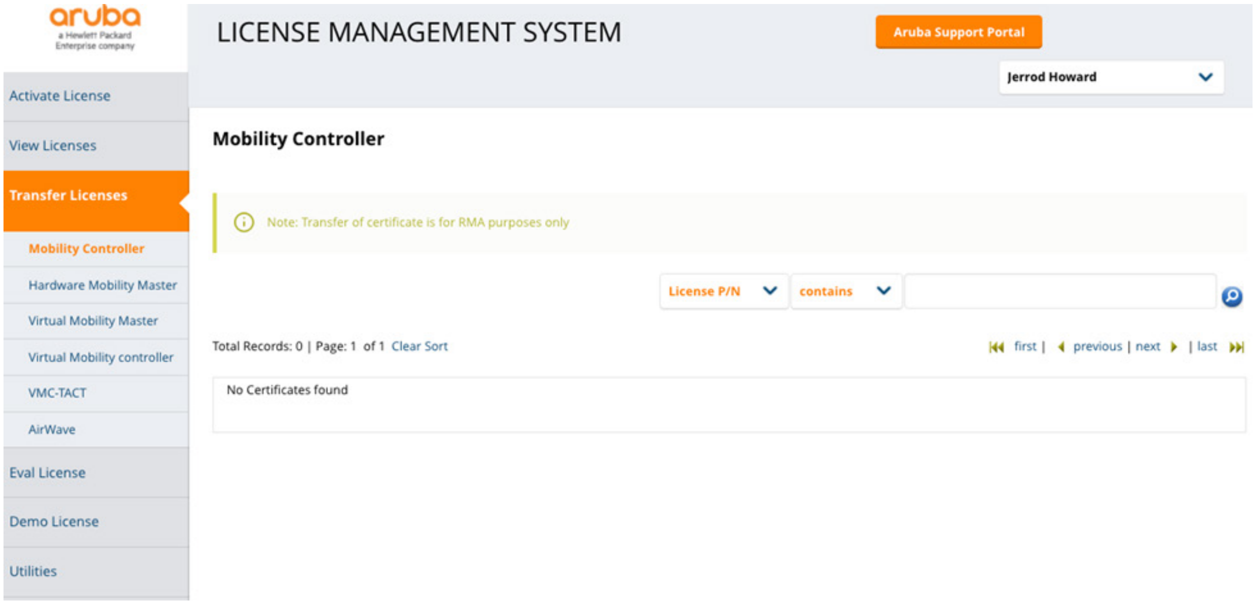

Mobility Controller

Note: Transfer of certificate is for RMA purposes only

Total Records: 0 | Page: 1 of 1 Clear Sort

No Certificates found

Figure 101 *Transferring License in ASP*

License Migration Tool

The ArubaOS 8 Migration Tool has the capability to communicate with MNP or ASP to communicate with controllers and automatically migrate their associated licenses as part of the upgrade process.

It’s important to note that the Migration Tool was designed specifically for scenarios involving an upgrade from ArubaOS 6 to ArubaOS 8. The Migration Tool cannot be used for license migration independent of an upgrade. If the Migration Tool is not being used for an upgrade then the existing licenses must be manually migrated via MNP or ASP, otherwise it would be necessary to purchase new licenses as part of the migration.

During the Migration Tool setup process a prompt will appear asking in the licenses should be migrated. If **yes** is selected, a field will appear requesting an MNP or ASP username. The password will also need to be entered at end of the migration process setup.

Once the serial numbers of the controllers being migrated to ArubaOS 8 and the MM passphrase (if applicable) have been entered into the Migration tool the upgrade process may be initiated. The Migration Tool will log in to MNP or ASP, search for the associated serial numbers, and migrate the licenses accordingly. Either all of the licenses will be migrated to the MM, or they will migrate according to their controller serial numbers.

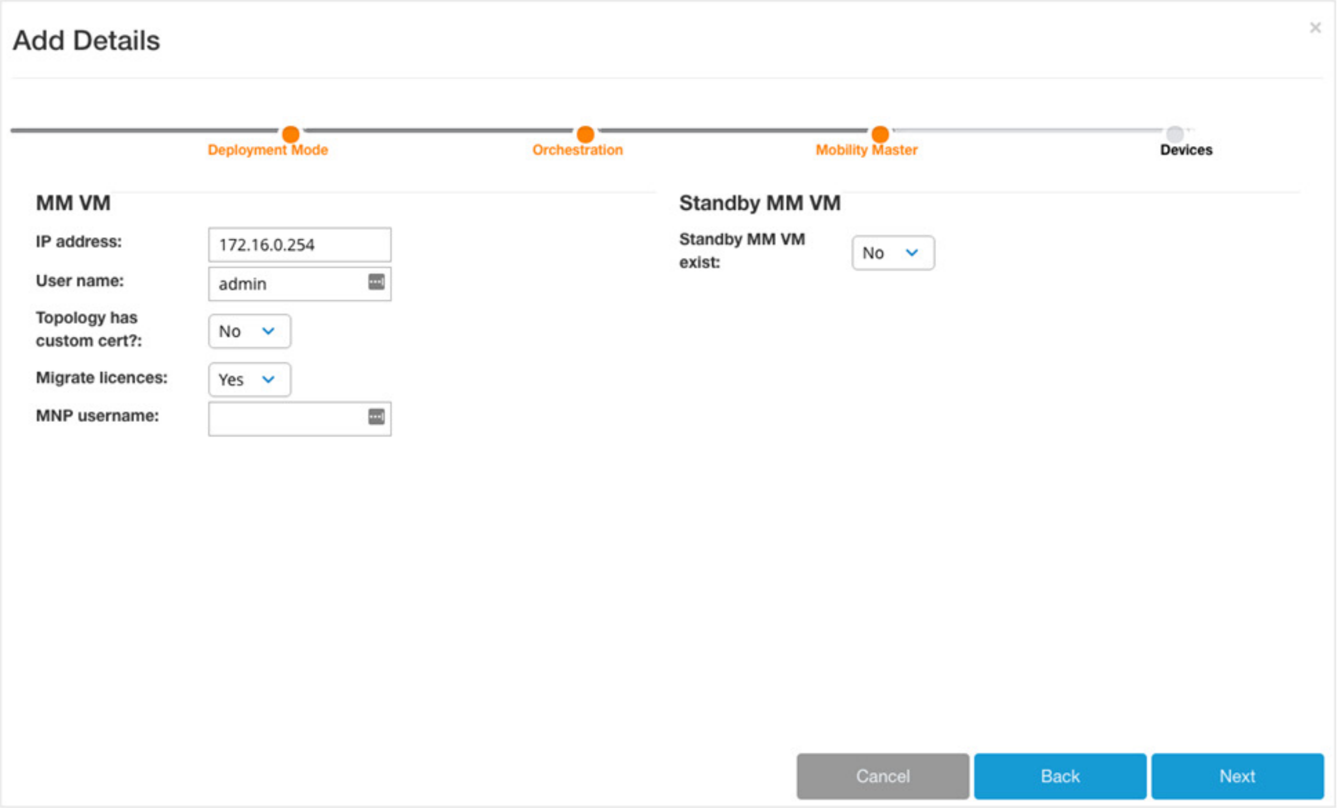
Add Details

MM VM

IP address: 172.16.0.254

User name: admin

Topology has custom cert?: No

Migrate licences: Yes

MNP username:

Standby MM VM

Standby MM VM exist: No

Figure 102 *Adding Details for License Migration in ASP*

License Installation

License Components

In order to install licenses on a controller or MM the *license key* generated from the MNP or ASP activation is required. The key will generate an alphanumeric string which enables a specific license type, feature, or subscription. Licenses can be installed on a MM (MM), MM Controller (MCM), or Standalone MC (MC) through the GUI or the CLI.

License Installation via GUI

The MM serves as the master licensing server for all controllers and services. Licenses for the devices under a particular MM are installed at the MM level. To install licenses on an MM, access the GUI and navigate to **MM>System>Licensing tab**, click **MM Licenses**, and then click the blue **+** sign to add a new license.

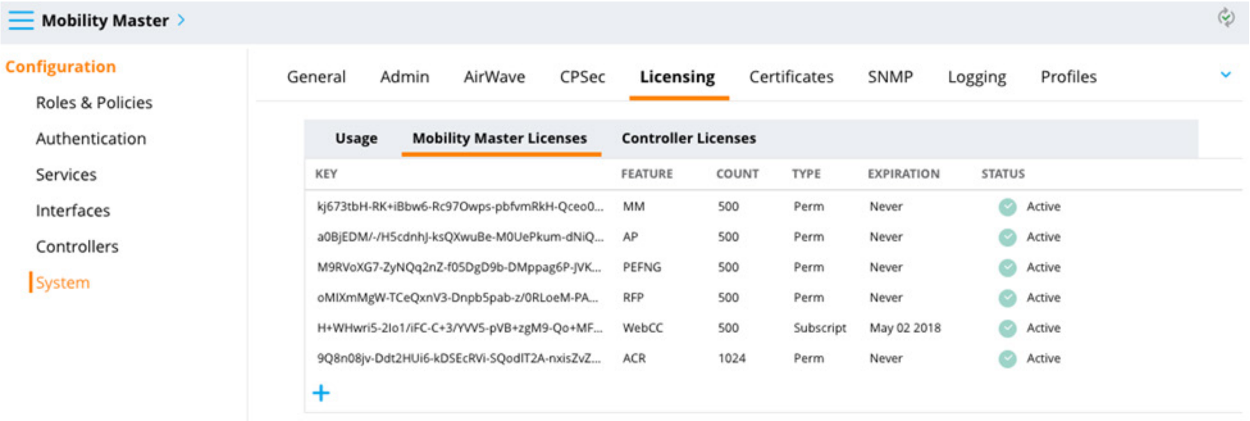

KEY	FEATURE	COUNT	TYPE	EXPIRATION	STATUS
kj673tbH-RK+iBbw6-Rc97Owps-pbfvmRkH-Qceo0...	MM	500	Perm	Never	✓ Active
a0BjEDM/-/H5cdnhJ-ksQXwuBe-M0UePkum-dNiQ...	AP	500	Perm	Never	✓ Active
M9RVoXG7-ZyNQq2nZ-f05DgD9b-DMppag6P-JVK...	PEFNG	500	Perm	Never	✓ Active
oMIXmMgW-TCeQxnV3-Dnpb5pab-z/0RLoeM-PA...	RFP	500	Perm	Never	✓ Active
H+WHwri5-2Io1/iFC-C+3/YVV5-pVB+zgM9-Qo+MF...	WebCC	500	Subscript	May 02 2018	✓ Active
9Q8n08jv-Ddt2HUi6-kDSEcRVi-SQodIT2A-nxisZvZ...	ACR	1024	Perm	Never	✓ Active

Figure 103 *MM License Installation Using the GUI*

All installed licenses as well as their quantities can be seen from this page. In addition, licenses can be deleted if necessary. The MM licensing services are capable of managing all licenses including dynamic provisioning of MC-VA licensing. This functionality is unique to the MM and allows for a single installed MC-VA license pool to be split up between multiple VMCs under that MM on demand. E.g., a single MC-VA-250 license pool could be spread out over 2 MCs or 20 VMCs under the same MM so long as the total number of APs terminated across all controllers does not exceed 250 APs. License allocation under a MM is completely flexible and can be easily adapted according the needs of a given deployment,

The license installation location and process in the GUI is the same for MCM and Standalone controllers as it is for MCs and VMCs. The main difference is that instead of selecting the MM section **MC>Configuration>System>Licensing** should be selected. Next, choose the **Inventory** drop down screen and click on the blue + sign to add a new license.

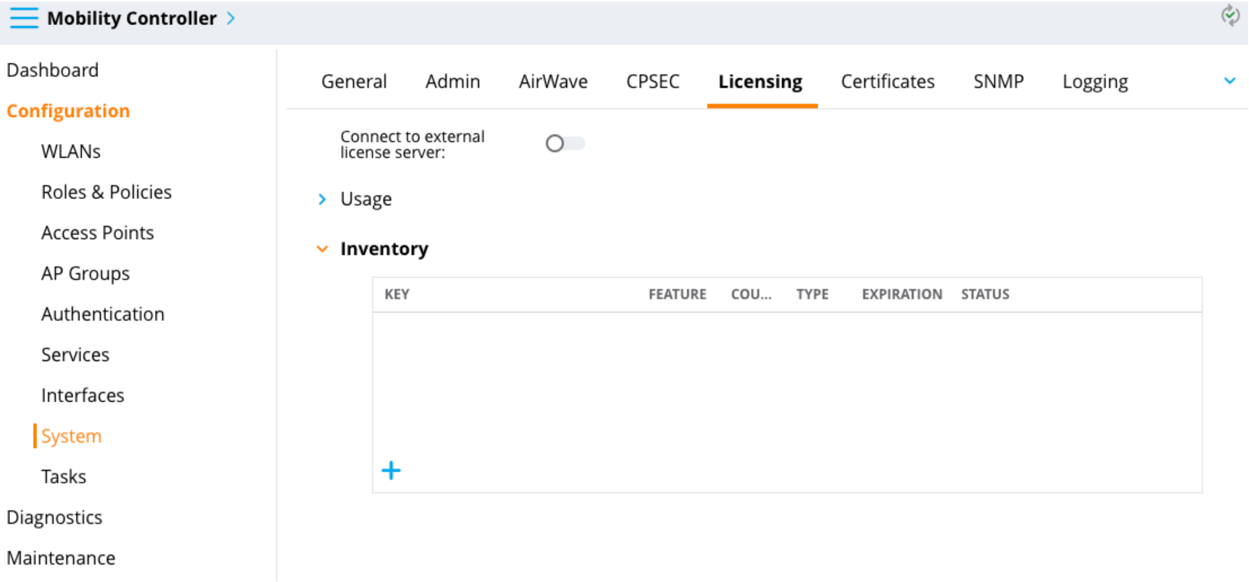

KEY	FEATURE	COU...	TYPE	EXPIRATION	STATUS

Figure 104 *MCM and Standalone License Installation Using the GUI*

In both cases, a prompt will appear with the serial number, license passphrase (in the case of a virtual appliance or VMC), and a window to paste one or more licenses to apply to the desired controller.

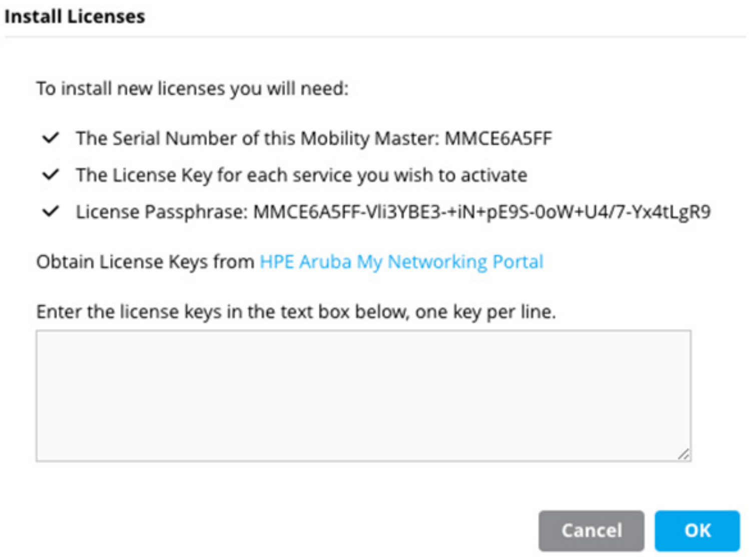
Install Licenses

To install new licenses you will need:

- ✓ The Serial Number of this Mobility Master: MMCE6A5FF
- ✓ The License Key for each service you wish to activate
- ✓ License Passphrase: MMCE6A5FF-Vli3YBE3-+iN+pE9S-0oW+U4/7-Yx4tLgR9

Obtain License Keys from HPE Aruba My Networking Portal

Enter the license keys in the text box below, one key per line.

Figure 105 *License Installation Dialogue Window*

Once the license key(s) have been added, select the blue **OK** button to complete the installation process. In most cases, if the controllers are having licenses installed on them for the first time they will need to be rebooted in order for the licenses to take effect. A reboot is not typically required for future license additions where additional quantities of existing licenses are being installed. Generally the page indicate whether or not the controllers require a reboot. A license flagged with an “R” designation means that a reboot is required.

License Installation via CLI

Installing licenses via the CLI can be accomplished by logging in to the MM, MCM, or Standalone controller via SSH and entering the following command:

```
#License add <insert-license-key-here>
```

License inventory can be displayed with the following command:

```
#show licensing
```

Mobility Master License Pools

Another feature that is unique to the MM (MM) is the ability to create *license pools*. *License pools* allow the MM to manage smaller pools of the larger global license pool as well as apply licenses to controllers or groups of controllers under that MM. Allocating licenses in this manner allows for more granular control over how many licenses each controller or group of controllers are allowed to consume. Use cases for this feature include situations where the global network administrator wants to limit how many APs a location can deploy so that their license supply is not exhausted in order to avoid creating issues at other sites.

License Pools are not available in the MCM and Standalone Controller Modes.

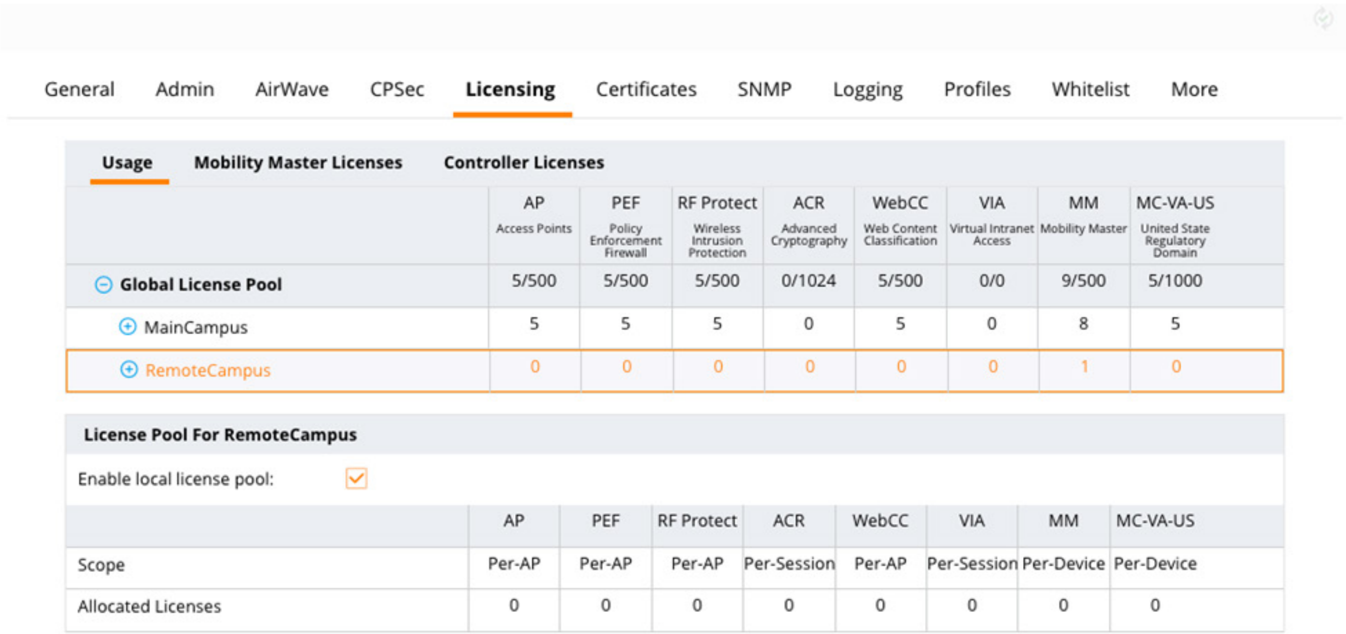

	AP Access Points	PEF Policy Enforcement Firewall	RF Protect Wireless Intrusion Protection	ACR Advanced Cryptography	WebCC Web Content Classification	VIA Virtual Intranet Access	MM Mobility Master	MC-VA-US United State Regulatory Domain
⊖ Global License Pool	5/500	5/500	5/500	0/1024	5/500	0/0	9/500	5/1000
⊕ MainCampus	5	5	5	0	5	0	8	5
⊕ RemoteCampus	0	0	0	0	0	0	1	0

License Pool For RemoteCampus

Enable local license pool: ☑

	AP	PEF	RF Protect	ACR	WebCC	VIA	MM	MC-VA-US
Scope	Per-AP	Per-AP	Per-AP	Per-Session	Per-AP	Per-Session	Per-Device	Per-Device
Allocated Licenses	0	0	0	0	0	0	0	0

Figure 106 *Creating License Pools*

License pools are defined by the nodes of the MM. In the example below, there is a “Main Campus” as well as a “Remote Campus” group of controllers. The global license pool consists of 500 licenses for AP, PEF, RF Protect, etc. The network administrators for this example scenario want to create a limit of 50 APs at the Remote Campus.

To set up the desire limit, they would click on the desired node and check the **Enable Local License Pool** box. Next, they will click on each license type to assign a quantity of licenses that are allowed for the Remote Campus node. This process will need to be repeated for each license type required for the node. In this case there are 50 APs requiring licenses, meaning that 50 licenses of each type will be being deducted from the Global License Pool. In our example, the “RemoteCampus” node cannot use more than 50 APs of licensing from the MM global pool. The remaining 450 licenses are available in the Global License Pool for allocation to other nodes.

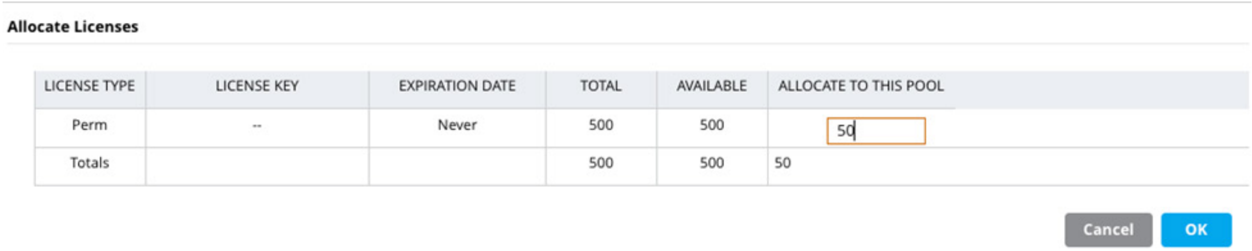

Allocate Licenses

LICENSE TYPE	LICENSE KEY	EXPIRATION DATE	TOTAL	AVAILABLE	ALLOCATE TO THIS POOL
Perm	--	Never	500	500	50
Totals			500	500	50

Figure 107 *Allocating Licenses to Pools*

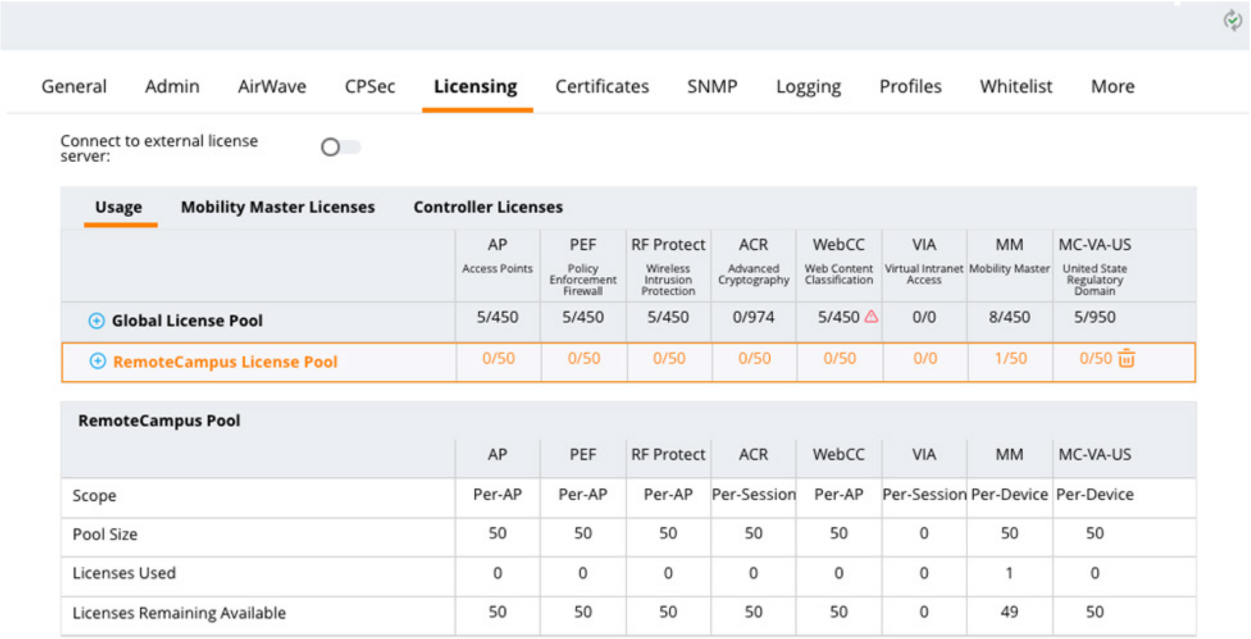

Connect to external license server: ○

	AP Access Points	PEF Policy Enforcement Firewall	RF Protect Wireless Intrusion Protection	ACR Advanced Cryptography	WebCC Web Content Classification	VIA Virtual Intranet Access	MM Mobility Master	MC-VA-US United State Regulatory Domain
⊕ Global License Pool	5/450	5/450	5/450	0/974	5/450 ⚠	0/0	8/450	5/950
⊕ RemoteCampus License Pool	0/50	0/50	0/50	0/50	0/50	0/0	1/50	0/50

RemoteCampus Pool

	AP	PEF	RF Protect	ACR	WebCC	VIA	MM	MC-VA-US
Scope	Per-AP	Per-AP	Per-AP	Per-Session	Per-AP	Per-Session	Per-Device	Per-Device
Pool Size	50	50	50	50	50	0	50	50
Licenses Used	0	0	0	0	0	0	1	0
Licenses Remaining Available	50	50	50	50	50	0	49	50

Figure 108 *License Pool Dashboard*

Controller Reference Architectures

Introduction

This chapter addresses the design decisions and best practice options for implementing an end-to-end Aruba mobile first architecture for a typical enterprise network. This chapter focuses on architecture design recommendations and explains the various configurations and considerations that are needed to build each architecture. Reference architectures are provided for small, medium, and large buildings as well as large campuses. For each architectural model the following topics are discussed:

- Recommended modular local area network (LAN) designs
- MC cluster placement
- Design considerations and best practices
- Suggested switch and wireless platforms

The information provided in this chapter is useful for network architects responsible for greenfield designs, network administrators responsible for optimizing existing networks, and network planners requiring a template that can be followed as their network grows. The scope of this chapter applies to environments from small offices with fewer than 32 APs up to large campuses supporting up to 10,000 APs.

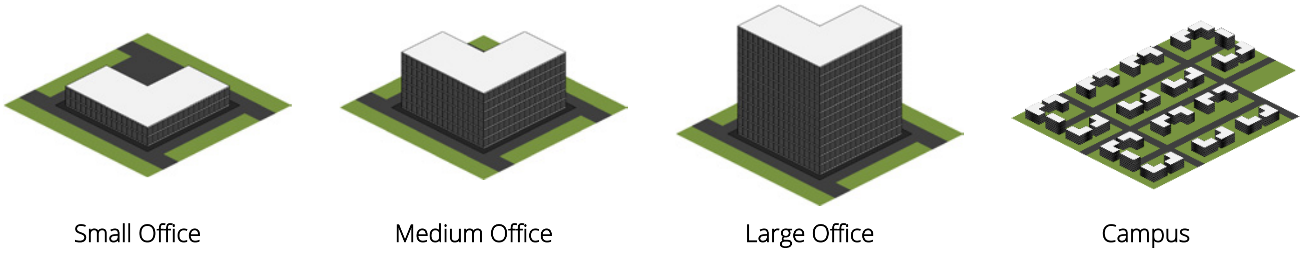

Figure 109 *Scope of Designs*

NOTE

This chapter does not provide step-by-step configuration examples or vertical specific wireless designs. Detailed configuration examples are provided by Aruba Solutions Exchange (ASE) while vertical specific designs are provided in separate VRD documents.

Design Priciples

The foundation of each reference architecture provided in this chapter is the underlying modular LAN design model that separates the network into smaller more manageable modular components. A typical LAN consists of a set of common interconnected layers such as the core, distribution, and access layers which form the main network along with additional modules that provide specific functions such as Internet, WAN, Wireless, and server aggregation.

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

1. Modules can be easily replicated which allows for growth and scalability
2. Modules can be added and removed with minimum impact to other layers as network requirements evolve
3. Modules allow the impact of operational changes to be constrained to a smaller subset of the network
4. Modules compartmentalize the netowrk into specific fault domains providing fault tolerance

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

Modular Designs

The modular design that is selected for a specific LAN deployment is dependent on numerous factors. As an industry best practice, networks are built using either a 2-tier or 3-tier modular LAN design which differ from each other by the inclusion or exclusion of a distribution layer. The additional distribution layer resides between the core and access layers to provide aggregation and routing:

- **2-tier Modular LAN** – Collapses the core and distribution layers into a single layer. The switches in the core/distribution layer perform a dual role by providing aggregation to the access layer modules as well as performing IP routing functions
- **3-tier Modular LAN** – Utilizes a dedicated distribution layer between the core and access layers. The distribution layer switches provide aggregation to the access layers and are connected directly to the core. Distribution layer switches are commonly deployed in larger networks and are typically used to connect different modules such as wireless, WAN, internet, and servers

NOTE

For purposes of this chapter the terms "aggregation layer" and "distribution layer" are interchangeable.

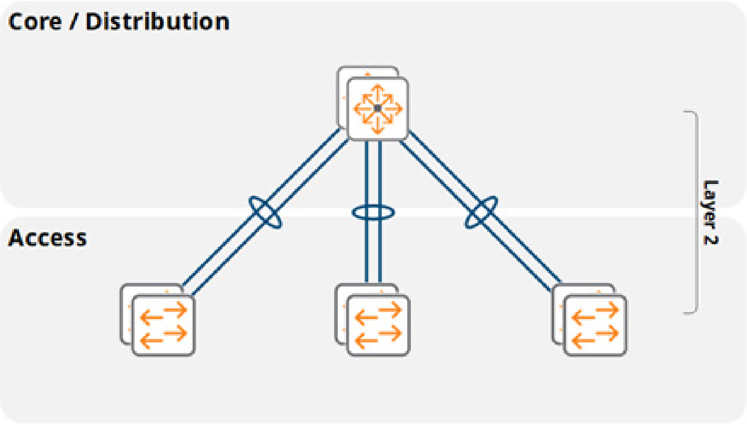

Figure 110 *2-Tier Modular LAN*

A 2-tier modular LAN is well suited for small buildings with few wiring closets and access switches. The access layer VLANs are extended between the access layer switches and the core/distribution layer switches using 802.1Q trunking. The core/distribution switches include IP interfaces for each VLAN and operate as the default gateway for the access layer hosts.

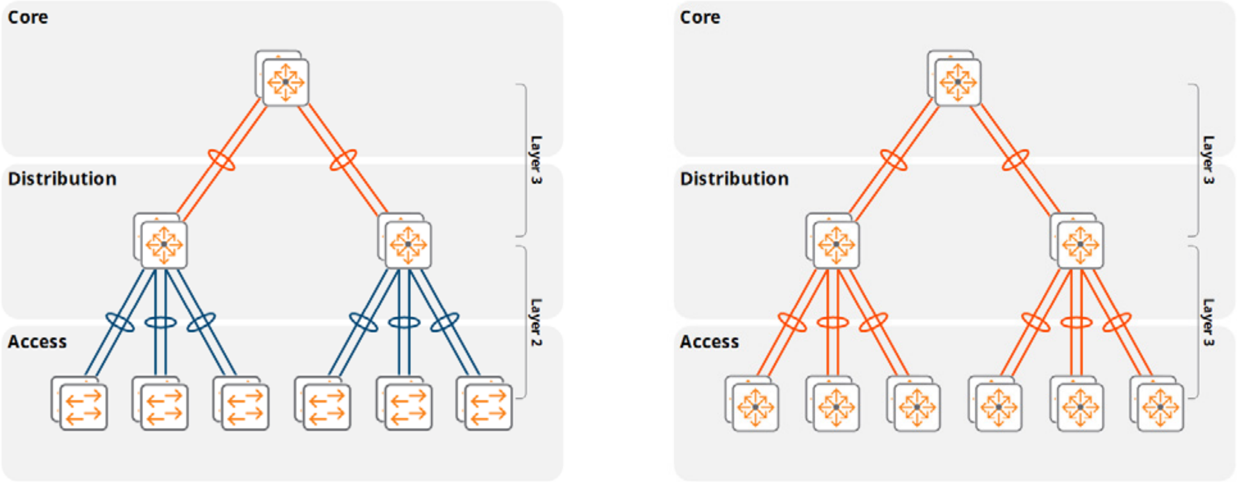

Figure 111 *Layer 2 and Routed Access 3-tier LAN Designs*

In a 3-tier modular design, the IP routing functions are distributed between the core and aggregation layers. Depending on the design they may be extended to the access layers as well. All 3-tier LAN designs will implement IP routing between the aggregation and core switches using a dynamic routing protocol such as Open Shortest Path First (OSPF) for reachability and address summarization:

- **Layer 2 access layer** – All VLANs from the access layer are extended to the aggregation layer switches using 802.1Q trunking. The aggregation switches provide layer 3 interfaces (VLAN interfaces or SVIs) for each VLAN and provide reachability to the rest of the IP network

- **Routed access layer** – IP routing is performed between the aggregation and access layer switches (as well as between the aggregation and core layers). In this deployment model, each access layer switch or stack provides reachability to the rest of the IP network

The Aruba Mobile First Architecture supports both designs allowing our customers to leverage the benefits of Aruba solutions with either network design.

LAN Aggregation Layer

When designing and planning a mobile first network the decision of whether or not to deploy an aggregation layer hinges on several key factors:

1. The number of access layer switches that need to be connected. Eventually the number of SFP/ SFP+/QSFP ports required to connect the access layer will exceed the physical port capacity of the core switches. Adding an additional layer between the core and access layers provides aggregation reducing the number physical ports required in the core
2. The structured wiring design of the building. Intermediate distribution frames (IDFs) in larger buildings typically connect to main distribution frames (MDFs) via fiber at strategic locations within the building. Each MDF typically connects to a server room or data center.
   - Aggregation switches are often required in MDFs due to limited fiber capacity between the MDFs and main server room or data center
   - When multi-mode fiber is deployed aggregation switches allow the IDFs to be connected when the combined fiber lengths (IDF+MDF+server room) exceed the distance specifications for fiber optic connections
3. MDFs provide ideal locations for aggregation layer switches as they typically aggregate the fiber connections from the access layer and provide connectivity to the core deployed in the main server room or data center
4. Network manageability, stability, and scalability concerns dictate that specific fault domains should be introduced into the network. This is typically achieved by implementing IP routing between the core and respective aggregation layers. Designing a network in this manner ensures that the core is isolated from layer 2 faults or operational changes originating from other layers or modules
5. Reducing layer 2 and layer 3 processing load on the core. As a network grows the MAC address table sizes and IP protocol processing overhead increases proportionately. The inclusion of an aggregation layer offloads the layer 2 learning and IP protocol processing overhead from the core to the respective aggregation layer switches. The aggregation layer then becomes the layer 2 and layer 3 demarcation points for the clients allowing the core to be dedicated to IP routing functions

Wireless Module Aggregation Layer

A dedicated aggregation layer will typically be introduced for the wireless module once the number of wireless and dynamically segmented client host addresses exceeds a specific threshold. As wireless and dynamically segmented client traffic is tunneled from the APs and access layer switches to the MC cluster, the MAC learning and IP processing overhead is incurred by the first hop router for those VLANs. In a 2-tier modular network design this overhead is incurred by the aggregation and core layer devices. In a 3-tier modular network design the overhead is incurred by the core. The addition of a dedicated wireless module aggregation layer alleviates the MAC learning and IP processing overhead from the core and shifts it to a dedicated wireless aggregation layer providing stability, fault isolation, and scalability.

As a general best practice Aruba recommends implementing a dedicated wireless aggregation layer when the total number of IPv4 and IPv6 addresses from both wireless and dynamically segmented clients exceeds 4,096. Doing so safeguards the network from future growth and ensures the core layer is not overwhelmed as new classes of devices are added to the network or IPv6 is introduced which can double or triple number of host IP addresses.

The limit to wireless module scalability is dependent on the scaling capabilities of the aggregation layer switches deployed for the wireless module. Switches are designed to support a specific number of hosts which includes the necessary table sizes and processing power to perform layer 2 (MAC) and layer 3 (ARP and neighbor discovery) learning and table maintenance.

The latest generation of Ethernet switches from Aruba can comfortably scale to support 64,000 addresses (both IPv4 and IPv6) before the switches become overwhelmed. As a best practice networks should be designed so that the total number of host addresses per wireless module does not exceed the capacity of the wireless module aggregation switches. Larger wireless networks that grow beyond 64,000 IPv4 or IPv6 host addresses require additional wireless modules consisting of an aggregation layer and MC cluster.

Switch Series	Maximum IPv4 Addresses	Maximum IPv6 Addresses	Maximum IPv4 and IPv6 Addresses
Aruba 3810 Series	25,000	12,500	12,500
Aruba 5400R Series	25,000	12,500	12,500
Aruba 8320 Series	14,000	14,000	14,000
Aruba 8400 Series	64,000	64,000	128,000

Table 21 *Aruba Switch Series Capacities*

Scaling requirements for native IPv4 deployments can be easily calculated as each host is assigned a single IPv4 address. A wireless module using Aruba 8400 series aggregation switches can comfortably scale to support up to 64,000 IPv4 hosts before exceeding the capacity of the aggregation layer switches. The number of hosts than can be supported for a dual stack (IPv4 and IPv6) or native IPv6 deployments is more challenging to calculate as each IPv6 host can be assigned multiple IPv6 addresses (link-local address plus one or more global addresses).

Therefore the total number of IPv6 addresses that are assigned per host determines the maximum overall number of hosts that can be supported within each wireless module:

- **Native IPv6** – Assuming each host is assigned one global IPv6 address, each wireless module can support a maximum of 32,000 wireless with dynamically segmented dual stack hosts. Each host will include a mandatory link-local address so the maximum number of IPv6 address the aggregation switch can support will be half of its usual capacity
- **Dual Stack** – Assuming each host is assigned one IPv4 address and two global IPv6 addresses, each wireless module can support a maximum of 21,000 wireless and dynamically segmented hosts. As each host will include a mandatory link-local address, the maximum number of IPv6 address the aggregation switch can support will be a third of its usual capacity

NOTE

Strategies and architectures for scaling beyond 64,000 host addresses are discussed in the campus reference architecture section. An Aruba mobile first architecture can scale to support up to 100,000 clients per MM by implementing multiple MC clusters each with their own aggregation layer.

Wireless Module Redundancy

One important aspect of an Aruba mobile first redundant design is the connectivity of the wireless module that contains the MCs. The cluster of controllers terminates the AP management and control tunnels as well as the wireless and dynamically segmented client tunnels. To enable redundancy each cluster consists of a minimum of two MCs and can potentially scale up to four or twelve cluster members (depending on the controller model). As a best practice, each cluster must contain members of the same model.

Each of the MCs in the cluster connects to a pair of Aruba switches using dynamic port-channels forming a link aggregated connection (LAG). Link Aggregation Control Protocol (LACP) is enabled to verify peer availability and provide layer 2 loop prevention. The MCs are connected to core or wireless aggregation layer switches depending on whether a 2-tier or 3-tier hierarchical network design has been selected for the deployment in addition to the number of wireless and dynamically segmented hosts.

Redundancy within the wireless module is provided at multiple layers:

- **ArubaOS 8 Clustering** – Each AP and client establishes a tunnel to a primary and secondary MC within the cluster. This ensures that a network path is available to APs and clients in the event of a live upgrade or MC outage
- **Device and Link Redundancy** – Each MC is connected to two Aruba switches that support network virtualization functionality (NVF) if they are core switches or LAG if they are wireless aggregation switches. This ensures that a network path is always available to the MCs, APs, and clients in the event of a switch outage or link failure
- **Path Redundancy** – Link Aggregation Control Protocol (LACP) is part of the IEEE 802.3ad standard and ensures that all paths are fully redundant. LACP is an active protocol that

allows switch peers to detect the operational status of peers devices and their connected ports

- **First Hop Router Redundancy** – The network must ensure that packets will continue to be forwarded in the event of a default gateway failure. First hop router redundancy is natively provided by Aruba Switches supporting NVF without the need for implementing first-hop routing redundancy protocols such as VRRP

MC ports in the LAG are distributed between pairs of Aruba switches implementing NVF for all mobile first reference architectures. The model chosen for the switches that the MCs connect to will be depend on the 2-tier or 3-tier hierarchical network design selected for the deployment as well as the number of wireless clients that are supported. Switches supporting the wireless module can be stack of Aruba 3810Ms, a pair of Aruba 5400Rs configured for virtual switching framework (VSF), or a pair of Aruba 8320s/8400s configured for MC-LAG.

Aruba 3810M Switches

The figure below demonstrates how a cluster of MCs is connected to a stack of Aruba 3810M switches deployed in the core, core/aggregation, or wireless aggregation layer. The Aruba stacking architecture virtualizes both the control and data planes allowing the 3810M stack of switches to forward traffic as well as be configured and managed as a single virtual switch.

In this example two or more 1 Gigabit or 10 Gigabit Ethernet ports from each MC are configured as a LAG and are distributed between the Aruba 3810M switches in the stack. The switch ports are configured as a dynamic port-channel on the Aruba MCs and LACP trunks on the Aruba 3810M switches.

First-hop router redundancy for the cluster management and client VLANs is natively provided by the stack of Aruba 3810M switches that provide the default gateway for each VLAN. One Aruba 3810M switch in the stack operates in a commander role while a second switch operates as a standby. The switch roles can be automatically or manually assigned. The commander switch provides IP forwarding during normal operation and the standby switch provides backup in the event that the commander switch fails.

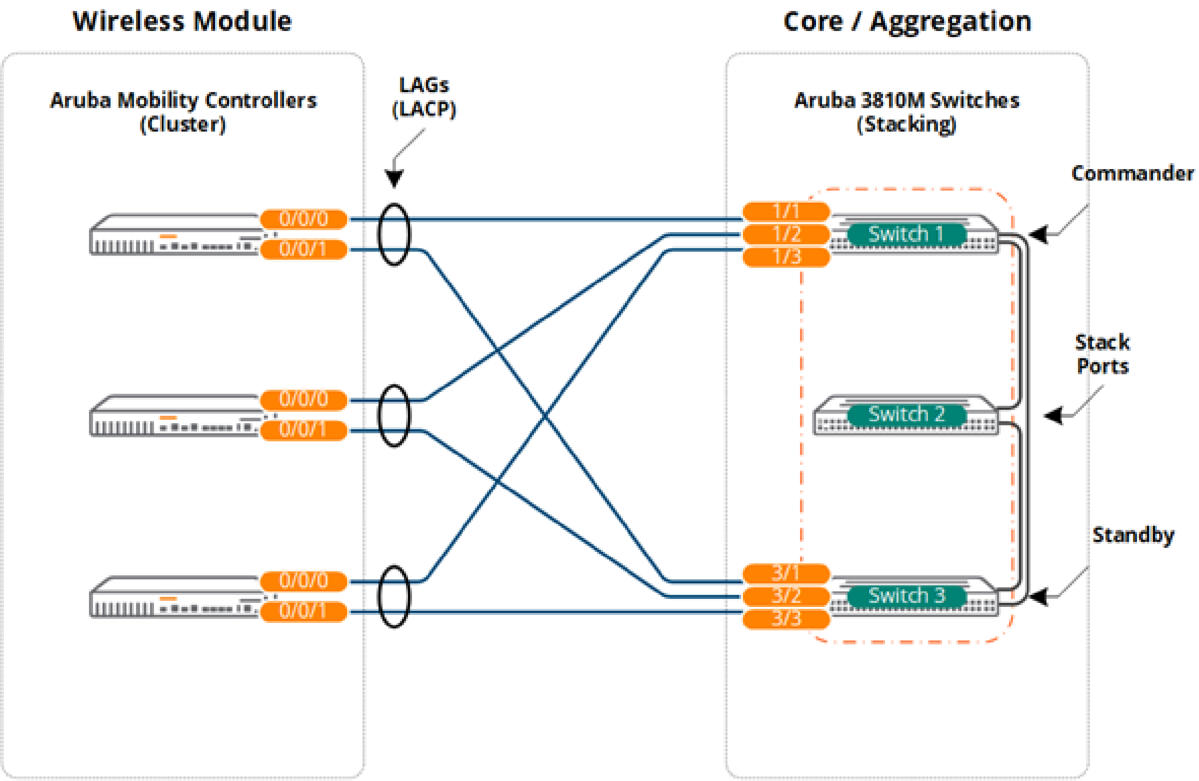

Figure 112 *Core/Aggregation using Stacking*

Aruba 5400R Switches

Figure 112 demonstrates how a cluster of MCs is connected to a pair of Aruba 5400R switches configured for VSF that have been deployed in either the core or wireless aggregation layer. The Aruba VSF architecture virtualizes both the control and data planes allowing all the pair of 5400R switches to forward traffic and be configured and be managed as a single virtual switch.

In this example two or more 1 Gigabit, 10 Gigabit, or 40 Gigabit Ethernet ports from each MC are configured as a LAG and are distributed between the pair of Aruba 5400R switches. The switch ports are configured as a dynamic port-channel on the Aruba MCs and LACP trunks on the Aruba 5400R switches.

First-hop router redundancy for the cluster management and client VLANs is natively provided by the VSF pair of Aruba 5400R switches that provide the default gateway for each VLAN. One Aruba 5400R switch operates in a “commander” role while the second switch operates as a “standby”. The switch roles can be automatically or manually assigned. The “commander” switch provides IP forwarding during normal operation while the “standby” switch provides backup in the event that the “commander” switch fails.

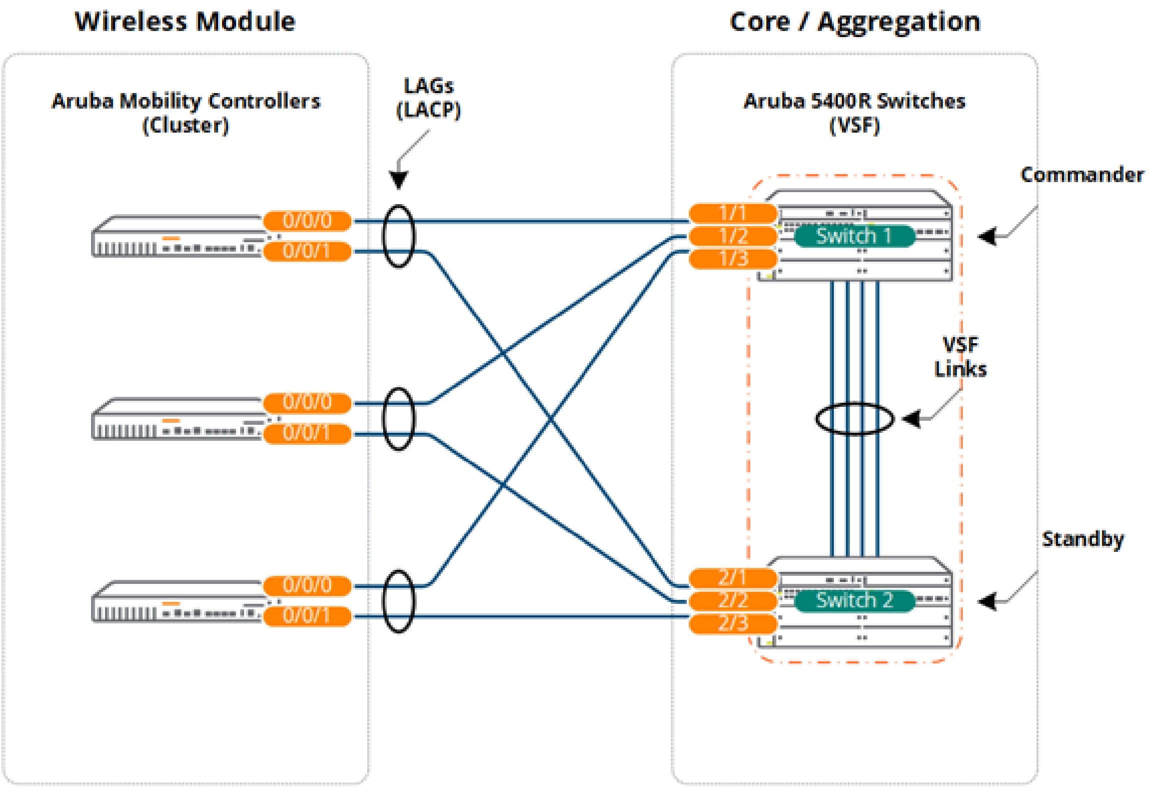

Figure 113 *Core/Aggregation using Virtual Switching Framework*

Aruba 8320/8400 Switches

The figure below demonstrates how a cluster of MCs is connected to a pair of Aruba 8320 or 8400 switches that have been configured for Multi-Channel LAG and are deployed in the core or wireless aggregation layer. The Aruba MC-LAG architecture virtualizes data planes allowing all the pair of 8320/8400 switches to forward traffic as a single virtual switch. Unlike the Aruba stacking or VSF architectures, each 8230/8400 is configured and managed independently.

In this example two or more 1 Gigabit, 10 Gigabit, or 40 Gigabit Ethernet ports from each MC are configured as a LAG and are distributed between the pair of Aruba 8320/8400 switches. The switch ports are configured as a dynamic port-channel on the Aruba MCs and MC-LAG on the Aruba 8320/8400 switches.

First-hop router redundancy for the cluster management and client VLANs is natively provided by the MC-LAG pair of Aruba 8320/8400 switches which provide the default gateway for each VLAN. The active gateway feature is enabled for each VLAN providing IP forwarding and failover on both switches.

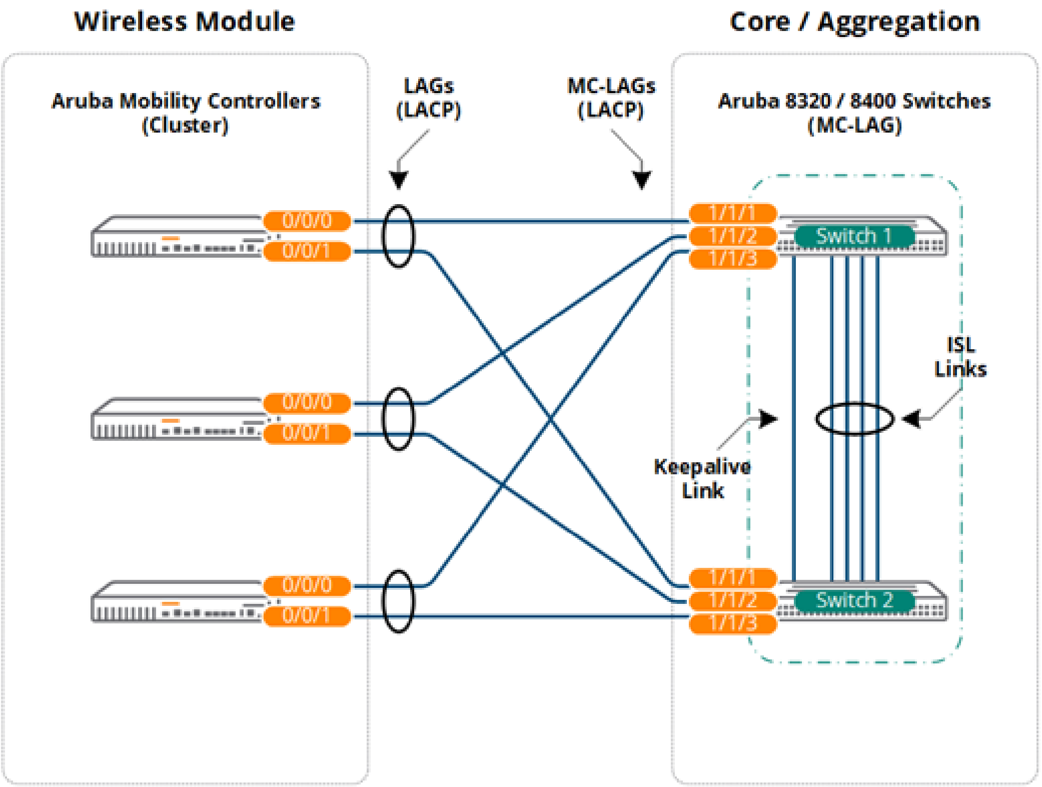

Figure 114 *Core/Aggregation using Multi-Chassis LAG*

Reference Architectures

This section includes mobile first reference architectures for small, medium, and large buildings as well as campuses consisting of multiple buildings of varying size. A scenario is provided for each architecture to provide a foundation upon which the modular network and wireless module design are based. Each architecture also builds upon the previous design adding additional layers as the access layer and client counts are increased.

Small Office

Scenario

The following reference design is for a small office consisting of a single floor. The building includes one MDF/server room and one IDF that connects to the MDF using multi-mode fiber. The building supports up to 150 employees and requires 15 802.11ac Wave 2 Access Points to provide full 2.4GHz and 5GHz coverage.

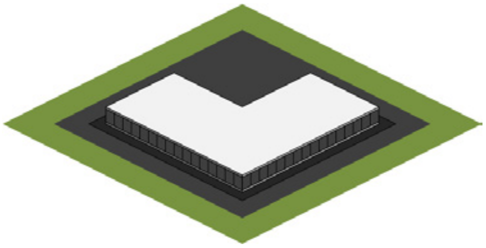

Building Characteristics:

- 1 Floor / 20,000 sq. ft. Total Size
- 150 x Employees / 300 x Concurrent IPv4 Clients
- 15 x 802.11ac Wave 2 Access Points
- 1 x Combined Server Room / Wiring Closet (MDF)
- 1 x Wiring Closet (IDF)

Figure 115 *Small Office Characteristics*

The building only has two wiring closets and therefore does not require an aggregation layer between the core and access layer. This building will implement a 2-tier modular network design where the access layer switches and modules connect directly to a collapsed core/aggregation layer. This 2-tier modular network design can also accommodate small buildings with a larger square footage and additional floors if required.

The following is a summary of the modular network architecture and design:

LAN Core/Aggregation:

- Cluster or stack of switches with mixed ports:
  - ➢ SFP/SFP+ (Access Layer Interconnects)
  - ➢ 10/100/1000BASE-T Ports (Module Connectivity)
- IP routing
- Layer 2 Link Aggregation to Access layer devices and Module Connectivity

LAN Access:

- A stack of two or more switches per wiring closet
  - ➢ SFP/SFP+ (core/aggregation layer interconnects)
  - ➢ 10/100/1000BASE-T with PoE+ (Edge Ports)
- Layer 2 link aggregation to core/aggregation layer devices
- 802.11ac Wave 2 APs

The number of APs required for this hypothetical scenario was calculated based on the buildings square footage along with the wireless density and capacity requirements. It was determined that 15 APs would be appropriate assuming that each AP provides 1,200 square feet of coverage. Each AP in this scenario is supporting 20 clients.

The actual number of APs and their placement for a production environment should be determined using a site survey that accounts for the density requirements for each individual coverage area.

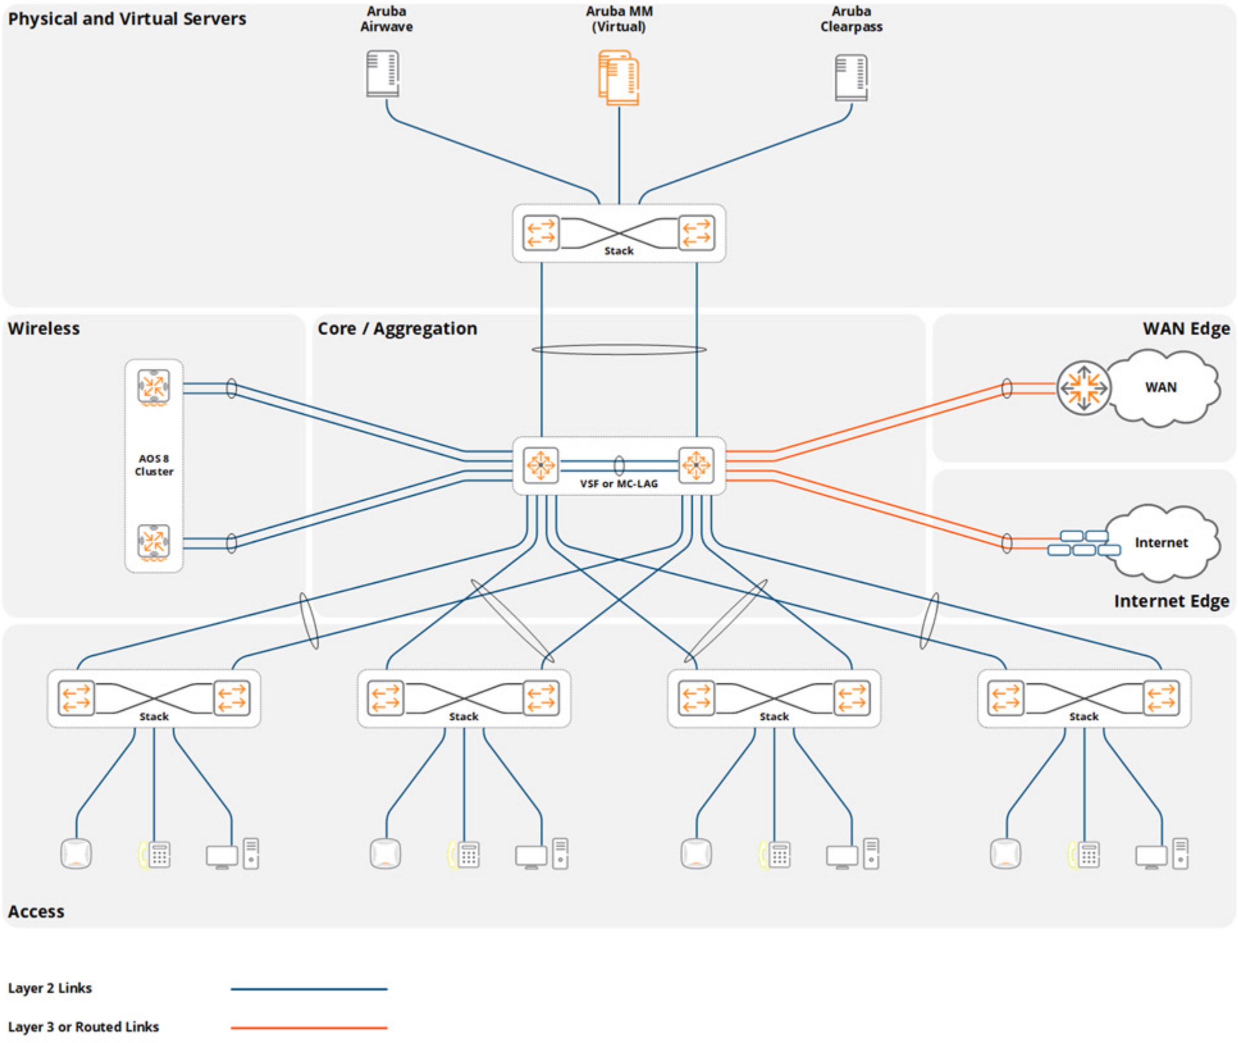

Figure 116 *Small Office 2-Tier Modular Network Design*

Considerations and Best Practices

Wireless LAN Components

Aruba offers both controller-less and controller-based deployment options for small deployments. A controller-less architecture is provided using Aruba Instant Access Points (IAPs) while a controller-based architecture is provided using MCs and Campus APs. Both deployment options are valid for this reference design, however this guide focuses specifically on a controller-based architecture.

The small building in this scenario includes various wireless components which are either deployed in the wireless module or server room. A MM and a single cluster of MCs is required to accommodate the AP and client counts. The exact number of cluster members is determined by the hardware or Virtual MC model that has been selected. For redundancy purposes the MC cluster consists of a minimum of two MCs. Each cluster member needs to provide adequate capacity and performance to operate the wireless network in the event of a single MC failure.

Table 22 below provides a summary of these components:

Component	Description	Notes
Aruba MM (MM)	Virtual Appliance	1 Required, 2 Recommended
Aruba MCs	Hardware or Virtual Appliances	2 Minimum (Clustered)
Aruba Access Points	802.11ac Wave 2 Access Points	15 Required
Aruba ClearPass	Virtual Appliance	Recommended

Table 22 *Small Building Wireless LAN Components*

While the number of 802.11ac Wave 2 APs required for this design is relatively small, Aruba recommends implementing a MM to take advantage of specific features that are required to provide mission-critical wireless services when wireless is the primary access medium. The addition of a MM to the design provides centralized configuration and monitoring, supports features including clustering, AirMatch, and Live Upgrades, and provides centralized application support (UCC and AppRF).

While a controller-based solution can be deployed without a MM, it is not a recommended best practice. If it is not feasible to deploy a MM the MCs can optionally be deployed as a pair of standalone devices and be configured for master redundancy. However, a deployment model implemented in such a fashion will not support clustering and therefore will lack specific features such as fast failover, live upgrades, AirMatch, and centralized application support.

Redundancy

Redundancy for a small building reference architecture is provided across all layers. The redundancy built into the 2-tier modular network design that establishes the foundation network is what determines the level of redundancy that is provided to the modules. Often the cost of an outage is the key driver in implementing to provide network redundancy. Most small networks use dual power supplies and often use a stack of switches as their primary redundancy mechanism.

For this scenario the MM and MC cluster members are deployed within a server room and are directly connected to the core/aggregation switches. To provide full redundancy, two virtual MMs and one cluster of hardware or virtual MCs is required:

- Aruba MM (MM):
  - ➢ Two virtual MMs
  - ➢ L2 master redundancy (Active / Standby)
- Hardware MCs (MCs):
  - ➢ Single cluster of hardware MCs
  - ➢ Minimum of two cluster members
- Virtual MCs (MCs):

  - Single cluster of virtual MCs
  - Minimum of two cluster members
  - Separate virtual server hosts
- Access Points
  - AP Master pointing to the cluster’s VRRP VIP
  - Fast failover using cluster’s redundancy functionality

Figures 117 and 118 provide detailed examples of how the virtual and hardware cluster members are connected to the core/aggregation layer. Hardware MCs are directly connected to the core/aggregation layer switches via two or more 1 Gigabit Ethernet ports configured in a LAG group. The LAG port members are distributed between core/aggregation layer stack members.

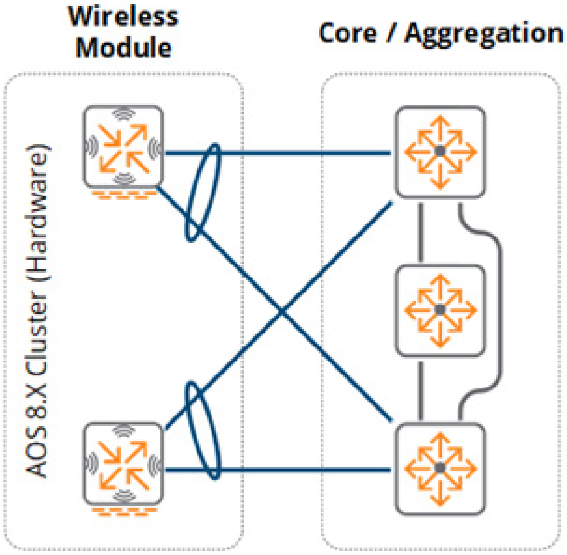

Figure 117 *Hardware MC Cluster – Core/Aggregation Layer*

VMCs are logically connected to a virtual switch within the virtual server host. The virtual host server is directly connected to the core/aggregation switches via two or more 1 Gigabit or 10 Gigabit Ethernet ports implementing 802.3ad link aggregation or a proprietary load-balancing/failover mechanism. Each port is distributed between core/aggregation layer switch stack members.

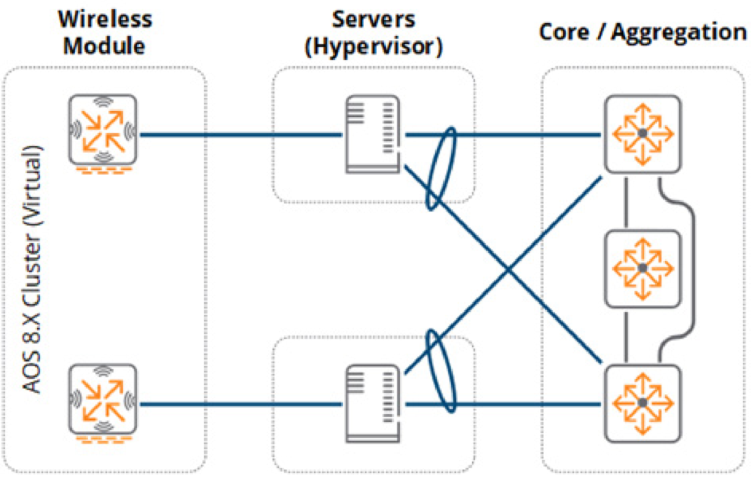

Figure 118 *Virtual MC Cluster – Core/Aggregation Layer*

The MMs are deployed in a similar manner to the VMC clusters. Each virtual server host supports one VMM operating in active/standby mode. While a single MM can be implemented for a small building there are no additional licenses required to implement a standby. The only network overhead for such a model would be the additional CPU, memory, and storage utilization on the virtual server host.

NOTE

Redundancy for virtual servers is hypervisor-dependent. To safeguard the network against link, path, and node failures, the hypervisor may implement 802.3ad link aggregation or a proprietary load-balancing/failover mechanism.

Virtual MCs

Virtual MCs can optionally be deployed for a small building environments. If VMCs are deployed then the virtual server infrastructure must be scaled accordingly to provide the necessary CPU and memory resources to each VMC in the cluster:

1. Each VMC should be deployed across different virtual server hosts. This design requires two virtual server hosts
2. Uplinks between the virtual server host and the core/aggregation layer must be scaled accordingly to support the wireless and dynamically segmented client throughput requirements. The throughput of cluster will be limited by the Ethernet PHYs installed on the virtual server host

Redundancy between the virtual server host and its peer switches can use standard 802.3ad link aggregation or a proprietary hypervisor specific load-balancing and failover mechanism. Each hypervisor supports specific load-balancing and failover mechanisms such as active/standby, round-robin load-balancing, or link aggregation. The appropriate redundancy mechanism should be selected to support the specific implementation requirements.

Scalability

For this scenario there are no specific LAN scalability considerations that need to be taken into account. The core/aggregation and access layers can easily accommodate the APs and client counts without modification or derivation from the base design. A wireless aggregation layer can be added in the future if necessary to accommodate additional APs and clients that are added to the network.

Wireless module scalability is also not a concern as the MMs can be expanded and additional cluster members added over time to accommodate additional APs, clients, and switching capacity as the network increases in size.

As a best practice, Aruba recommends implementing the MM-VA-50 MM and a cluster of two hardware or virtual MCs for such a small building design per the platform suggestions. The MM selected for this design can scale to support 50 APs, 500 clients, and 5 MCs.

Virtual LANS

For this design the core/aggregation layer provides layer 2 transport (VLAN extension via 802.1q trunking) and terminates all the VLANs from the access layer and wireless module with layer three interfaces. Aruba recommends using tagged VLANs throughout the network.

The wireless module consists of one or more client VLANs depending on the security and policy model. For a single VLAN design, all wireless and dynamically segmented clients are assigned a common VLAN with roles and policies determining the appropriate level of network access for each client. The single VLAN is extended from the core/aggregation layer switches to each physical or virtual MC cluster member. Additional VLANs can be added and extended as required. E.g., the mobile first design may require separate VLANs to be assigned to wireless and dynamically segmented clients for policy compliance purposes.

A minimum of two VLANs are required between the core/aggregation layer and each MC cluster member. One VLAN is dedicated for management and MM communications while the second VLAN is used for client traffic. All VLANs are common between cluster members to permit seamless mobility. The core/aggregation layer switches are configured with layer three interfaces and addressing to operate as the default gateway for each VLAN. First-hop router redundancy is natively provided by the Aruba stacking architecture.

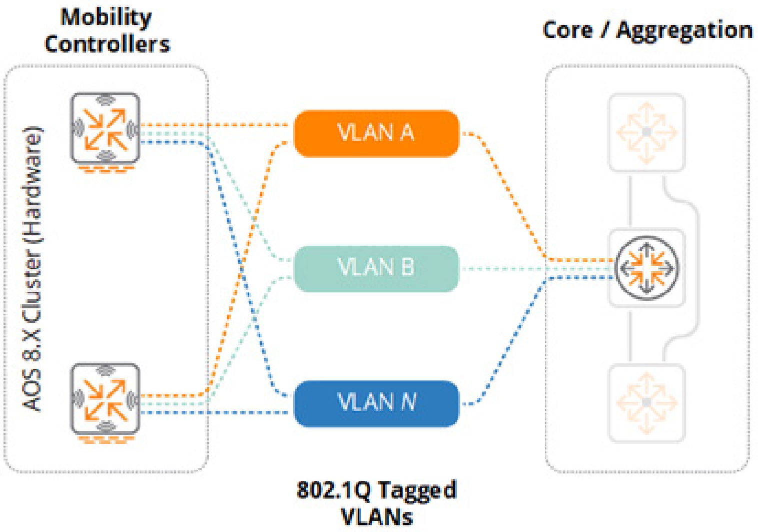

Figure 119 *Hardware MC Cluster – VLANs*

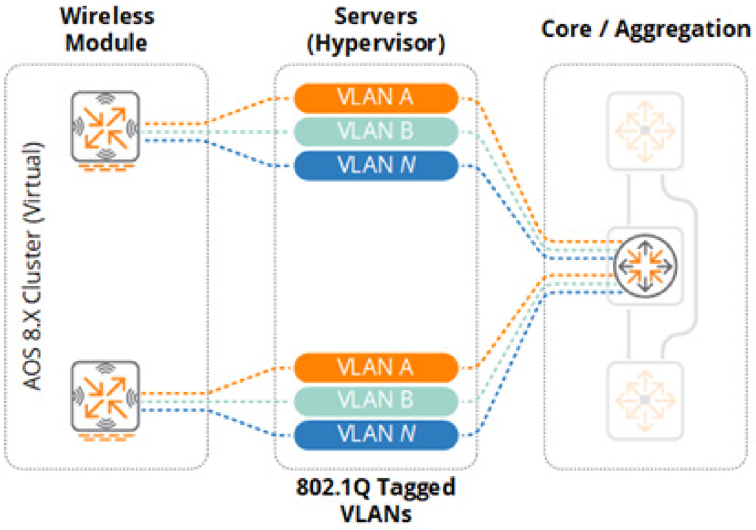

Figure 120 *Virtual MC Cluster – VLANs*

As a best practice Aruba recommends implementing unique VLAN IDs within the wireless module. This allows an aggregation layer to be introduced in the future without disrupting the other layers within the network. It also facilitates the creation of smaller layer 2 domains which is critical to reducing layer 2 instability. Operational changes, loops, or misconfigurations originating from other layers or modules in the network can adversely impact the wireless module unless the network has been properly segmented with appropriately sized layer 2 domains.

Platform Suggestions

Figure 121 below provides platform suggestions for a small building scenario supporting 15 APs and 300 concurrent clients. A “good, better, and best suggestion” is made based on feature, performance, and scaling capabilities. These are suggestions are scenario-specific and may be substituted at your own discretion.

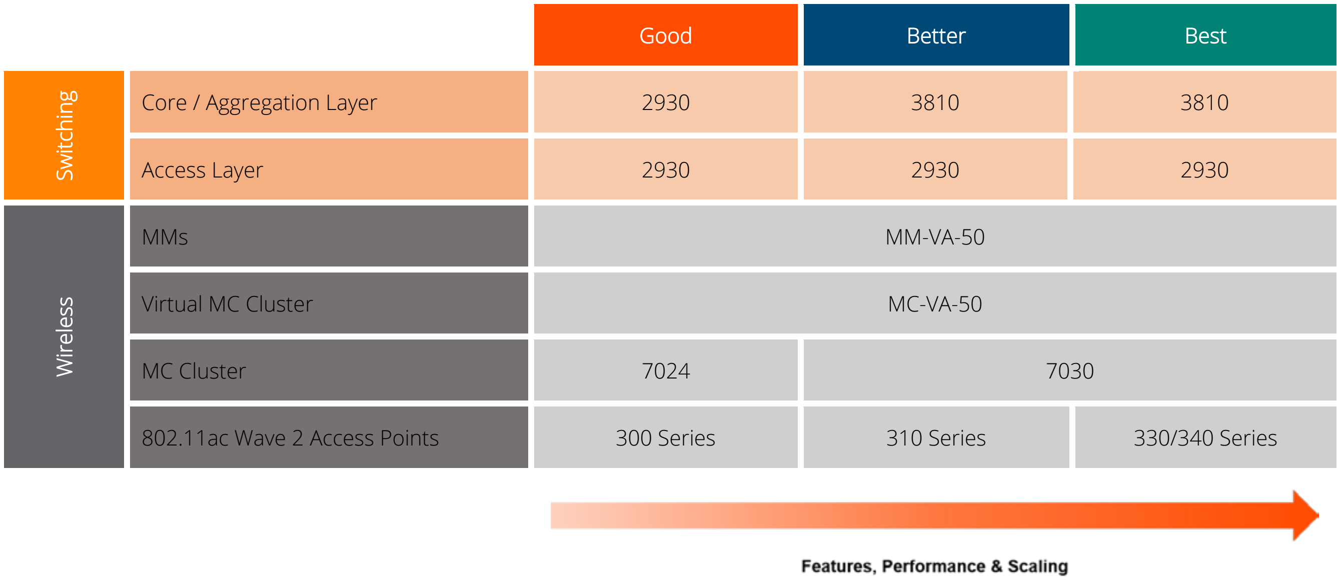

		Good	Better	Best
Switching	Core / Aggregation Layer	2930	3810	3810
	Access Layer	2930	2930	2930
Wireless	MMs	MM-VA-50		
	Virtual MC Cluster	MC-VA-50		
	MC Cluster	7024	7030	
	802.11ac Wave 2 Access Points	300 Series	310 Series	330/340 Series

Features, Performance & Scaling

Figure 121 *Small Building Platform Suggestions*

Medium Office

Scenario

The following reference design is for a medium office building with six floors. The building includes a data center which is connected to an MDF on each floor through single-mode fiber. Each floor includes three IDFs which connect to the MDF through multi-mode fiber. The building supports up to 1,500 employees and requires 120 802.11ac Wave 2 APs to provide full 2.4 and 5GHz coverage.

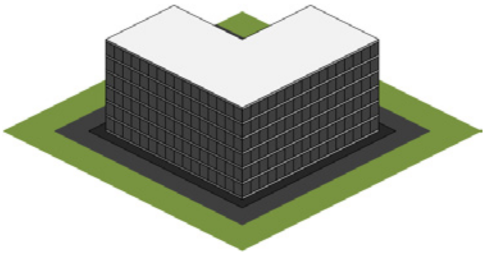

Building Characteristics:

- 6 Floors / 150,000 sq. ft. Total Size
- 1,500 x Employees / 3,000 x Concurrent IPv4 Clients
- 120 x 802.11ac Wave 2 Access Points
- 1 x Computer Room
- 1 x MDF per floor (6 total)
- 2 x IDFs per floor (12 total)

Figure 122 *Medium Office Characteristics*

As this design implements a structured wiring design using MDFs and IDFs, an aggregation layer is required to connect the access layer. This building will also implement a 3-tier modular network design where the access layer switches connect to aggregation layer switches in each MDF which in turn connect directly to the core. This modular network design also includes an additional aggregation layers for the computer room which facilitates scalability, aggregation, and fault domain isolation.

The list below provides an outline of the modular network architecture and design:

  - **LAN Core:** A cluster of switches with fiber ports:
  - SFP/SFP+ (module connectivity)
- IP routing to aggregation layer devices and modules

LAN Aggregation:

- A stack of two switches with fiber ports for each MDF:
  - SFP/SFP+/QSFP+ (core and access layer interconnects)
- IP routing to core layer devices
- Layer 2 link aggregation to access layer devices

LAN Access:

- A stack of two or more switches for each MDF and IDF:
  - SFP/SFP+ (aggregation layer interconnects)
  - 10/100/1000BASE-T with PoE+ (edge ports)
- Layer 2 link aggregation to aggregation layer devices
- 802.11ac Wave 2 APs

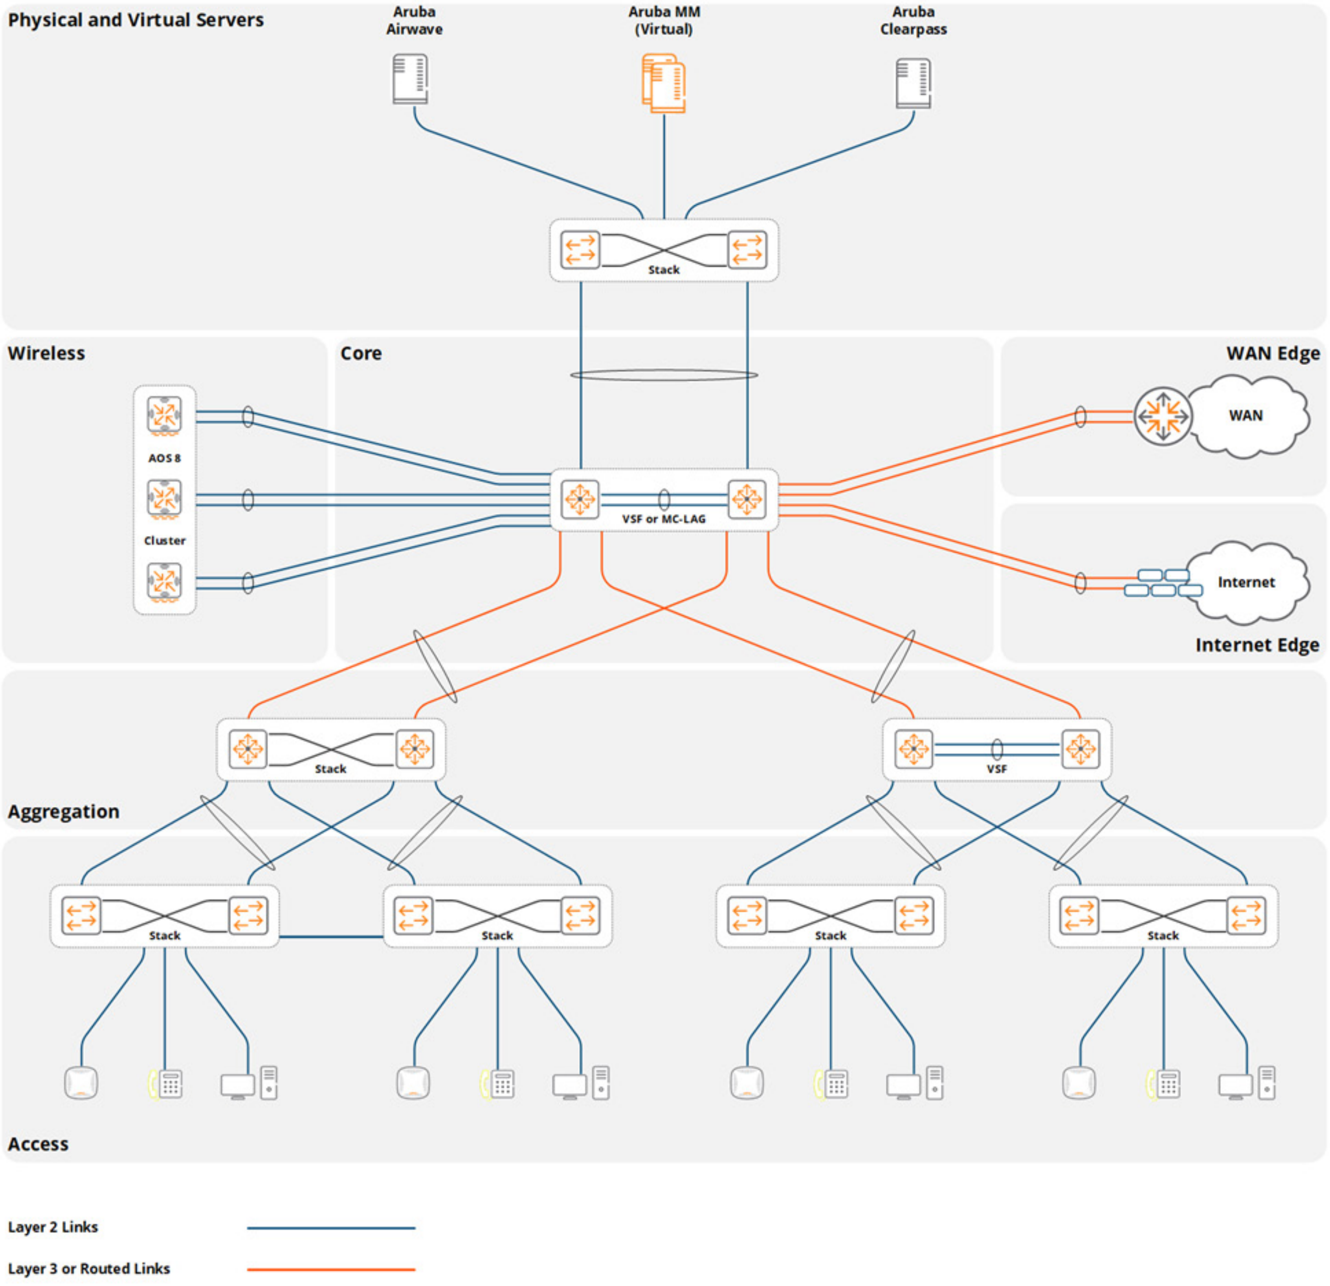

Figure 123 *Medium Office – 3-Tier Modular Network Design*

The number of APs required for this hypothetical scenario was calculated based on the building's square footage, wireless density, and capacity requirements. It was determined that 120 APs would be required based on the assumption that each AP will provide coverage for 1,200 square feet and support 25 clients.

The actual number of APs and their placement for a production environment should be determined using a site survey that accounts for the density requirements for each individual coverage area.

Considerations and Best Practices

Wireless LAN Components

The medium building in this scenario includes various wireless components which are either deployed in the wireless module or the server room. To accommodate the AP and client counts, an MM and a single cluster of MCs is required. The number of cluster members is determined by the MC (either hardware or virtual) model that is selected. The MC cluster consists of a minimum of two MCs for redundancy purposes. Each member provides adequate capacity and performance to allow the wireless network to continue to function in the event of a single MC failure.

Table 23 provides a summary of these components:

Component	Description	Notes
Aruba MM (MM)	Virtual Appliance	1 MM is required, 2 are recommended
Aruba MCs	Hardware or Virtual Appliances	A minimum of 2 MCs are required (Clustered)
Aruba Access Points	802.11ac Wave 2 Access Points	120 APs are required
Aruba ClearPass	Virtual Appliance	Recommended

Table 23 *Medium Building – Wireless LAN Components*

Redundancy

Redundancy for a medium building reference architecture is provided across all layers. The redundancy built into the 3-tier modular network design that establishes the foundation network determines the level of redundancy that is provided to the modules. Aruba recommends using NVF functions (stacking or MC LAG) to provide network and link redundancy as well as redundant power supplies to maximize network availability and resiliency.

For this scenario the MM and cluster members are deployed within a computer room and connect directly to the core or computer room aggregation switches. Two VMMs and one cluster of hardware or virtual MCs is required to fully enable redundancy:

- Aruba MM (MM):
  - Two virtual MMs
  - L2 master redundancy (active/standby)
- Hardware MCs (MCs):
  - Single cluster of hardware MCs
  - Minimum of two cluster members

- Virtual MCs (VMCs):
    - Single cluster of virtual MCs
    - Minimum of two cluster members
    - Separate virtual server hosts
- Access Points
    - AP Master pointing to the cluster's VRRP VIP address
    - Fast failover using built in cluster redundancy

Figures 124 and 125 provide detailed examples for how the virtual and hardware cluster members are connected to their respective layers. Hardware MCs are directly connected to the core layer switches via two or more 1 Gigabit or 10 Gigabit Ethernet ports configured in a LAG group. The LAG port members being distributed between redundant core/aggregation switches.

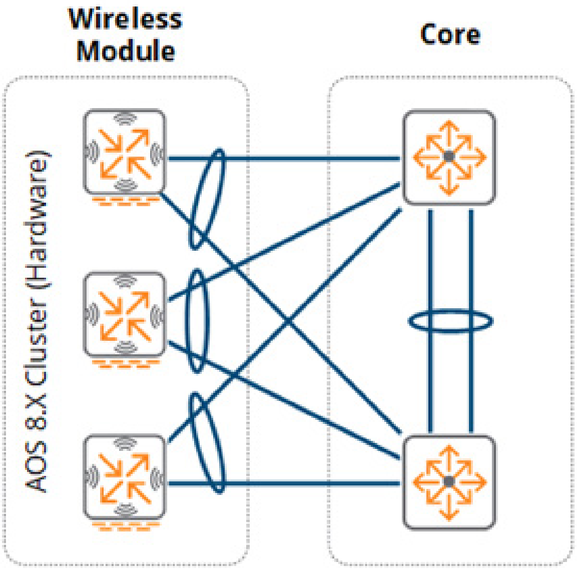

Figure 124 *Hardware MC Cluster – Core Layer*

VMCs are logically connected to a virtual switch within the virtual server host. The virtual host server is directly connected to the computer room aggregation switches via two or more 1 Gigabit or 10 Gigabit Ethernet ports implementing 802.3ad link aggregation or a proprietary load-balancing and failover mechanism. Each port is distributed between redundant computer room aggregation switches.

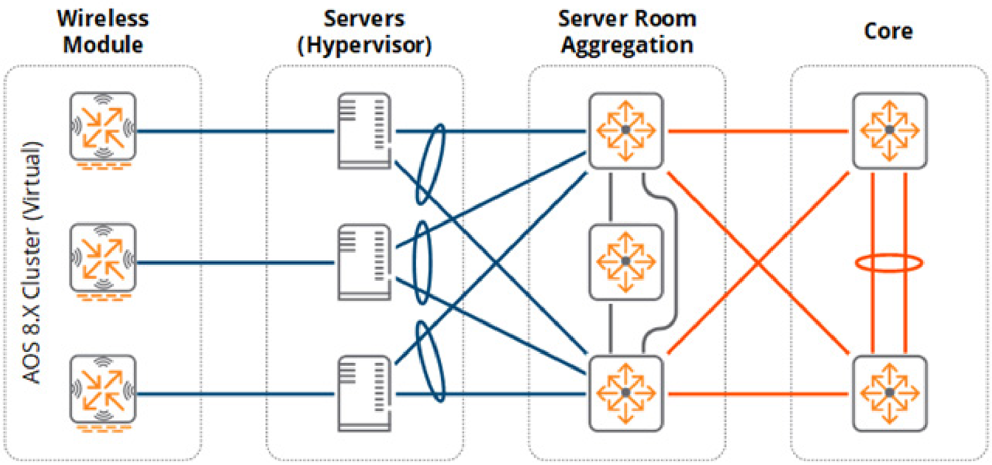

Figure 125 *Virtual MC Cluster – Computer Room Aggregation Layer*

The MM(s) are deployed in a similar manner to the cluster of VMCs. Each virtual server host supports one virtual MM operating in an active/standby mode.

NOTE

Redundancy for virtual servers is dependent on the hypervisor. To provide against link, path and node failures, the hypervisor may implement 802.3ad link aggregation or a proprietary load-balancing and failover mechanism.

Virtual MCs

For medium building deployments VMCs may also be deployed as an alternative to hardware MCs. If VMCs are deployed, the virtual server infrastructure must be scaled accordingly to provide the necessary CPU and memory resources to support each virtual MC in the cluster:

1. Each VMC in the cluster should be deployed across different virtual server hosts. For this design two virtual server hosts are required.
2. Uplinks between the virtual server host and the computer room aggregation layer must be scaled accordingly to support the wireless and dynamically segmented client throughput requirements. The throughput of cluster will be limited by the Ethernet PHYs installed on the virtual server host.

Redundancy between the virtual server host and its peer switches can use either standard 802.3ad link aggregation or a proprietary hypervisor specific load-balancing and failover mechanism. Each hypervisor supports specific load-balancing and failover mechanisms such as active/standby, round-robin load-balancing, or link aggregation. The appropriate redundancy mechanism should be selected to support the specific implementation requirements for each site.

Scalability

For this scenario there are no specific LAN scalability considerations that need to be taken into account. The core, aggregation, and access layers can easily accommodate the APs and client counts without modification or derivation from the design. A wireless aggregation layer can be added in the future as additional APs and clients are added to the network.

Wireless module scaling is also not a concern as the MMs can be expanded and additional cluster members added over time to accommodate additional APs, clients, and switching capacity as the network scales.

For this medium building design Aruba recommends implementing the MM-VA-500 MM and a cluster of two or more hardware or virtual MCs per the platform suggestions. The MM selected for this design can scale to support 500 APs, 5,000 clients, and 50 MCs.

Virtual LANs

In the medium office design the core or computer room aggregation layer terminates all VLANs from the MCs. The VLANs are extended from the MCs to the core or computer room aggregation layer using 802.1Q trunking. Aruba recommends using tagged VLANs wherever possible to provide additional loop prevention. The wireless module consists of one or more user VLANs depending on the security and policy model. For a single VLAN design, all wireless and dynamically segmented clients are assigned to a common VLAN ID. Roles and policies determine the level of access each user is provided on the network. The single VLAN is extended from the core or computer room aggregation layer switches to each physical or virtual MC cluster member. Additional VLANs can be added and extended as required. E.g., the mobile first design may require separate VLANs to be assigned to wireless and dynamically segmented clients for policy compliance reasons.

A minimum of two VLANs are required between the core or computer room aggregation layer and each MC cluster member. One VLAN is dedicated for management and MM communications while the second VLAN is mapped to clients. All VLANs are common between cluster members to permit seamless mobility. The core or computer room aggregation layer switches have VLAN-based IP interfaces defined and operate as the default gateway for each VLAN. First-hop router redundancy is natively provided by the Aruba stacking architecture.

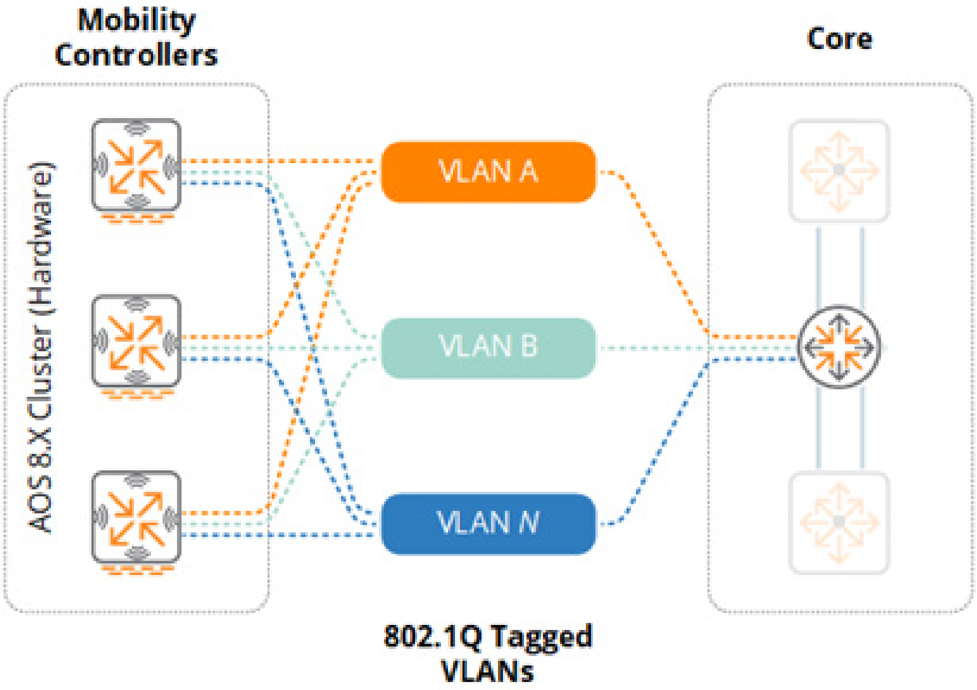

Figure 126 *Hardware MC Cluster – VLANs*

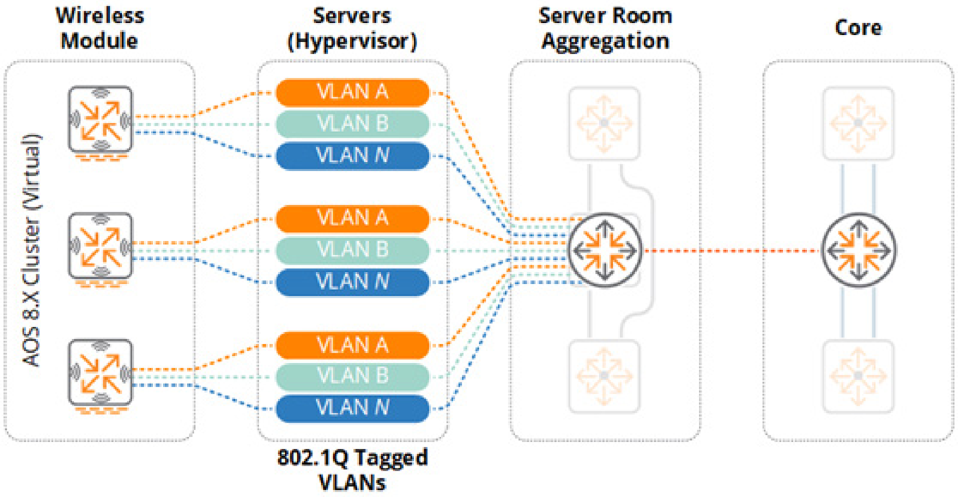

Figure 127 *Virtual MC Cluster – VLANs*

As a best practice, Aruba recommends implementing unique VLAN IDs within the wireless module. This allows for an aggregation layer to be introduced in the future without disrupting the other layers within the network. This also allows for the creation of smaller layer 2 domains. Segmenting the network in this manner reduces layer 2 instability and protects the wireless module from operational changes, loops, or misconfigurations originating from other layers or network modules.

Platform Suggestions

The figure below provides platform suggestions for the medium building scenario supporting 120 APs and 3,000 concurrent clients. A good, better, and best suggestion is made based on features, performance, and scalability. These are suggestions based on the described scenario and may be altered according to the discretion of network administrators.

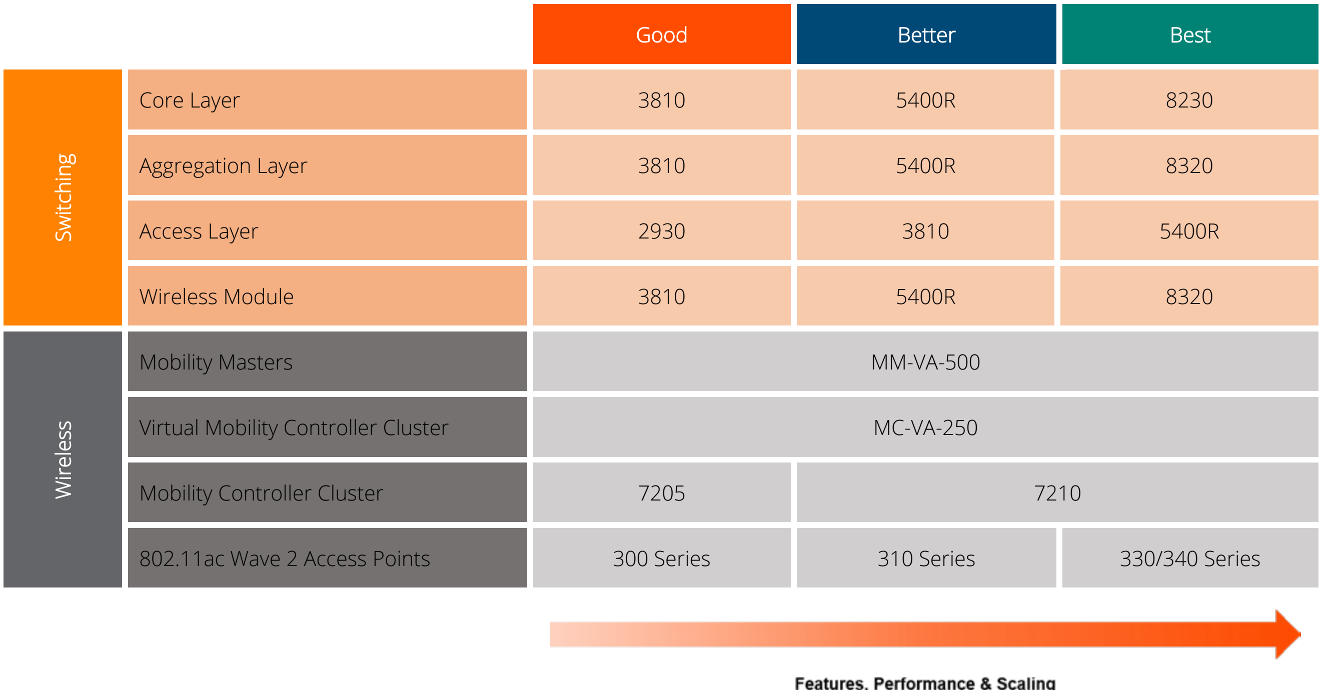

		Good	Better	Best
Switching	Core Layer	3810	5400R	8230
	Aggregation Layer	3810	5400R	8320
	Access Layer	2930	3810	5400R
	Wireless Module	3810	5400R	8320
Wireless	Mobility Masters	MM-VA-500		
	Virtual Mobility Controller Cluster	MC-VA-250		
	Mobility Controller Cluster	7205	7210	
	802.11ac Wave 2 Access Points	300 Series	310 Series	330/340 Series

Features, Performance & Scaling

Figure 128 *Medium Building Platform Suggestions*

Large Office

Scenario

The following reference design is for a large office consisting of 12 floors. The building includes a data center which connects via single-mode fiber to an MDF on each floor. Each floor also includes three IDFs which connect to the MDF via multi-mode fiber. The building supports up to 3,000 employees and requires 300 802.11ac Wave 2 APs to provide full 2.4 and 5GHz coverage.

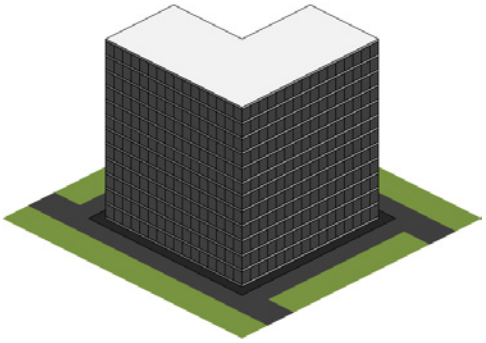

Building Characteristics:

- 12 Floors / 360,000 sq. ft. Total Size
- 3,000 x Employees / 6,000 x Concurrent IPv4 Clients
- 300 x 802.11ac Wave 2 Access Points
- 1 x Computer Room
- 1 x MDF per floor (12 total)
- 2 x IDFs per floor (24 total)

Figure 129 *Large Office Characteristics*

The large building design implements a structured wiring design using MDFs and IDFs which requires an aggregation layer to connect the access layer. The building implements a 3-tier modular network design where the access layer switches connect via aggregation layer switches in each MDF which in turn connect directly to the core. The modular network design also includes additional aggregation layers for the computer room and wireless modules for scalability, aggregation, and fault domain isolation.

The list below provides an outline of the modular network architecture and design:

LAN Core:

- A pair of redundant switches with a mix of 10G and 40G fiber ports:
  - SFP/SFP+/QSFP+ (aggregation layer interconnects)
- IP routing to aggregation layer devices and modules
- Optional NVF Functions (MC LAG)

LAN Aggregation:

- A stack of two switches with fiber ports per MDF:
  - SFP/SFP+/QSFP+ (core and access layer interconnects)
- NVF Functions (MCLAG/VSX)
- IP routing to core layer devices

LAN Access:

- A stack of two or more switches per MDF and IDF:
  - SFP/SFP+ (aggregation layer interconnects)
  - 10/100/1000BASE-T with PoE+ (edge ports)
- Layer 2 link aggregation to access layer devices
- 802.11ac Wave 2 APs

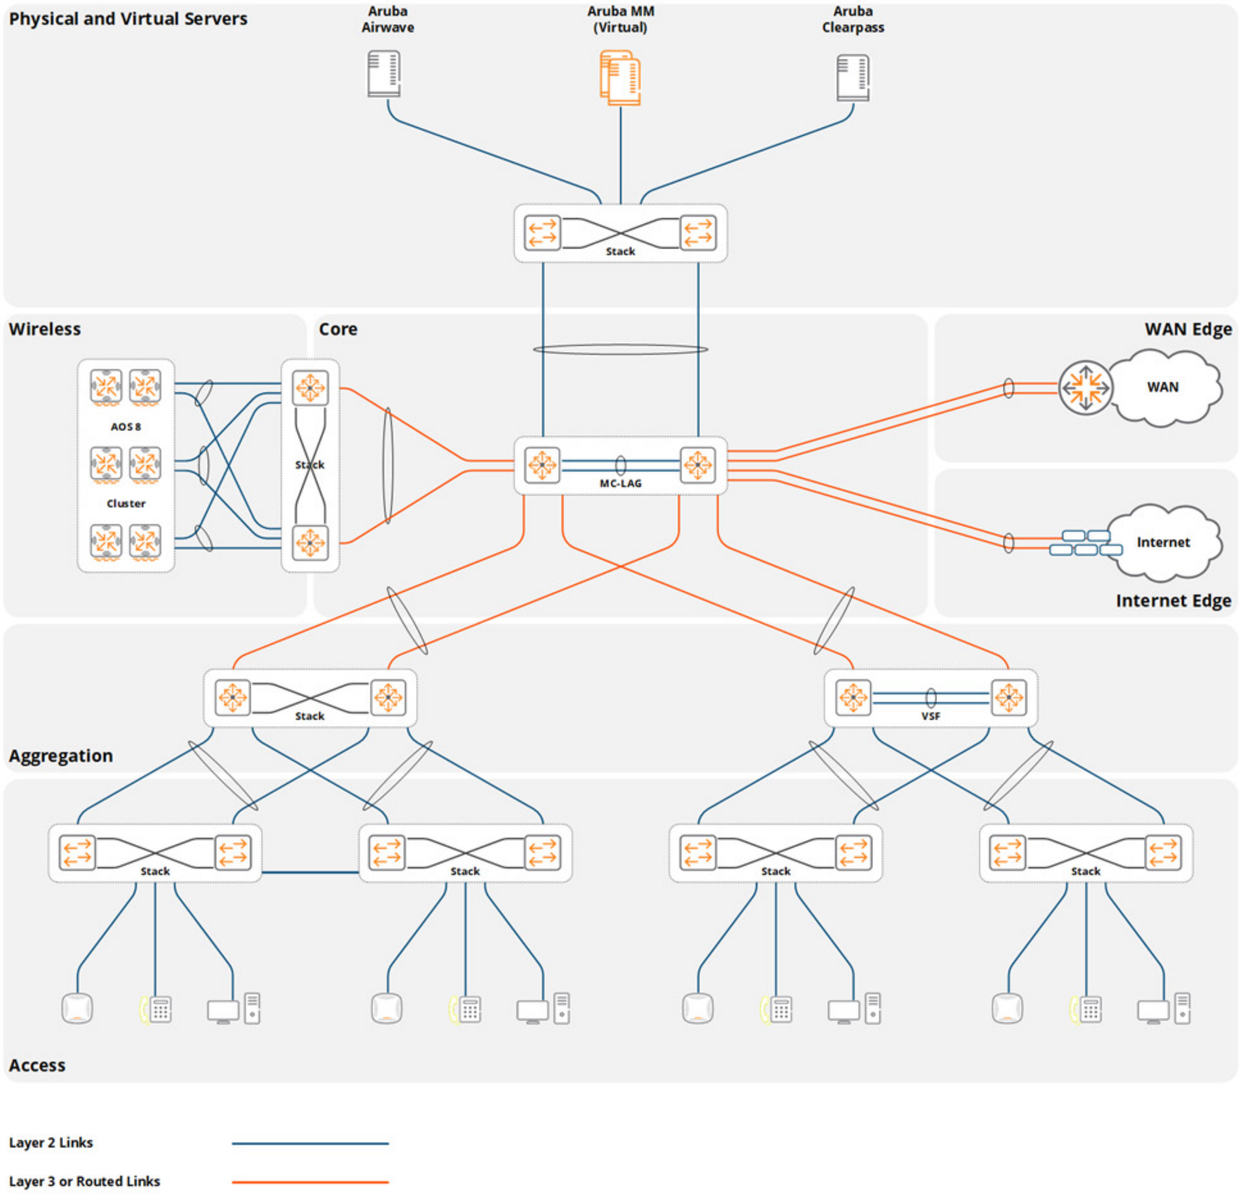

Figure 130 *Large Office – 3-Tier Modular Network Design*

The number of APs required for this hypothetical scenario was calculated based on the square footage, wireless density, and capacity requirements for the building. It was determined that 300 APs would be required based on an assumption of each AP providing 1200 square feet of coverage and supporting 20 clients.

The actual number of APs and their placement for a production deployment should be determined using a site survey which takes into account the density requirements for each coverage area.

Considerations and Best Practices

Wireless LAN Components

The large building in this scenario includes various wireless components which are either deployed in the wireless module or the server room. To accommodate the AP and client counts for this scenario an MM and a single cluster of MCs is required. The number of cluster members is determined by the hardware or virtual MC model that is selected. The MC cluster consists of a minimum of two MCs for redundancy purposes. Each member provides adequate capacity and performance to allow the wireless network to continue to function in the event of a single MC failure.

Table 24 provides a summary of these components:

Component	Description	Notes
Aruba MM (MM)	Hardware or Virtual Appliances	2 required
Aruba MCs	Hardware or Virtual Appliances	2 minimum (clustered)
Aruba Access Points	802.11ac Wave 2 Access Points	300 required
Aruba Airwave	Hardware or Virtual Appliance	Recommended
Aruba ClearPass	Hardware or Virtual Appliance	Recommended

Table 24 *Large Building – Wireless LAN Components*

Redundancy

Redundancy for a large building architecture is provided across all layers. The redundancy built into the 3-tier modular network design that establishes the foundation network determines the level of redundancy that is provided to the modules. Aruba recommends using NVF functions (stacking or MC LAG) to provide network redundancy as well as using redundant links and power supplies to maximize network availability and resiliency. The Aruba 8400 provides the maximum redundancy capabilities of any Aruba Switch and is recommended for use in the Core, Aggregation, and Wireless Aggregation layers.

For this scenario the MM and mobility cluster members are deployed within a computer room and are directly connected to the wireless aggregation or computer room aggregation switches. Two hardware or virtual MMs and one cluster of hardware or virtual MCs is required to fully enable redundancy:

- Aruba MM (MM):
  - ➢ Two hardware or virtual MMs
  - ➢ L2 master redundancy (active/standby)
- Hardware MCs (MCs):

- Single cluster of hardware MCs
  - Minimum of two cluster members
- Virtual MCs (MCs):
  - Single cluster of virtual MCs
  - Minimum of two cluster members
  - Separate virtual server hosts
- Access Points
  - AP Master pointing to the cluster's VRRP VIP
  - Fast failover using built in cluster redundancy

Figures 131 and 132 below provide detailed examples for how the virtual and hardware cluster members are connected to their respective layers. Hardware-based MCs are directly connected to the core layer switches via two or more 10 gigabit Ethernet ports configured in a LAG. The LAG port members are distributed between redundant wireless aggregation switches.

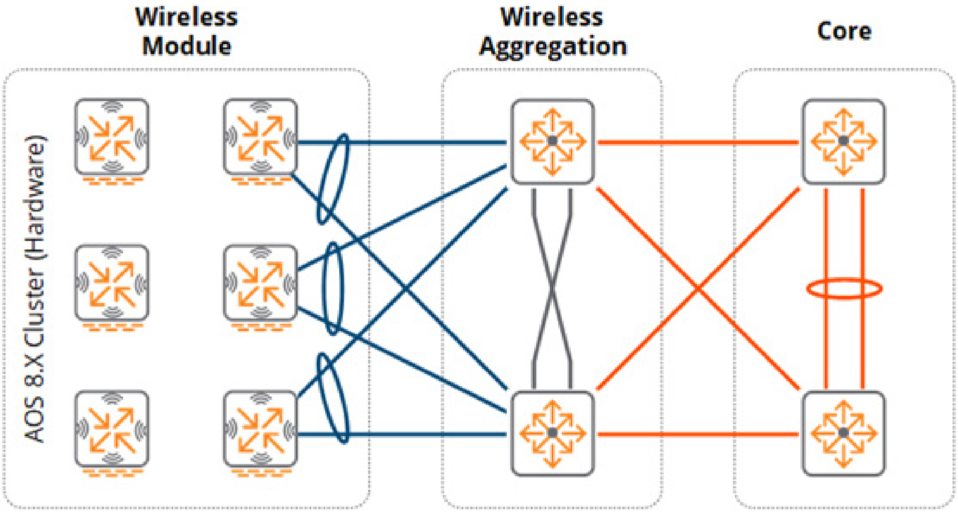

Figure 131 *Hardware MC Cluster – Wireless Aggregation Layer*

VMCs are logically connected to a virtual switch within the virtual server host. The virtual host server is directly connected to the computer room aggregation switches via two or more 10 gigabit Ethernet ports implementing 802.3ad link aggregation or a proprietary load-balancing and failover mechanism. Each port is distributed between redundant computer room aggregation switches.

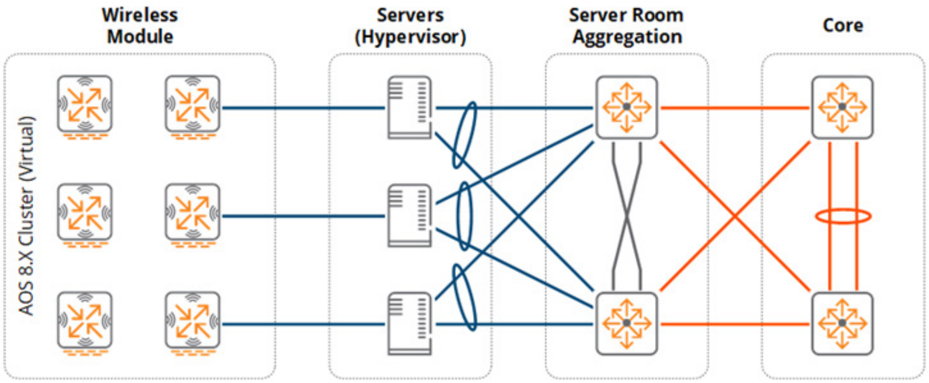

Figure 132 *Virtual MC Cluster – Computer Room Aggregation Layer*

The MM(s) are deployed in a similar manner as the cluster of VMCs. Each virtual server host supports one virtual MM operating in an active/standby mode.

NOTE

Redundancy for virtual servers is hypervisor-dependent. To provide against link, path, and node failures, the hypervisor may implement either 802.3ad link aggregation or a proprietary load-balancing and failover mechanism.

Virtual MCs

VMCs may be optionally deployed for large building deployments. If VMCs are deployed the virtual server infrastructure must be scaled accordingly to provide the necessary CPU and memory resources for each VMC in the cluster:

1. Each VMC should be deployed across different virtual server hosts. Two virtual server hosts are required for the large office design.
2. Uplinks between the virtual server host and the computer room aggregation layer must be scaled accordingly to support the wireless and dynamically segmented client throughput requirements. The throughput of cluster will be limited by the Ethernet PHYs installed on the virtual server host.

Redundancy between the virtual server host and its peer switches can use standard 802.3ad link aggregation or a proprietary hypervisor specific load-balancing and failover mechanism. Each hypervisor supports specific load-balancing and failover mechanisms such as active/standby, round-robin load-balancing, and link aggregation. The appropriate redundancy mechanism should be selected to support the specific implementation requirements of the deployment.

Scalability

The large office design includes a wireless aggregation layer to accommodate 6,000 wireless IPv4 hosts on the network. As a general best practice Aruba recommends considering a wireless aggregation layer once the combined IPv4 and IPv6 host count exceeds 4,094. The wireless aggregation layer is needed if hardware MCs are deployed and are connected directly to the core layer. If VMCs are deployed then the computer room aggregation switches provide the aggregation functionality.

Future growth is not a concern as the MMs can be easily expanded and additional cluster members can be added over time to accommodate additional APs, clients, and switching capacity. For this large building design, Aruba recommends implementing the MM-HW-5K or MM-VA-5K MM and a cluster of two or more hardware or virtual MCs. The MM selected for this design can scale to support 5,000 APs, 50,000 clients, and 500 MCs.

Virtual LANs

In the large office design the wireless module aggregation layer terminates all layer 2 VLANs from the MCs. The VLANs are extended from the MCs to their respective aggregation layer switches using 802.1Q trunking. Aruba recommends using tagged VLANs wherever possible to provide additional loop prevention.

The wireless module consists of one or more user VLANs depending on the security and policy model that has been implemented. For a single VLAN design, all wireless and dynamically segmented clients are assigned to a common VLAN ID with roles and policies determining the level of access each user is provided on the network. The single VLAN is extended from the respective aggregation layer switches to each physical or virtual MC cluster member. Additional VLANs can be added and extended as required. For example, a mobile first design may require separate VLANs to be assigned to wireless and dynamically segmented clients for policy compliance reasons.

A minimum of two VLANs are required between each aggregation layer switch and the respective MC cluster members. One VLAN is dedicated for management and MM communications while the second VLAN is mapped to clients. All VLANs are common between cluster members to enable seamless roaming functionality for clients. The aggregation layer switches have defined VLAN-based IP interfaces and operate as the default gateway for each VLAN. First-hop router redundancy is natively provided by the Aruba clustering and stacking architecture.

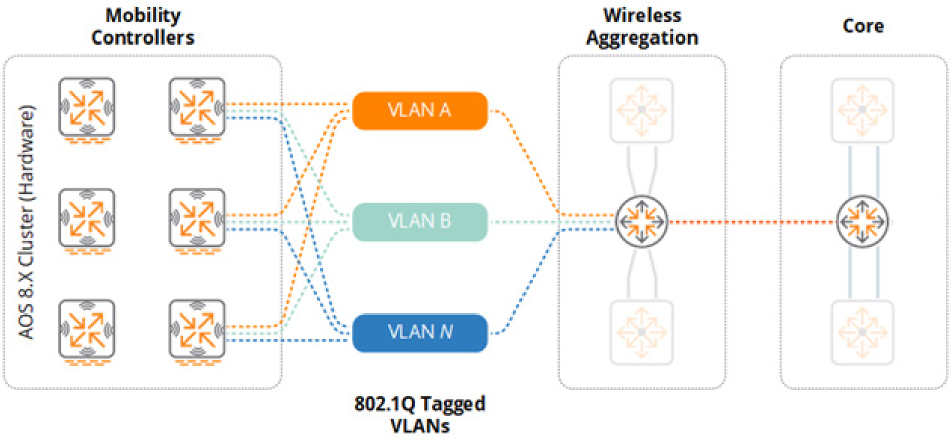

Figure 133 *Hardware MC Cluster – VLANs*

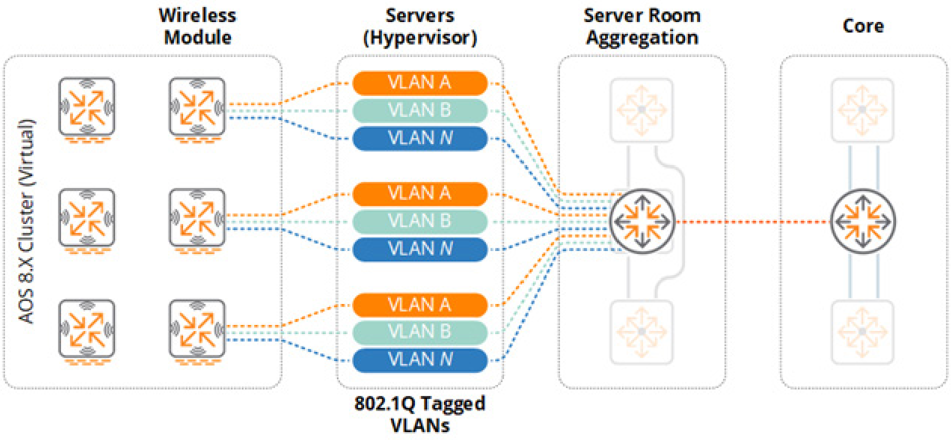

Figure 134 *Virtual MC Cluster – VLANs*

As a best practice Aruba recommends implementing unique VLAN IDs within the wireless module. This allows for the introduction of an aggregation layer in the future without disrupting the other layers within the network. This also allows for the creation of smaller layer 2 domains. Segmenting the network in this manner reduces layer 2 instability and protects the wireless module from operational changes, loops, or misconfigurations originating from other layers or network modules.

Platform Suggestions

Figure 135 provides platform suggestions for the large building scenario based on the assumption of supporting 300 APs and 6,000 concurrent clients. A good, better, and best suggestion is made based on features, performance, and scalability. These are suggestions based on the described scenario and may be altered according to the discretion of network administrators.

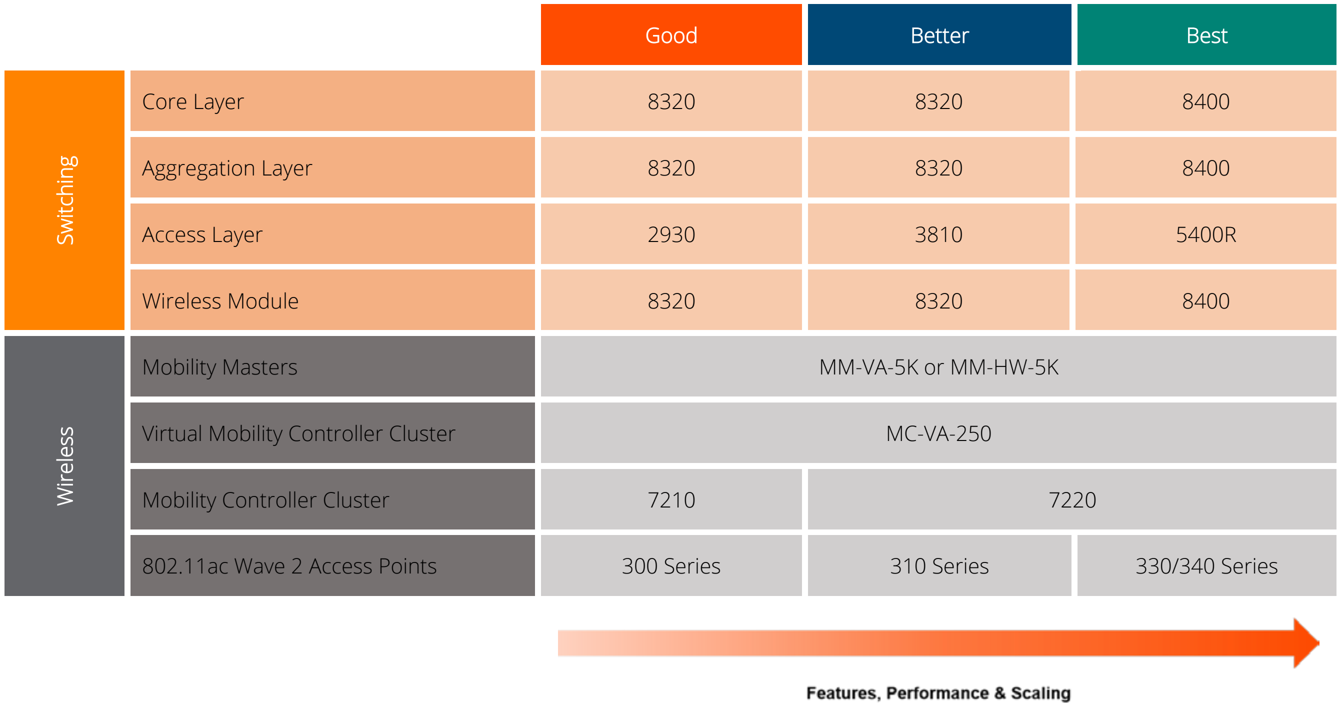

		Good	Better	Best
Switching	Core Layer	8320	8320	8400
	Aggregation Layer	8320	8320	8400
	Access Layer	2930	3810	5400R
	Wireless Module	8320	8320	8400
Wireless	Mobility Masters	MM-VA-5K or MM-HW-5K		
	Virtual Mobility Controller Cluster	MC-VA-250		
	Mobility Controller Cluster	7210	7220	
	802.11ac Wave 2 Access Points	300 Series	310 Series	330/340 Series

Features, Performance & Scaling

Figure 135 *Large Building Platform Suggestions*

Campus

The following reference design is intended for a campus deployment which consists of multiple buildings of varying size and two datacenters. Each building in the campus implements its own 2-tier or 3-tier modular network which connects to a campus backbone. The campus in this scenario needs to support 64,000 concurrent dual-stack clients and requires 6,000 802.11ac Wave 2 APs.

A key decision that needs to be taken into account for a campus deployment is where to place the MC clusters (wireless module). Due to the challenging scalability requirements, a campus of this size will generally require multiple clusters of MCs which can either be centralized in the datacenters or strategically distributed between the buildings. The clusters in both cases are managed by hardware or virtual MMs deployed across the datacenters.

Both centralized and distributed MC deployment models are valid for campus deployments with each model supporting different mobility needs. As seamless mobility can only be provided between APs managed by a common cluster, the mobility requirements are usually the determining factor influencing the cluster deployment model that is selected.

Traffic flows are another factor that needs to be taken into account when determining cluster placement. If user applications are primarily hosted in the datacenter then a centralized cluster is an appropriate choice as the wireless and dynamically segmented client sessions are terminated within the cluster. Placing the cluster closer to the applications optimizes the traffic flows. If the primary applications are distributed between buildings in the campus, a distributed MC model may be a more efficient design.

Centralized Clusters

- Permits a larger mobility domain when ubiquitous indoor and outdoor coverage is required.
- Efficient when the primary applications are hosted in the cloud or datacenter.

Distributed Clusters

- Permits smaller mobility domains such as within buildings or between co-located buildings.
- Efficient when the primary applications are distributed or workgroup based.

Scenario 1 – Centralized Clusters

The Centralized Clusters reference design is appropriate for a campus such as a corporate headquarters with two datacenters implementing centralized clusters. The campus LAN implements a high-speed layer 3 backbone that interconnects each building to both datacenters. The campus needs to support 64,000 concurrent dual-stack wireless clients across 6,000 802.11ac Wave 2 APs. Each host in this example is assigned a single global IPv6 address from a stateful DHCPv6 server. To enable roaming between large groups of buildings indoor and outdoor APs with overlapping coverage will be assigned to the same MC cluster.

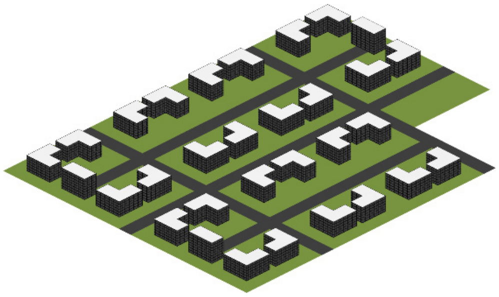

Campus Characteristics:

- 6,000 x 802.11ac Wave 2 Access Points
- 64,000 x Concurrent Dual-Stack Clients
- 2 x Datacenters with Layer 2 Extension

Figure 136 *Scenario 1 Campus Characteristics*

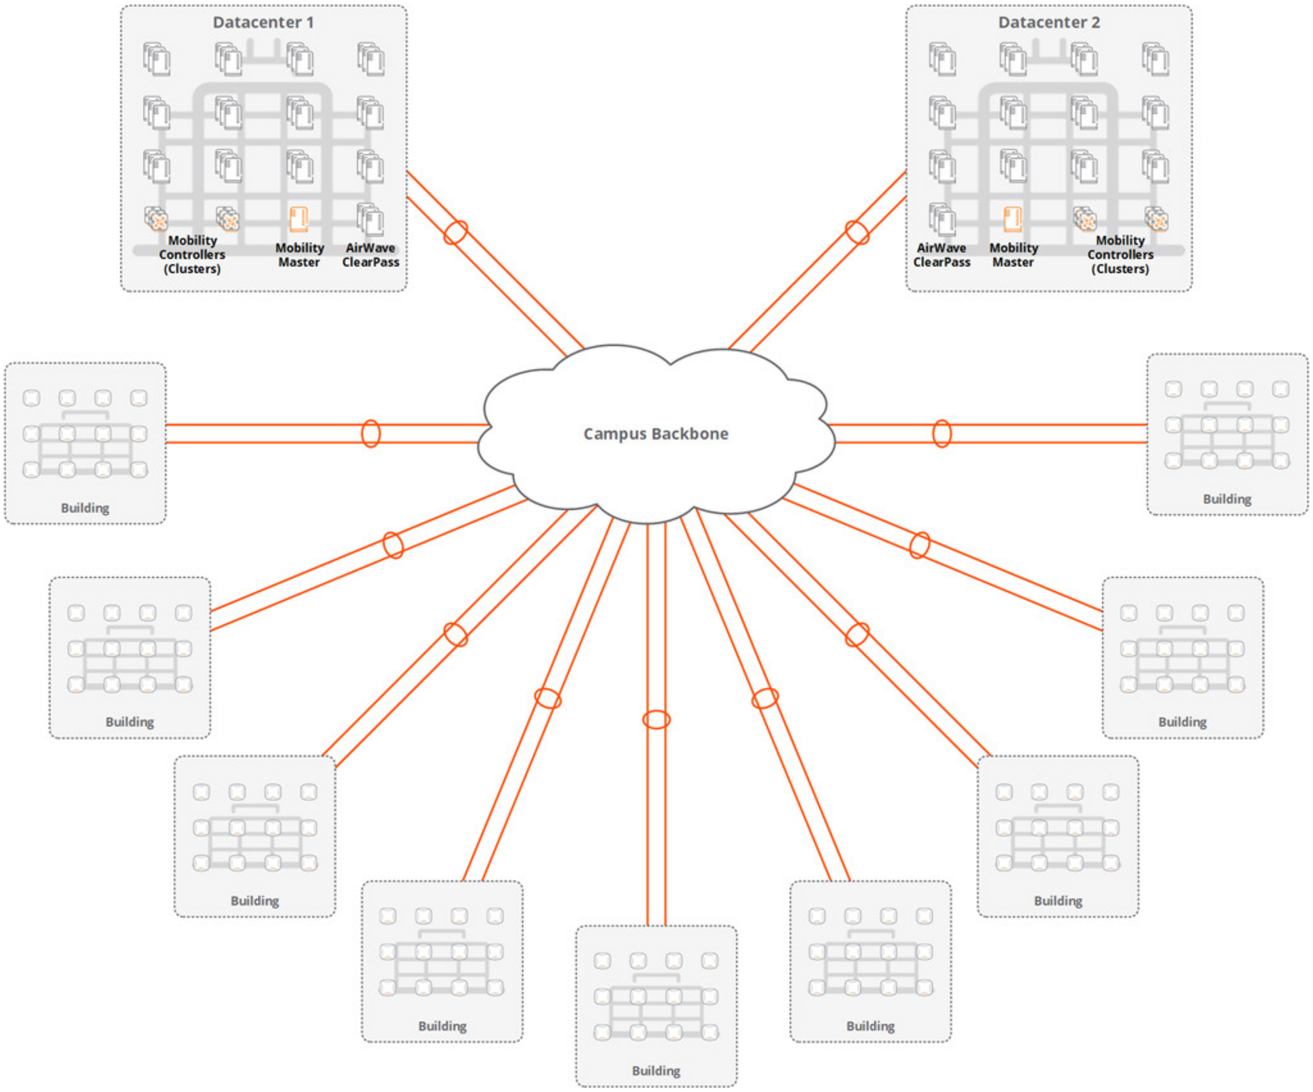

Figure 137 *Campus Modular Network Design – Centralized MC Clusters*

Wireless LAN Components

The campus in this scenario includes the MMs and clusters of MCs which are distributed across two datacenters. The number of MMs and MC clusters required to enable full redundancy will be influenced by the datacenter design. The datacenters can either support Layer 2 VLAN extensions or be separated at layer 3:

- **Layer 2 Extensions** – VLANs and their associated broadcast domains are common between datacenters.
- **Layer 3 Separation** – VLANs and their associated broadcast domains are unique for each datacenter.

When VLANs are extended between the datacenters, the MMs and MC cluster members can be split between the datacenters. Each datacenter hosts 1 MM and half of the MCs. Two MMs and two clusters of MCs are required to accommodate the AP and client counts for this scenario. To enable aggregation layer scalability as well as fault domain isolation, each cluster of mobility MCs is connected to separate Aruba 8400 series aggregation layer switches. Each aggregation layer can accommodate up to 32,000 IPv4 and 64,000 IPv6 host addresses.

Component	Description	Notes
Aruba MM (MM)	Hardware or Virtual Appliances	2 required
Aruba MCs	Hardware or Virtual Appliances	2 clusters
Aruba Access Points	802.11ac Wave 2 Access Points	6,000 required
Aruba Airwave	Hardware or Virtual Appliance	Recommended
Aruba ClearPass	Hardware or Virtual Appliance	Recommended

Table 25 *Wireless LAN Components – Layer 2 Extension*

A different approach is required when datacenters are separated at layer 3. To support the AP and client counts and maintain full redundancy an active/standby model is implemented. In such a design each datacenter hosts an equal quantity of MMs and MCs:

1. **MMs** – Two MMs are hosted per datacenter implementing layer 2 and layer 3 master redundancy. Layer 2 master redundancy is provided between MMs within each datacenter while layer 3 master redundancy provides redundancy between datacenters.
2. **MC Clusters** – Two clusters of MCs are hosted per datacenter. The APs are configured with a primary LMS and backup LMS to determine their primary and secondary cluster assignments. Fast failover is provided within the primary cluster while a full bootstrap is required to failover between the primary and secondary clusters.

Each cluster of MCs is connected to separate Aruba 8400 series aggregation layer switches for aggregation layer scaling and fault domain isolation. Each aggregation layer switch can accommodate up to 32,000 IPv4 and 64,000 IPv6 host addresses. Each datacenter is separated at layer 3 requiring four wireless modules and wireless aggregation layers to accommodate an individual datacenter failure.

Component	Description	Notes
Aruba MM (MM)	Hardware or Virtual Appliances	4 required (L3 redundancy)
Aruba MCs	Hardware or Virtual Appliances	4 clusters (2 per datacenter)
Aruba Access Points	802.11ac Wave 2 Access Points	6,000 required
Aruba Airwave	Hardware or Virtual Appliance	Recommended
Aruba ClearPass	Hardware or Virtual Appliance	Recommended

Table 26 *Wireless LAN Components – Layer 3 Separation*

Roaming Domains

Seamless mobility is provided between APs managed by a common cluster in an ArubaOS 8 architecture. Each wireless and dynamically segmented client is assigned a primary UAC and S-UAC cluster member to provide fast failover in the event of a cluster member failure or live upgrade. Two clusters of MCs are required in order to ensure adequate scalability.

It is important to take into account that seamless roaming can only occur between APs managed by the same cluster. The following considerations need to be made:

1. APs in the same building must be managed by the same cluster. This ensures wireless client sessions are not interrupted as the clients roam within the building.
2. Indoor and outdoor APs in co-located buildings with overlapping coverage must be managed by the same cluster. This ensures client sessions are not interrupted as the clients roam within a building or between buildings.

APs in buildings that are geographically separated and do not have overlapping coverage can be distributed between clusters as required with attention being made to ensure AP and client capacity is as evenly distributed as possible:

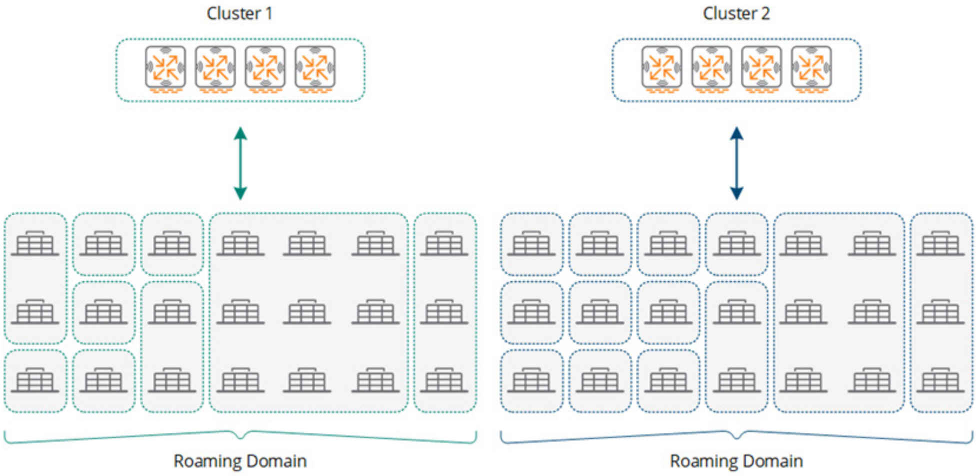

Figure 138 *Roaming Domains*

If the campus deployment supports both wireless and dynamically segmented clients it may be beneficial to deploy separate clusters.

Redundancy

In the Centralized Clusters scenario the datacenters are located in separate buildings which are connected to the campus backbone. The datacenters are interconnected using high-speed links ensuring there is adequate bandwidth capacity available to support the hosted applications and services that are hosted in each datacenter.

For a dual datacenter design the MMs and MC clusters are distributed between both datacenters. The wireless components can be deployed using several strategies to achieve redundancy which is depend on the datacenter design:

- **Layer 2 Extension** – If VLANs are extended between datacenters, the MMs and MC cluster members can be split between the datacenters. Each datacenter will host 1 MM and half of the cluster members.
- **Layer 3 Separation** – The MMs and MC cluster members are duplicated in each datacenter.

Layer 2 Extension

The layer 2 datacenter redundancy model is relatively straightforward as it operates in the same manner as a single datacenter deployment model. Each datacenter hosts a MM and half of the MCs of each cluster. The MMs are configured for L2 redundancy while AP and client load-balancing as well as fast failover functionality is provided by each cluster:

- Aruba MM (MM):
  - ➢ Two hardware or virtual MMs (one per datacenter)
  - ➢ L2 master redundancy (active/standby)
- Hardware MCs (MCs):
  - ➢ Two clusters of hardware MCs
  - ➢ Cluster members equally distributed between datacenters
- Access Points
  - ➢ AP Master pointing to the cluster's VRRP VIP address
  - ➢ Fast failover using cluster built-in redundancy
  - ➢ AP cluster assignment based on roaming requirements for each building

APs and clients will be load-balanced and distributed between cluster members residing in each datacenter by default. With the Centralized Cluster Campus design it is possible that APs and clients within a building will be assigned to cluster members in different datacenters.

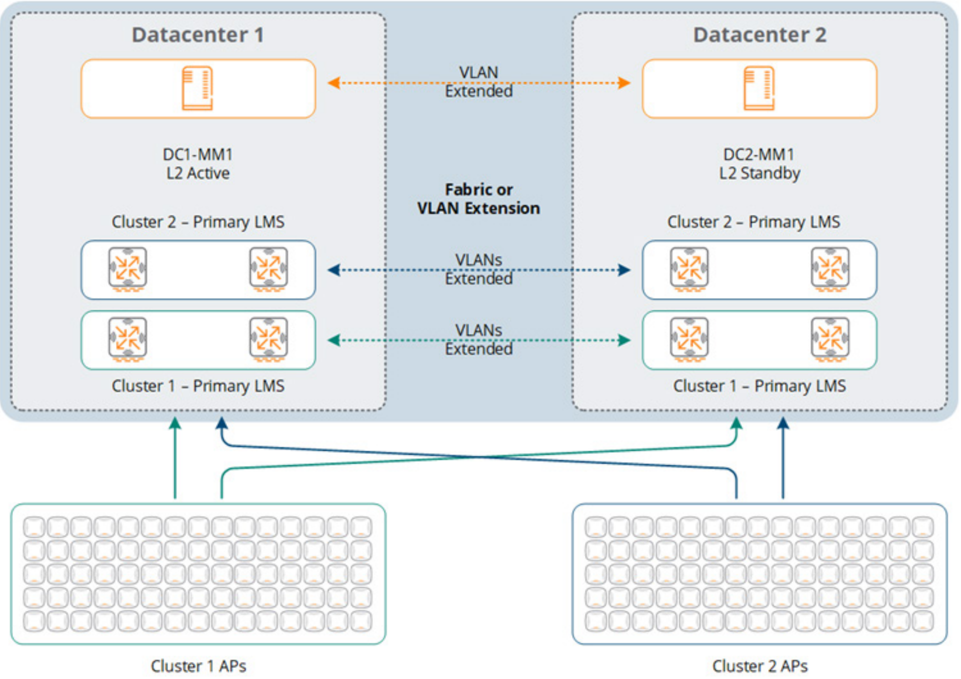

Figure 139 *Redundancy – Layer 2 Extension*

Layer 3 Separated

The layer 3 Separated datacenter redundancy model differs from Layer 2 Extension by duplicating the MMs and clusters within each datacenter. Each datacenter hosts two MMs and two clusters of MCs. The MMs are configured for L2 redundancy within the datacenter and L3 redundancy between datacenters. The APs within each building are assigned a primary and backup cluster using the primary and backup LMS. AP and client fast failover functionality is provided within each cluster while a full bootstrap is required to provide failover between clusters:

- Aruba MM (MM):
  - Four hardware or virtual MMs (two per datacenter)
  - L2 master redundancy (Active/Standby)
  - L3 master redundancy (Primary/Secondary)
- Hardware MCs (MCs):
  - Four clusters of hardware MCs (Primary/Secondary)
  - Cluster members duplicated between datacenters
  - Primary clusters alternating between datacenters
- Access Points
  - Primary and Backup LMS using the primary and secondary cluster VRRP VIP addresses
  - Fast failover using cluster built-in redundancy
  - Bootstrap failover between primary and secondary clusters

- AP cluster assignment based on roaming requirements for each building

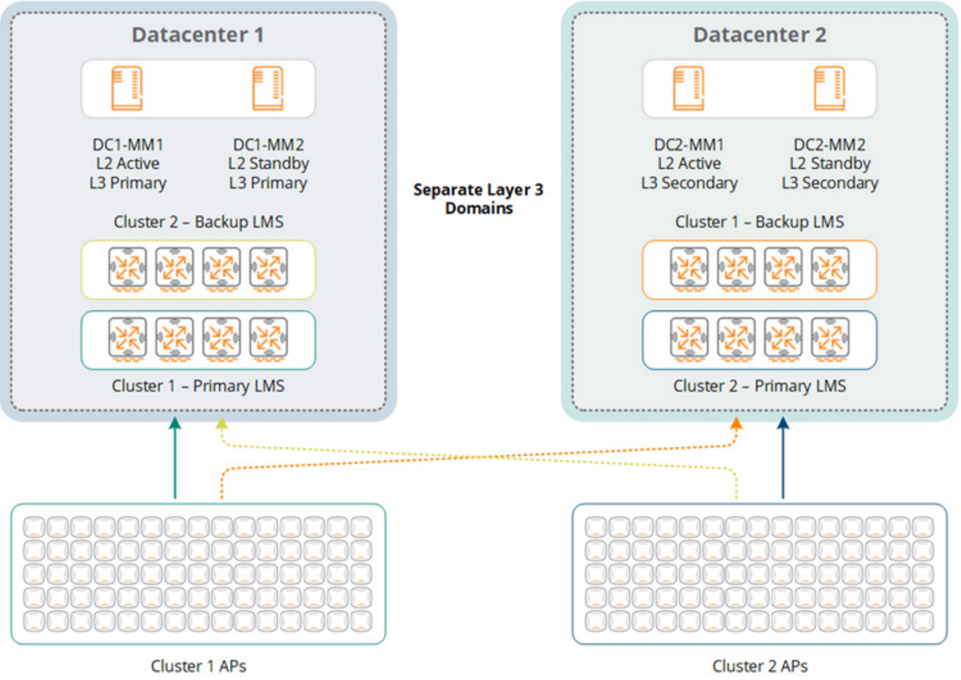

Figure 140 *Redundancy – Layer 3 Separation*

Scalability

Scalability is the primary concern for this campus scenario and the inclusion of a secondary datacenter with the datacenter deployment model can pose a challenge. Considerations had to be made for both the datacenter aggregation layer and the MC cluster design to accommodate the network growth and redundancy requirements.

Datacenter Aggregation Layer

Both datacenter deployment models require clusters of MCs that are connected to their respective datacenter aggregation layers. Two clusters of MCs are required to accommodate 64,000 concurrent dual-stack hosts with each supporting up to 32,000. Each IPv6 host in this example is assigned a single global IPv6 address.

Each cluster is connected to a separate Aruba 8400 series wireless aggregation layer switch due to the high number of clients that must be supported. This recommendation applies to both layer 2 extended and layer 3 separated datacenter designs:

- **Layer 2 Extension** – Requires two datacenter aggregation layers which are split between datacenters. Each wireless aggregation layer supports one cluster of MCs.
- **Layer 3 Separated** – Requires two datacenter aggregation layers per datacenter. Each wireless aggregation layer is connected to a primary or secondary cluster of MCs.

This datacenter aggregation layer design ensures that a single aggregation layer never exceeds more than 64,000 IPv4 or IPv6 host addresses during normal operation as well as provides sufficient capacity to continue normal operations in the event of a datacenter failure:

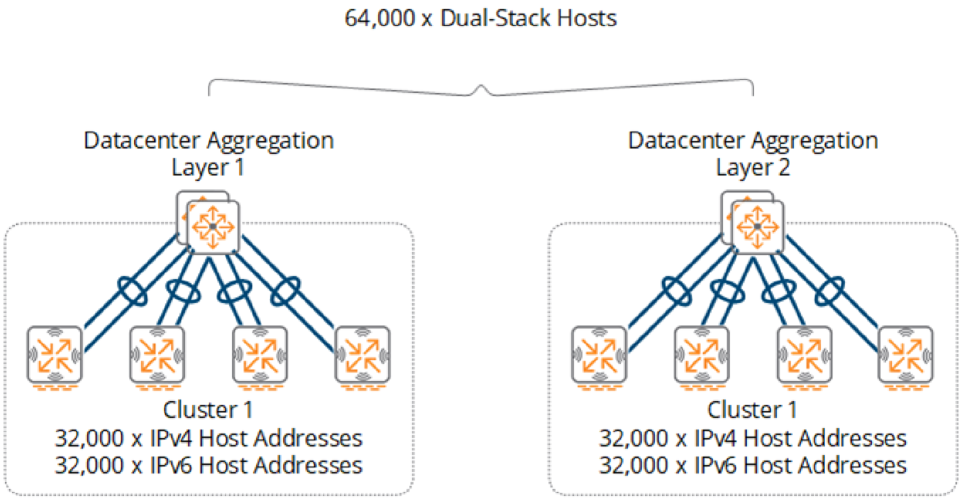
64,000 x Dual-Stack Hosts
Datacenter Aggregation Layer 1
Datacenter Aggregation Layer 2
Cluster 1
32,000 x IPv4 Host Addresses
32,000 x IPv6 Host Addresses
Cluster 1
32,000 x IPv4 Host Addresses
32,000 x IPv6 Host Addresses

Figure 141 *Datacenter Wireless Aggregation Layer 3 Separated Design*

MC Clusters

Scalability for each cluster is provided by selecting the appropriate controller model and determining the correct number members per cluster. The throughput capabilities of the chosen MC model are also a factor as each model has different switching capacities and physical interfaces. For this campus scenario the 7200 series controllers are recommended with each cluster consisting of four MCs.

While the VMCs can be selected for a campus deployment Aruba recommends hardware MCs be deployed for throughput and performance reasons. The fact that the hardware is dedicated guarantees that a specific level of performance can be provided.

MM

For the Centralized Cluster Campus design Aruba recommends implementing the MM-HW-10K or MM-VA-10K MM. Data switching throughput is not as big of a concern as with the MC clusters so either hardware or virtual MMs can be deployed.

The MM selected for this design should scale to support 10,000 APs, 100,000 clients, and 1,000 MCs. Implementing this level of capacity will ensure adequate support for the AP, client, and MCs while providing additional room for future growth. Additional clients and APs can be added as the campus grows by adding additional aggregation layers and MC clusters.

Scaling beyond 64,000 dual-stack clients for a centralized deployment model can be achieved by deploying additional MC clusters within the datacenter. In an ArubaOS 8 deployment a MM can scale to support up to 100,000 clients, 10,000 APs, 1,000 MCs. Additional scalability can be achieved by deploying additional MMs and MC clusters.

Virtual LANs

In a centralized cluster design the datacenter aggregation layer terminates all the VLANs from the MC cluster members. The datacenter architecture is what determines the VLAN design. In both designs the VLANs are extended from the MCs to their respective datacenter aggregation layer switches using 802.1Q trunking. The primary difference between the designs being the number of VLANs that are required.

Layer 2 Extension

When VLANs are extended between datacenters, each cluster implements its own unique VLAN IDs and broadcast domains. Each cluster consists of one or more user VLANs depending on the VLAN model that has been implemented. For a single VLAN design, all wireless and dynamically segmented clients are assigned to a common VLAN ID with roles and policies determining the level of access each user is provided on the network. Each cluster implements unique VLAN IDs.

The user VLANs are extended from the aggregation layer switches to each MC cluster member. A minimum of two VLANs are required between the datacenter aggregation layers and each MC cluster member. One VLAN is dedicated for management, cluster, and MM communications while the additional VLANs are mapped to clients. The VLANs are common between cluster members split between the datacenters to enable seamless mobility. The datacenter aggregation layer switches have VLAN based IP interfaces defined and operate as the default gateway for each VLAN. First-hop router redundancy is natively provided by VRRP or the Aruba clustering architecture.

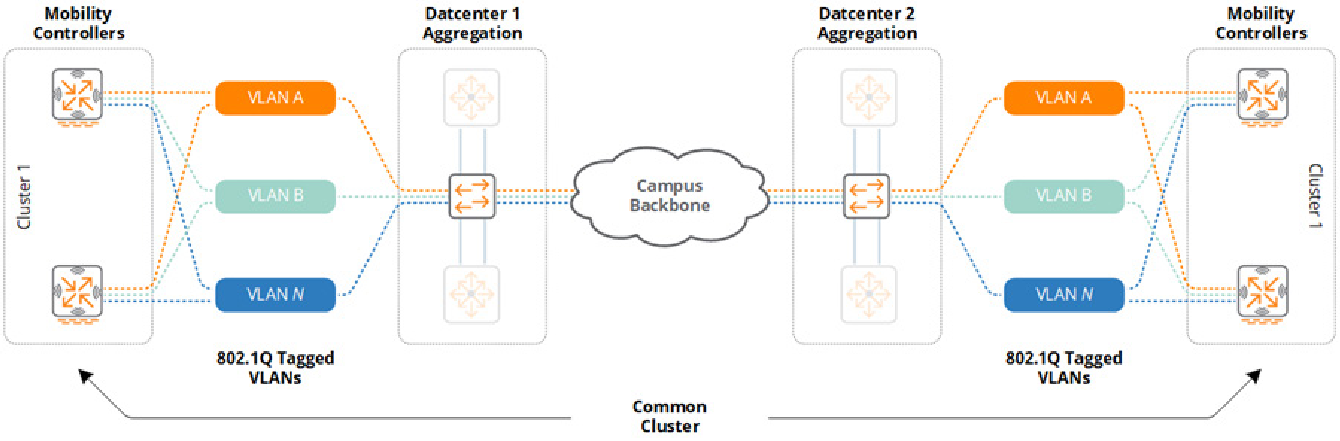

Figure 142 *Wireless and Dynamically Segmented Client VLANs – Layer 2 Extension*

Layer 3 Separation

When the datacenters are separated at layer 3, the VLANs are unique for each datacenter. The primary and secondary clusters in each datacenter each requiring their own unique VLAN IDs and broadcast domains. Each cluster consists of one or more user VLANs depending on the VLAN model that has been implemented. For a single VLAN design, all wireless and dynamically segmented clients are assigned to a common VLAN ID with roles and policies determining the level of access each user is provided on the network. Each cluster implements unique VLAN IDs.

The user VLANs are extended from the aggregation layer switches to each MC cluster member. A minimum of two VLANs are required between the datacenter aggregation layers and each MC cluster member. One VLAN is dedicated for management, cluster, and MM communications while the additional VLANs are mapped to clients. The VLANs are common between cluster members in each datacenter to permit seamless mobility. The datacenter aggregation layer switches have VLAN based IP interfaces defined and operate as the default gateway for each VLAN. First-hop router redundancy is natively provided by VRRP or the Aruba clustering architecture.

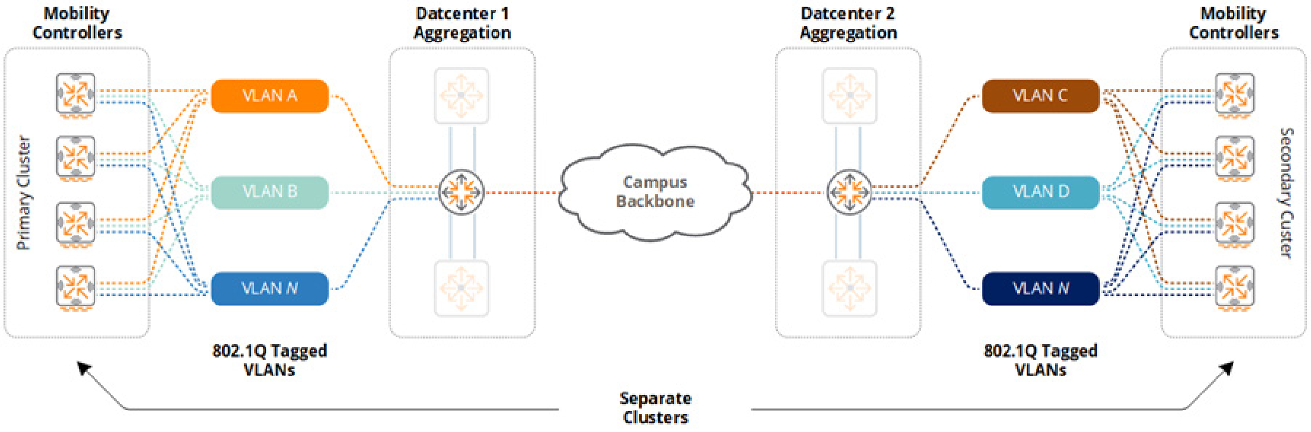

Figure 143 *Wireless and Dynamically Segmented Client VLANs – Layer 3 Separation*

One key difference between the two datacenter designs is client VLAN assignment and broadcast domain membership during a datacenter failure. While both models offer full redundancy, only the layer 2 VLAN extension model offers fast failover in the event of a datacenter outage:

1. **Layer 2 Extension** – Impacted clients maintain their VLAN ID and IP addressing after a datacenter failover. The APs, Aruba switches, and clients are assigned to a new cluster member in their existing cluster in the remaining datacenter.
2. **Layer 3 Separated** – Impacted clients are assigned a new VLAN ID and IP addressing after a datacenter failover. The APs, Aruba switches, and clients will be assigned to a secondary cluster member in the remaining datacenter.

Platform Suggestions

Figure 144 provides platform suggestions for the centralized cluster campus deployment scenario that supports 6,000 APs and 64,000 concurrent clients. Where appropriate a good, better, and best suggestion is made based for feature, performance, and scalability requirements. These are suggestions based on the described scenario and may be altered according to the discretion of network administrators.

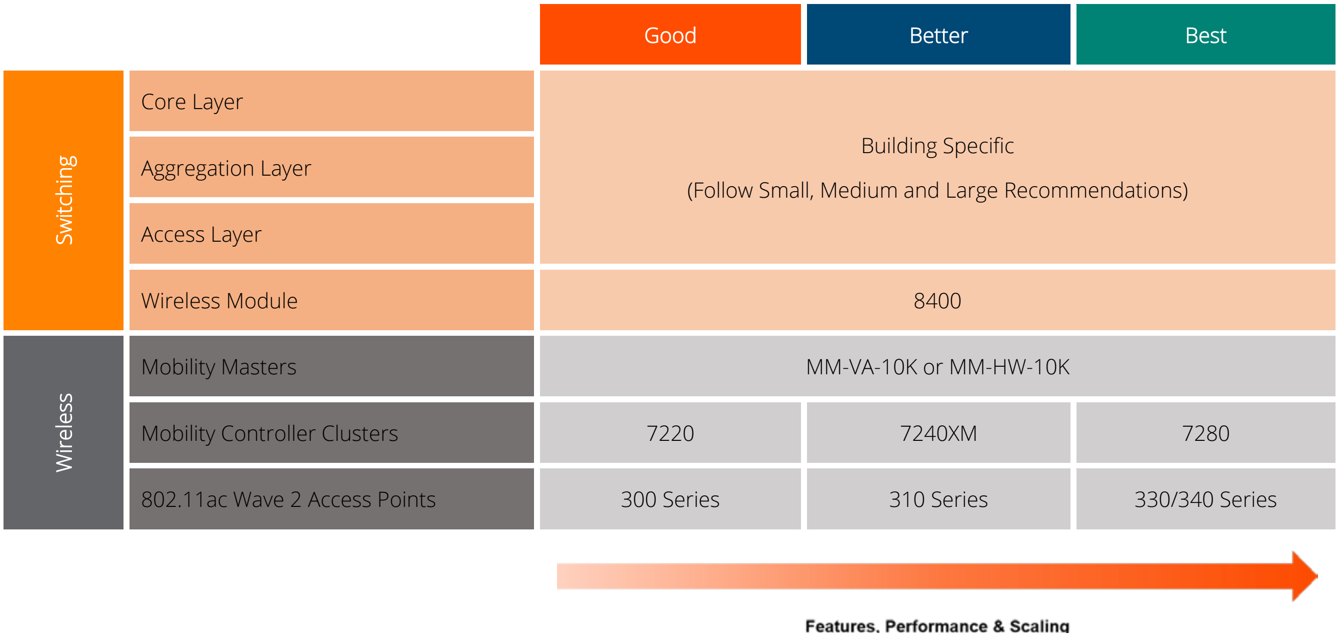

		Good	Better	Best
Switching	Core Layer	Building Specific (Follow Small, Medium and Large Recommendations)		
	Aggregation Layer			
	Access Layer			
	Wireless Module	8400		
Wireless	Mobility Masters	MM-VA-10K or MM-HW-10K		
	Mobility Controller Clusters	7220	7240XM	7280
	802.11ac Wave 2 Access Points	300 Series	310 Series	330/340 Series

Features, Performance & Scaling

Figure 144 *Centralized Cluster Campus Building Platform Suggestions*

As each building in the campus can vary in size each one will require its own 2-tier or 3-tier hierarchical network design. For that reason switching suggestions for core, aggregation, and access layers are not provided as these selections will be unique for each building. The individual building selections should be made following the small, medium, and large suggestions highlighted in the previous sections.

Scenario 2 – Distributed Clusters

The following reference design is for a large campus such as a university with 285 buildings distributed over a 900 acre site. Each building will implement its own 2-tier or 3-tier modular network design that connects to a common campus backbone. The university has 20,000 faculty, staff, and students with IPv4 and IPv6 clients. The university has deployed 3,500 802.11ac Wave 2 APs to provide adequate coverage.

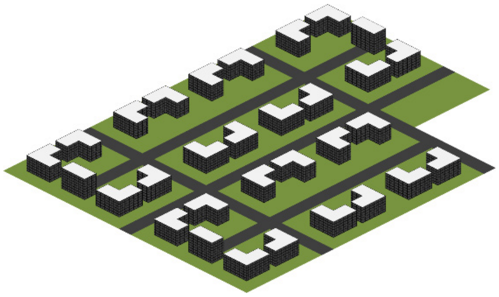

Campus Characteristics:

- 3,500 x 802.11ac Wave 2 Access Points
- 40,000 x Concurrent Clients (Native IPv4 and/or Dual-Stack)
- 1 x Datacenter

Figure 145 *Scenario 2 Campus Characteristics*

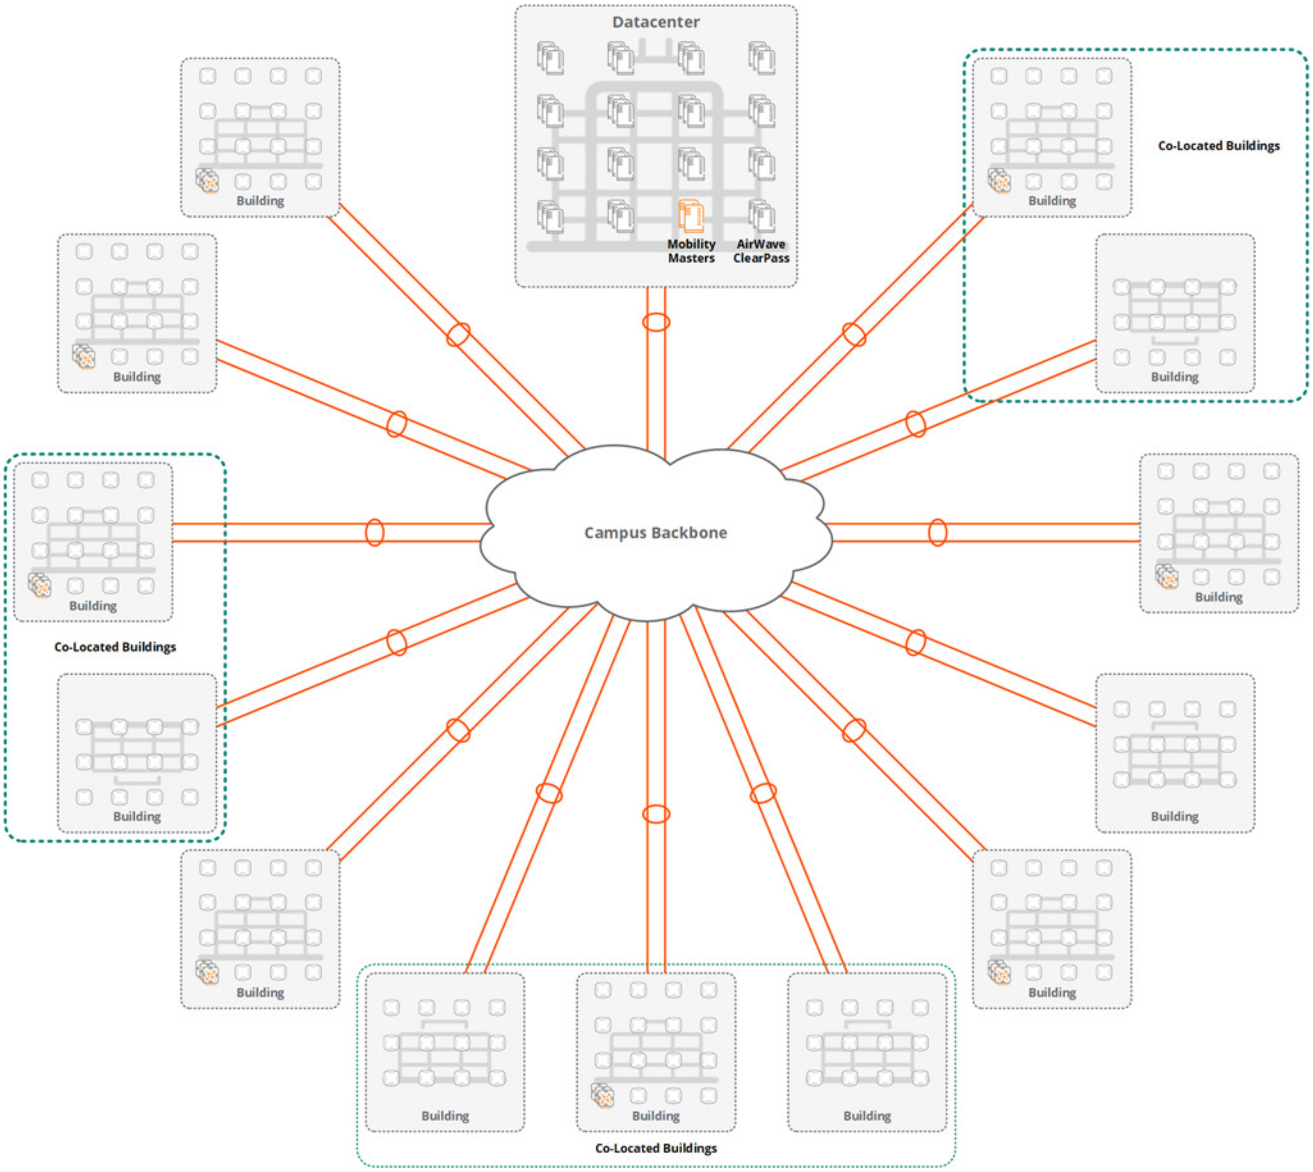

Figure 146 *Distributed Clusters Campus Architecture*

Wireless LAN Components

The campus in this scenario includes the MMs deployed in a datacenter and clusters of MCs that are distributed between buildings. The campus in this scenario includes a single datacenter, however multiple datacenters may exist in production deployments. If multiple datacenters exist, then the campus reference architecture detailed in previous chapters provides details for the MM deployments options that can be selected for layer 2 and layer 3 datacenter deployment models.

Unlike the previous campus example, the MC clusters described in this section are distributed between individual buildings rather than deployed in the datacenter. This means that wireless and dynamically segmented traffic is terminated within the buildings rather than the datacenter. Roaming can only occur within a cluster of MCs, therefore APs in co-located buildings require overlapping coverage and are serviced by a cluster of MCs strategically deployed in one of the buildings. APs and clients in isolated buildings need to be serviced by their own cluster of MCs.

The modular network design and MC cluster placement recommendations for each building in the campus follow the same recommendations provided for the small, medium, and large office reference designs. The MC clusters connect to their respective layers depending on the size of the building. As with the previous recommendations, a wireless aggregation layer is recommended when the wireless and dynamically segmented client count exceeds 4,096.

As the building grows in size the number of APs and hosts will vary. The MC clusters are customized for each building or co-located buildings to meet the specific AP, client, and throughput requirements. For ease of deployment, troubleshooting, and repair it is recommended to standardize common models of MCs for small, medium, and large buildings. A design may include specifying two or three different controller models depending on the range of building sizes requiring support.

Component	Description	Notes
Aruba MM (MM)	Hardware or Virtual Appliances	2 required
Aruba MCs	Hardware or Virtual Appliances	Varies
Aruba Access Points	802.11ac Wave 2 Access Points	3,500 required (distributed)
Aruba Airwave	Hardware or Virtual Appliance	Recommended
Aruba ClearPass	Hardware or Virtual Appliance	Recommended

Table 27 *Wireless LAN Components*

Roaming Domains

Seamless mobility is provided between APs managed by a common cluster in an ArubaOS 8 architecture. Each wireless and dynamically segmented client is assigned a primary UAC and secondary S-UAC cluster member to provide fast failover in the event of a cluster member failure or live upgrade.

This campus design includes both standalone and co-located buildings. Roaming is provided within each building as well as strategically between co-located buildings where overlapping coverage is provided. Co-located buildings provide indoor as well as outdoor coverage and enable roaming as faculty and students move between the buildings:

- **Standalone Buildings** – Serviced by a cluster of MCs deployed within each building. When necessary APs in small buildings are serviced by an MC cluster in a neighboring building.
- **Co-Located Buildings** – Serviced by a cluster of MCs strategically deployed in one of the co-located buildings. Each cluster services APs across two or more buildings.

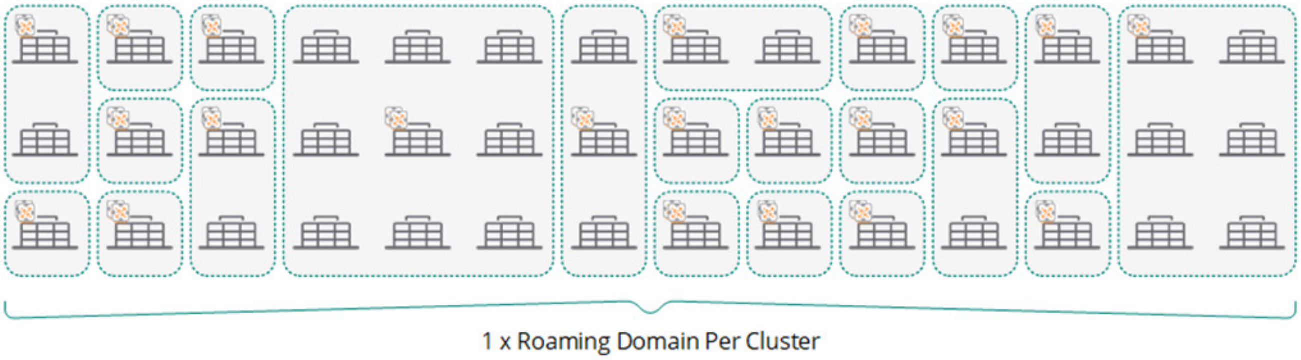

Figure 147 *Roaming Domains*

Redundancy

For this scenario the MMs are deployed within the datacenter and connect directly to separate datacenter aggregation switches. Redundancy within each building is provided by the modular network design and clusters of MCs. The MCs are deployed following the same recommendations provided for the small, medium, and large office reference designs:

- Aruba MM (MM):
  - ➢ Two hardware or virtual MMs
  - ➢ L2 master redundancy (Active/Standby)
- Hardware MCs (MCs):
  - ➢ Multiple clusters of hardware MCs
  - ➢ Minimum of two cluster members
- Virtual MCs (MCs):
  - ➢ Multiple clusters of virtual MCs
  - ➢ Minimum of two cluster members
- Access Points
  - ➢ AP Master pointing to the cluster's VRRP VIP address
  - ➢ Fast failover using cluster built-in redundancy

Additional redundancy between clusters can be achieved if desired by implementing the backup LMS option. This will allow APs in a building to failover to an alternative designated cluster in the event of a cluster or wireless aggregation layer failure. If such an event was to occur the APs will perform a full bootstrap to failover to the alternate cluster which will impact users. The alternate cluster and aggregation layer must also be scaled accordingly to accommodate the AP and client counts.

Scalability

The primary scalability concern for Distributed Cluster Campus scenario is MM scaling. For this campus design the total number of APs and clients need to be accommodated in addition to the total number of MCs which are distributed between buildings. The 285 buildings in this scenario will be serviced by 180 clusters which each have a minimum of two members. Clusters in some of the larger buildings may contain three or four cluster members as required.

For this campus design, Aruba recommends implementing the MM-HW-5K or MM-VA-5K MM. Either hardware or virtual MMs can be deployed since the number of distributed MCs is a key concern. The MM selected for this design needs to scale to support 5,000 APs, 50,000 clients, and 500 MCs. This will provide adequate capacity to support the AP, client, and MC counts while providing additional room for future growth. If a specific campus design requires additional MCs then the MM-HW-10K or MM-VA-10K MM can be selected. These MMs support up to 1,000 MCs each.

Virtual LANs

For a distributed cluster design the building core or wireless aggregation layer terminates all the VLANs from each building's wireless module. The wireless and dynamically segmented client VLANs are extended from the MCs to their respective building switches using 802.1Q trunking.

The wireless module consists of one or more user VLANs depending on the architectural model that is implemented. For a single VLAN design, all wireless and dynamically segmented clients are assigned to a common VLAN ID with roles and policies determining the level of access each user is provided on the network. The single VLAN is extended from the respective aggregation layer switches to each physical or virtual MC cluster member. Additional VLANs can be added and extended as required. E.g., a particular first design may require separate VLANs to be assigned to wireless and dynamically segmented clients for policy compliance purposes.

A minimum of two VLANs are required between the core or wireless aggregation layer switches in each building and each MC cluster member. One VLAN is dedicated for management and MM communications while the additional VLANs are mapped to clients. The VLANs are common between cluster members to permit seamless mobility within each building.

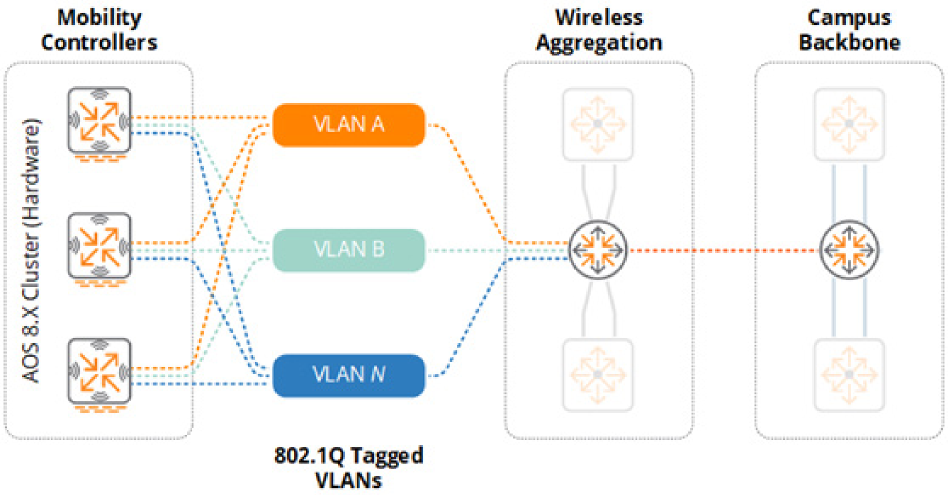

Figure 148 *Hardware MC Cluster – VLANs*

Each building may implement common VLAN IDs or unique VLAN IDs as required. Each building is layer 3 separated from the other buildings in the campus. This means that the VLAN IDs can be reused which simplifies the WLAN and dynamically segmented client deployment. However, each VLAN will require its own IPv4 and IPv6 subnet assignments.

Platform Suggestions

The distributed campus scenario requires switching and wireless components to be selected based on the specific requirements for each building. The component selection for each building should be based on the small, medium, and large suggestions highlighted in the previous sections. Each building in the campus will implement a 2-tier or 3-tier hierarchical network design with appropriate selections to meet each building's wired and wireless connectivity, performance, and redundancy requirements.

As previously mentioned, Aruba recommends standardizing on common models of MCs for the small, medium, and large buildings as doing so will simplify deployment, troubleshooting, and repair. A specific campus design may standardize on a common model of MC for all buildings or one model for each building size. The number of cluster members for each building should be adjusted to meet each building's redundancy and performance needs.

To support this distributed campus scenario, Aruba recommends the MM-VA-5K or MM-HW-5K which can scale to accommodate 5,000 Aps, 50,000 clients, and 500 MCs. The suggested MM models can meet the initial requirements to support 3,500 APs and 40,000 concurrent clients while allowing for future network scalability. Larger MMs such as the MM-VA-10K or MM-HW-10K are available to support larger distributed campuses if required. MM-VA-10K or MM-HW-10K are both capable of supporting 10,000 APs, 100,000 clients, and 1,000 MCs.

Migration to ArubaOS 8

Migration of Aruba deployments from ArubaOS 6 to ArubaOS 8 involves a few more steps and precautions than performing a simple controller image upgrade. This chapter covers topics including migration methods, best practice recommendations on when to choose a particular method over another, and outlines how typical ArubaOS 6 network topologies can be migrated to ArubaOS 8. Depending on the ArubaOS 6 topology, migration can either be performed manually or by using the Migration Tool.

NOTE

There is no automatic migration from 8.x standalone or MC Master to ArubaOS 8 MM

Migration Strategies

Migration Tool

The Migration Tool is a VM-based server that can be used to migrate an ArubaOS 6 deployment to ArubaOS 8. The Migration Tool GUI is used to supply IP addresses, credentials, and required roles for all of the controllers requiring migration. The tool then communicates with the MM, controllers, and VMware/KVM (if orchestrating), takes the required controller backups, upgrades the images to ArubaOS 8, and configures the controllers for communication with the MM.

Benefits

- Preserves existing WLAN configuration elements from ArubaOS 6 during migration which saves the time and effort required to possibly reconstruct them
- Multiple ArubaOS 6 deployment topologies (e.g. multiple master-locals) can be collapsed under a single MM
- Automates configuration backups, image downloads/upgrades, and license migration
- Supports both single and multi-phase migration approaches. E.g., the existing master can be used as a source for the MM configuration while other controllers are being migrated under the MM and the master itself may be migrated at a later stage
- Supports orchestration with both VMware and KVM

Supported Topologies for Migration

- Migrating Master-Local setup to MM or Master Controller Mode
- Migrating All-Master setup to MM
- Migrating to a stand-alone controller

NOTE

The Migration Tool retains all aspects of the previous configuration including old and unused configuration elements. For additional details on the Migration Tool, please refer to the ArubaOS Migration Guide.

Manual Migration

Manual migration involves taking backups from all controllers, rebuilding them by individually upgrading each one to ArubaOS8, and performing the initial setup to convert them to MM-managed controllers or standalone controllers. Conversion to a MM-managed controller requires having the MM installed, configured, and ready to accept controller connections. A manual migration may also be performed by standing up a MM in parallel, building the configuration, and then moving one controller at a time.

Benefits

- While the Migration Tool allows for a safe and easy migration path for ArubaOS 6 deployments with standard network configuration, the topology and feature benefits introduced by ArubaOS 8 may require new configuration elements post migration. In such cases, it may be more effective to bring up a MM in parallel to your existing ArubaOS 6 deployment and manually reconfigure elements of your WLAN to accommodate the new features
- Manual migration allows for small, incremental changes to be performed and tested, while allowing the existing network to keep running during the migration process
- Existing topologies may contain obsolete or deprecated features and manual migration allows for alternatives to be planned and configured accordingly
- If the existing configuration is very complex (e.g. numerous static GRE/IPsec tunnels, large mesh deployments, complex static channel plans, AP-specific settings, etc.), it may be more effective to migrate manually

Supported Topologies

Since manual migration involves individually preparing each controller for migration, a number of migration topologies are possible. For recommended topologies, please refer to the section on Migrating Different ArubaOS 6 Topologies later on in this chapter. Examples of topologies that can be migrated manually include the following:

- Master-Local to MM or MC Master (MCM)
- All-Masters to MM
- Master-Branch (BOC) to MM
- Master/standby-master to standalone/standby-standalone
- Standalone to MM, MC Master, or standalone
- Migrating to a standalone controller

Migration Tool vs. Manual Migration

There are no strict rules for determining whether to use the Migration Tool over manual migration or vice versa. The choice really comes down to context and the existing deployment's complexity. Existing deployments with a high degree of complexity may require a manual rebuild. The Migration tool may or may not be able to handle uncommon or very complex configuration elements (which may eventually get addressed over time with newer Migration Tool releases), so existing complexity will always need to be weighed against the capabilities of the tool. The points below provide a rough set of guidelines which can aid in determining which option is most appropriate for a given deployment:

- If migration is being performed primarily to support new topologies (or features that require moving to newer topologies) then it may be better to perform a manual migration
- The Migration Tool is suitable for collapsing multiple individual topologies under a single MM
- The Migration Tool is also suitable for deployments with basic Wi-Fi and guest features. However, deployments that use custom Captive Portal web pages and images may have to be rebuilt post migration

Migration Best Practice Recommendations

- Always backup everything in your existing topology prior to migration
- Always test the desired migration approach in a lab environment prior to migrating production deployments. The Migration Tool allows migration of local controllers without having to migrate the existing master controller which facilitates lab testing
- When lab testing, exercise caution when testing license migration via the "My Networking Portal" (MNP). There is a limitation of three license migrations
- If using Activate, make sure to update the ZTP settings if required, prior to migration

Migration Caveats

- Migration to ArubaOS 8 is not supported for 6000/M3, 3000, or 600 controller platforms. The prerequisite for migration is having 7000 and/or 7200 series controllers
- The 7000/7200 controller requirement for the master still applies for scenarios where only the local controllers need to be migrated and not the master. If the deployment has a master that is not from the 7000/7200 series, then either the master will need to be temporarily replaced with a 7000/7200 series controller before using the Migration Tool or the devices will require manual migration
- All controllers that are being migrated need to have their licenses in the same *My Networking Portal* account, otherwise license migration will not work. This applies to both manual Migration as well as the Migration Tool.

- All controllers that are being migrated must have a controller-IP and default gateway configured
- Deployments using custom captive portal web pages and images may have to be rebuilt after migration
- Only the 7030 and greater controller models can run as a MC Master in MC Master mode
- The 7024 and lower models can only be converted to MM managed, standalone, or MC Master managed controllers
- In scenarios where existing master-local deployments need to be migrated to a MC Master managed deployment the MC Master cannot terminate APs. If APs were previously terminating on the master, they will need to be accommodated either on the locals that moved under the MC Master or on a new controller
- ArubaOS 8 does not currently have a migration path to take a standalone or MC Master managed controller and bring it under a MM
- If a controller is repeatedly upgraded or downgraded between ArubaOS 6 and ArubaOS 8, subsequent migrations may fail due to temp files created on the controller that will cause a pre-migration check failure. If repeated upgrades or downgrades are required, the best solution is to capture a flash backup before the upgrade, then restore the backup before second or subsequent upgrades

General Migration Requirements

Controllers

The table below provides recommendations on the minimum controller model required for ArubaOS 8 migration:

Legacy ArubaOS 6 Controllers	APs	Clients	Minimum 7000/7200 Platform for ArubaOS 8 Migration	APs	Clients
6000/M3	512	8K	7210-7240	512-2K	16K-32K
3600	128	8K	7205	256	8K
3400	64	4K	7030	64	4K
3200	32	2K	7010	32	2K
651	16	512	7005/7008	16	1K
650	16	512	7005/7008	16	1K
620	8	256	7005/7008	16	1K

Table 28 *ArubaOS 8 Recommended Controllers*

Virtual Private Network Concentrators

MCs can be configured as VPNCs to function as an IPsec termination point in the data center for controllers in different geographical locations.

- From a topology standpoint, a VPNC is the hub with branch controllers as spokes
- From a configuration standpoint, the VPNC acts as another MC that is managed by the MM. VPNCs are placed under their own hierarchical node on the MM containing VPNC-specific configuration
- A VPNC may be backed up by a standby VPNC for redundancy purposes
- Though MCs could terminate their IPsec connections directly on the MM (provided that the MM is built according to SKU-mandated hardware specifications) it is highly recommended to terminate the controllers on a VPNC if user traffic from any branch site needs to be routed to the data center

Unsupported Access Points

The following APs are not supported under ArubaOS 8, as of ArubaOS 8.2.0.1. As always, please refer to the latest ArubaOS release notes to confirm supported hardware:

- AP-60
- AP-65
- AP-68
- AP-70
- AP-85
- AP-120/121
- AP-124/125
- AP-92/93 (supported up to ArubaOS 8.2)

ArubaOS 6 Topology Migrations

This section describes common ArubaOS 6 topologies being used in production environments and provides corresponding ArubaOS 8 topology migration recommendations.

For each topology recommendation, the following details are included:

- Description
- Advantages and disadvantages
- Migration requirements

- Migration procedure (manual)

Master and Standby Master

In this ArubaOS 6 design, a master controller terminates all the APs in the network. This active master is supported by a standby master using Virtual Router Redundancy Protocol (VRRP) for redundancy. High Availability (AP Fast Failover) is configured on the master meaning APs terminate their active tunnels on the active master in addition to establishing standby tunnels to the standby master.

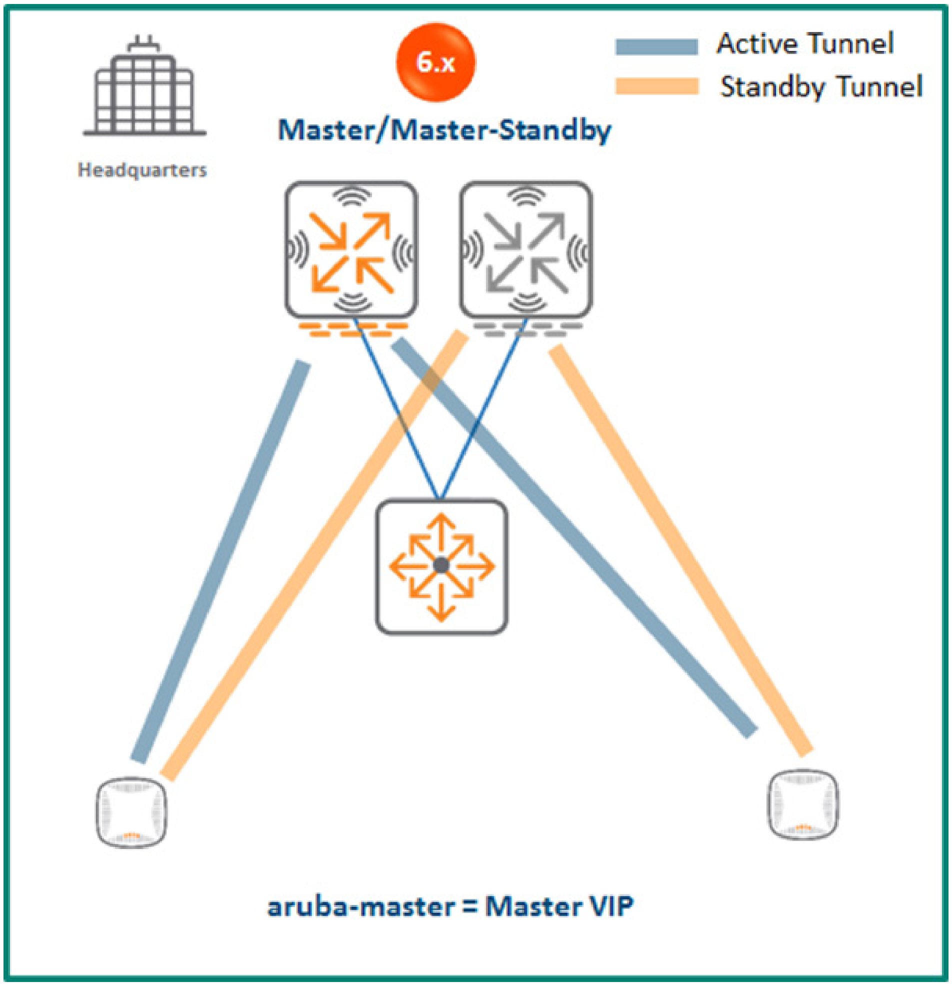

Figure 149 *ArubaOS 6 Master/Standby Architecture*

Since VRRP is being used for master failure detection and the master-standby master design does not support the inter-controller heartbeat feature of AP Fast Failover, failure detection will not be sub-second. I.e. APs will wait for eight missed heartbeats to the master before triggering failover to the standby master. However, the failover process will be instant and simultaneous for all APs unlike traditional VRRP failover which requires APs to re-bootstrap upon failover.

MM Terminating MCs

Topology

To implement this ArubaOS 8 design a MM must first be deployed and configured. The ArubaOS 6 master and standby master controllers become MCs managed by the MM. The controllers can form a cluster for redundancy as well as AP and client load balancing purposes. The controller that is elected as the cluster leader will determine how APs and clients are load balanced in the cluster.

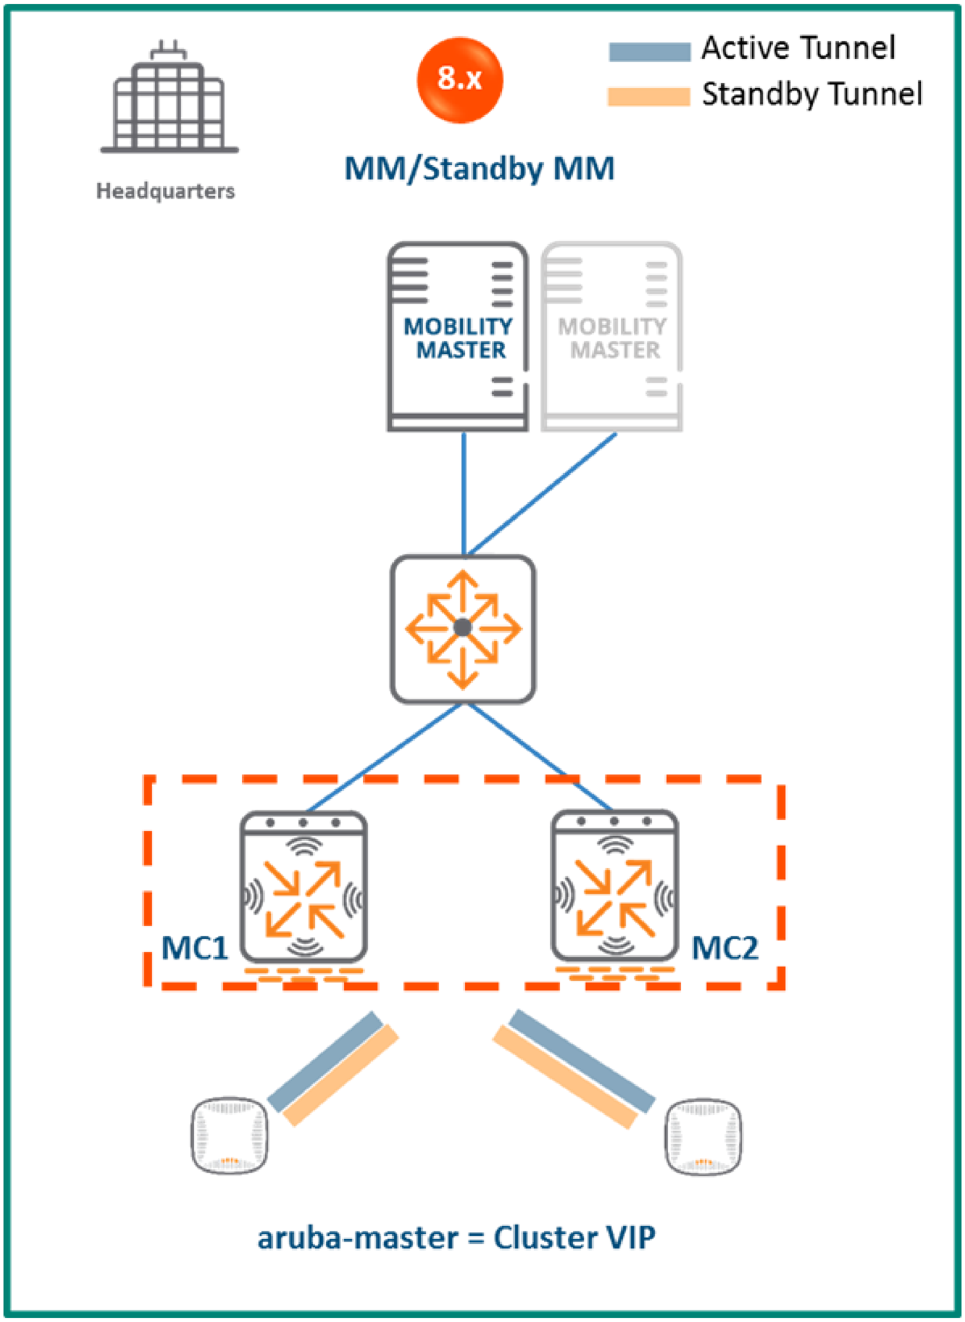

Figure 150 *MM Terminating MCs*

Configuration Hierarchy

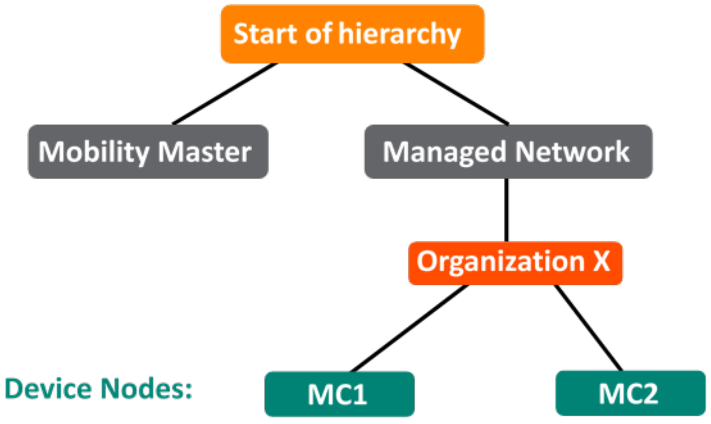

Figure 151 *MM Terminating MCs Configuration Hierarchy*

Design Benefits

- **Maximize benefits** - The MM Terminating MCs design is ideal for fully leveraging the capabilities of ArubaOS 8
- **Scalability** - New controllers can be easily added and managed by the MM
- **Ease of migration** - If an existing deployment has multiple topologies they can all be migrated under the MM into their own nodes in the hierarchy
- **Management** - Centralized configuration and management of controllers
- **Hierarchical configuration model** - Configuration is performed on nodes (logical folders such as USA, West Coast, California, Santa Clara, etc.) and each controller is assigned to a particular node depending on location and context
- **Clustering** - Having controllers in a cluster enables seamless client failover between controllers with no impact to user experience. The clustering feature also facilitates client roaming as well as AP and client load balancing. Clustering is required to support Live Upgrades.
- **Live Upgrade** - Upgrade a cluster of controllers in real time without end users experiencing any loss of connectivity or performance. There is no need for scheduling maintenance cycles for network upgrades
- **AirMatch** - RF intelligence is centralized on the MM which significantly improves the RF management and interference mitigation capabilities of the WLAN
- **REST API support**
- **Multi-version support** - Flexibility of upgrading individual controllers to test out new features or bug fixes. Controllers in a cluster need to run the same ArubaOS software version and may be upgraded together

- **In-service software module upgrades** - Loadable Service Modules (LSMs) such as UCC, AirGroup, AppRF etc. can be updated during runtime removing the need to schedule any maintenance cycles

Design Caveats

- The MM does not terminate APs. APs can only be terminated on a MC

Migration Requirements

- Migration requires the purchase of virtual MM capacity licenses or the purchase of a hardware MM (and optionally a backup hardware MM)
- If a backup MM is available then the licenses on each MM will be aggregated and synchronized across
- Other licenses such as AP and PEF need to be migrated manually or via the "My Networking Portal"

Migration Options

- Migration can occur manually or via the Migration Tool
- Manual migration steps are detailed below. To perform migration using the Migration Tool, please refer to the ArubaOS 8 Migration Guide

Migration Strategy

Existing ArubaOS 6 Deployment

- Active and standby master
- APs terminating on the active master with standby master as backup

New ArubaOS 8 Deployment

- MM managing controllers MC1 and MC2
- APs terminating on MC1 and MC2

Migration Procedure

Manual migration requires a complete rebuild of the existing ArubaOS 6 topology from by going through the following steps:

1. Deploy the MM and perform initial setup
2. Configure licensing on the MM

3. Create a configuration hierarchy on MM and whitelist the active and standby master MAC addresses
4. Repeat step 1 if a backup MM is being installed as well
5. Configure MM redundancy if a backup MM has been installed. Going forward, the MM VIP will be used for configuration management
6. Configure clustering between the controllers and enable AP load balancing
7. Create a VIP between the cluster member IPs and optionally create VIPs for RADIUS COA
8. Create an AP group and SSID. GUI: **Managed Network>(select node)>AP Groups.** GUI: **Managed Network>(select node)>Tasks>Create a new WLAN**
9. Whitelist the APs on the MM by populating the CPSec whitelist table (including mapping the APs to the appropriate AP group). GUI: **Managed Network>(select node)>Configuration>Access Points>Whitelist**
10. Back up the existing configuration on the ArubaOS 6 master controllers. GUI: **Maintenance>Backup Flash**
11. Upgrade the image on the active master to ArubaOS 8 and reboot it. GUI: **Maintenance>Image Management**
12. Provision the master to be managed by MM via the CLI setup dialog. The master will now become MC1
13. Repeat steps 11-12 to convert the standby master to ArubaOS 8 as MC2
14. Change **aruba-master** to point to the cluster VIP
15. The APs that were previously terminating on the master will find the cluster VIP, upgrade their images, terminate on MC1 or MC2 (depending on how the cluster leader load balances the APs), and broadcast the configured SSID
16. Connect a wireless client to the SSID to test connectivity
17. Optionally, test seamless client failover by running a voice/video application and disconnecting the user's active controller

Standalone MC and Standby Standalone

Topology

This ArubaOS 8 design consists of a standalone MC backed up by another standalone MC. As in the Master and Standby Master ArubaOS 6 design, VRRP is used between the two standalone controllers in an active-standby configuration. Similarly, High Availability (AP Fast Failover) is configured between the controllers so that APs terminate their tunnels on the active standalone controller in addition to setting up a standby tunnel to the standby standalone controller.

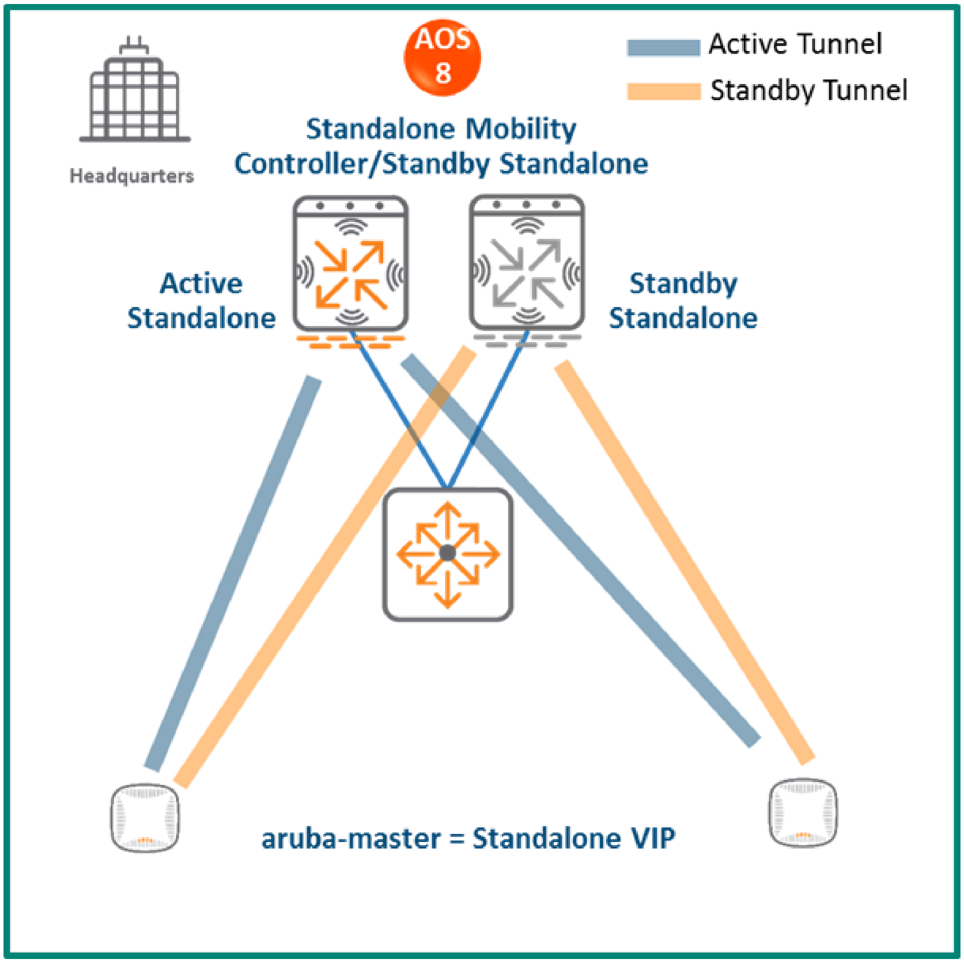

Figure 152 *Standalone MC and Standby Standalone Topology*

Just as with the ArubaOS 6 Master and Standby Master design, the AP Fast Failover detection is not sub-second (i.e. APs will wait for eight missed heartbeats to the master), however the failover itself occurs quickly due to the APs already having standby tunnels to the standby standalone controller. The standalone becomes the new active controller in the event of a failure.

Design Benefits

- No additional hardware is required for migration
- Multi-threaded CLI
- Auto-completion of profile names

Design Caveats

- APs can only terminate on the active standalone controller
- VRRP and AP Fast Failover are configured, however inter-controller heartbeats for AP Fast Failover is not supported in this design. AP Fast Failover detection will not be sub-second

since the failover depends on VRRP latency. Upon detection the actual failover itself will be quick and simultaneous for all APs due to their existing standby tunnels

Migration Requirements

Licenses such as AP and PEF need to be migrated manually or via the "My Networking Portal"

Migration Options

No migration tool support. Migration can only be performed manually.

Migration Strategy

Existing ArubaOS 6 Deployment

- Active and standby master
- APs terminating on the active master with standby master as backup

New ArubaOS 8 Deployment

- Active standalone and standby standalone controllers
- APs with the active and standby tunnels terminating on the active and standby controllers respectively

Migration Procedure

Manual migration requires a complete rebuild of the existing ArubaOS 6 topology.

1. Backup the existing configuration on the ArubaOS 6 masters
2. Upgrade the image on the active master to ArubaOS 8 and reboot the controller
3. Provision the active master as an ArubaOS 8 standalone controller via the CLI setup dialog. The master will now become an ArubaOS 8 standalone controller
4. Repeat steps 2-3 to convert the standby master into an ArubaOS 8 standalone controller
5. Configure licensing on the active standalone controller. The standby standalone controller will inherit licenses from the active standalone once database synchronization is configured as part of step 6
6. Configure master redundancy between the two standalone controllers. A VIP will be created between MC1 and MC2 as a result of the VRRP configuration. Going forward, configuration management will occur through the VIP
7. Create an AP group and SSID under the MM node (/mm in CLI). This will push the common configuration to both standalone controllers

8. Configure AP Fast Failover for both standalone controllers
9. Whitelist APs under **MM>Configuration>Access Points>Whitelist**
10. Change **aruba-master** to point to the standalone VIP
11. The APs will then find the VIP (i.e. active standalone controller), upgrade their images, terminate their tunnels on the VIP, and broadcast the configured SSID
12. Connect a wireless client to the SSID and test connectivity
13. Optionally, test client failover by disconnecting the active standalone controller

Master and Single Local

In this ArubaOS 6 design, a master controller is managing a local controller. The same controller models are recommended for the master and local. There can be two variations of this design:

- **Redundancy Model** (also known as an Active-Standby model) - APs terminate on the local controller and the master provides redundancy for the local. High Availability (AP Fast Failover) is configured between the controllers so that in the event the APs lose connectivity to the local controller, they can instantly failover to the master.
- **Capacity Model** (also known as an Active-Active model) **-** This is an alternative single-master, single-local design where the master, in addition to managing the local, also shares the AP load with the local. High Availability (AP Fast Failover) is configured between the controllers such that when one controller goes down, its APs can seamlessly failover to the other controller.

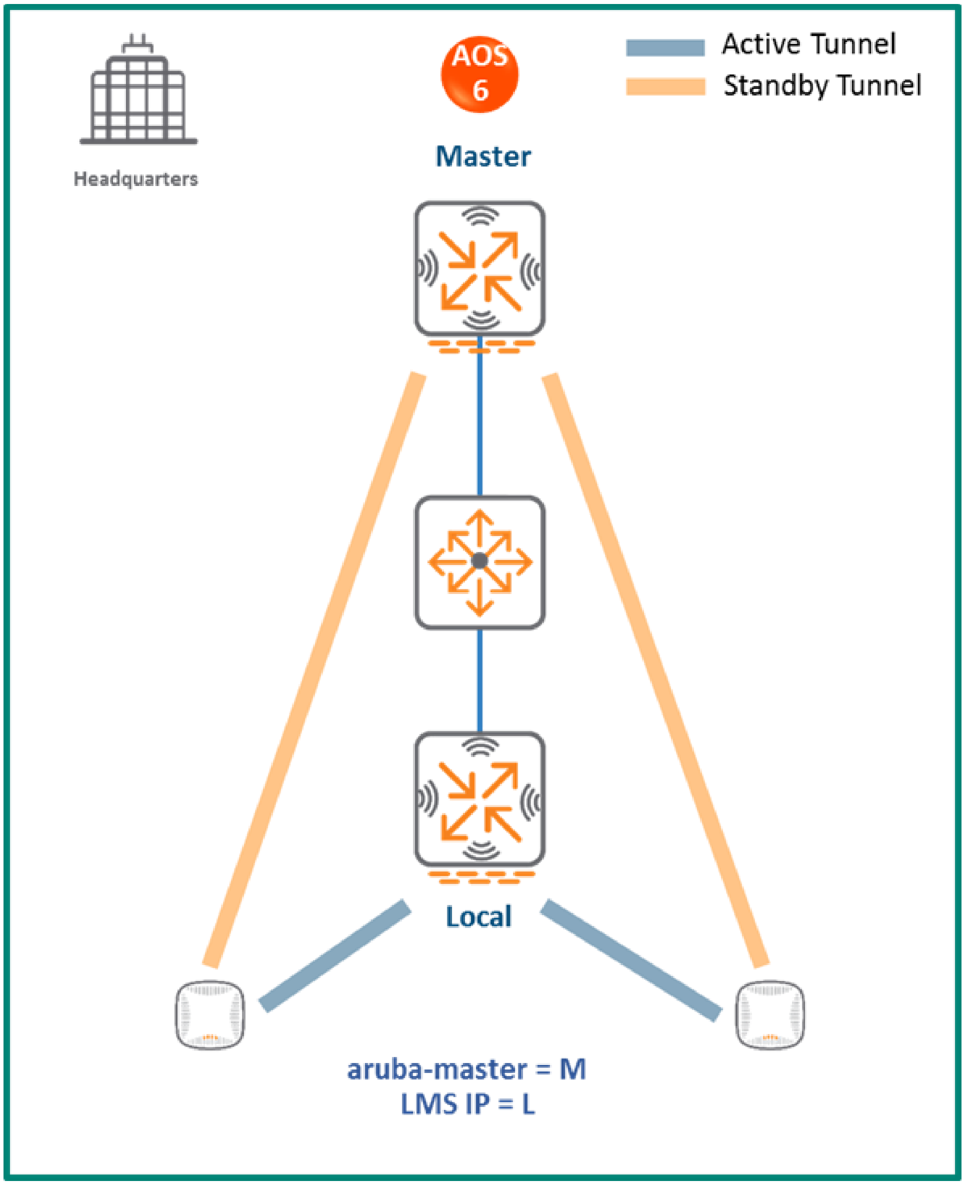

Figure 153 *Master and Single Local*

In both designs:

- Each controller needs to have enough capacity to accommodate the number of APs that could potentially failover from the second controller. In the redundancy model, each controller typically terminates APs at up to 80% of the controller capacity. In the capacity model, each controller typically terminates APs at up to 40% of the controller capacity
- The AP Fast Failover detection is not sub-second (i.e. APs will wait for eight missed heartbeats to the master) however the failover itself occurs quickly since all the APs already have standby tunnels built to the standby standalone controller. The standby standalone controller becomes the new active controller upon failover

MM Terminating MCs

Topology

In this ArubaOS 8 design, a MM is initially deployed and configured. The ArubaOS 6 master and local controllers become MCs managed by the MM. The controllers can form a cluster for redundancy and AP/client load balancing purposes. The controller that is elected as the cluster leader will decide how APs and clients are load balanced in the cluster.

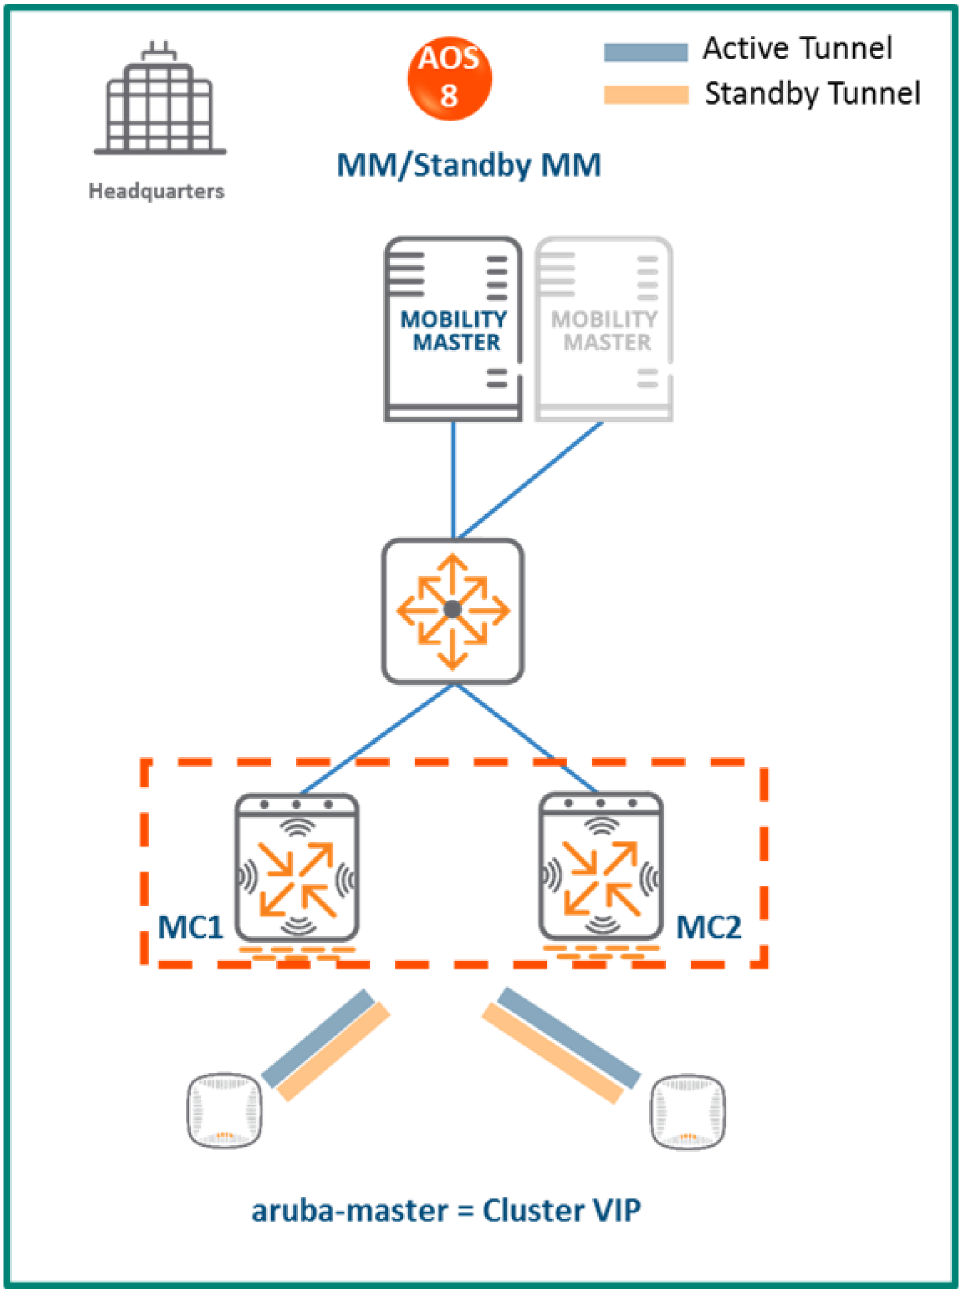

Figure 154 *MM Terminating MCs Topology*

Configuration Hierarchy

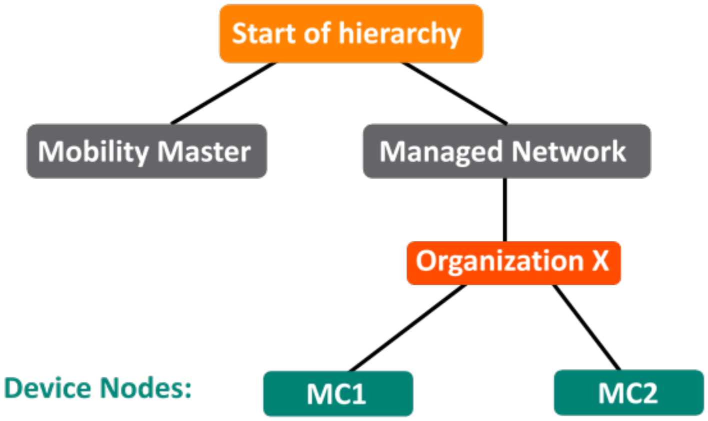

Figure 155 *MM Terminating MCs Configuration Hierarchy*

Design Benefits

- **Maximize benefits** - The MM Terminating MCs design is ideal for fully leveraging the capabilities of ArubaOS 8
- **Scalability** - New controllers can be easily added and managed by the MM
- **Ease of migration** - If an existing deployment has multiple topologies they can all be migrated under the MM into their own nodes in the hierarchy
- **Management** - Centralized configuration and management of controllers
- **Hierarchical configuration model** - Configuration is performed on nodes (logical folders such as USA, West Coast, California, Santa Clara, etc.) and each controller is assigned to a particular node depending on location and context
- **Clustering** - Having controllers in a cluster enables seamless client failover between controllers with no impact to user experience. The clustering feature also facilitates client roaming as well as AP and client load balancing. Clustering is required to support Live Upgrades
- **Live Upgrade** - Upgrade a cluster of controllers in real time without end users experiencing any loss of connectivity or performance. There is no need for scheduling maintenance cycles for network upgrades
- **AirMatch** - RF intelligence is centralized on the MM which significantly improves the RF management and interference mitigation capabilities of the WLAN
- **REST API support**
- **Multi-version support** - Flexibility of upgrading individual controllers to test out new features or bug fixes. Controllers in a cluster need to run the same ArubaOS software version and may be upgraded together

- **In-service software module upgrades** - Loadable Service Modules (LSMs) such as UCC, AirGroup, AppRF etc. can be updated during runtime removing the need to schedule any maintenance cycles

Design Caveats

- The MM does not terminate any APs. APs can only be terminated on MCs.

Migration Requirements

- Requires purchase of virtual MM capacity licenses or purchase of hardware MM (and optionally a backup hardware MM)
- If you have a backup MM, then the licenses on each MM will be aggregated and synchronized across both MMs
- Other licenses such as AP and PEF need to be migrated manually or via the "My Networking Portal"

Migration Options

- Migration can occur manually or via the Migration Tool
- Manual migration steps are detailed below. To perform migration using the Migration Tool, please refer to the ArubaOS 8 Migration Guide

Migration Strategy

Existing ArubaOS 6 Deployment

- Master and local
- APs terminating on the local with master as backup

New ArubaOS 8 Deployment

- MM backed up by a standby MM
- MM managing controllers MC1 and MC2
- APs terminating on MC1 and MC2

Migration Procedure

Manual migration requires a complete rebuild of the existing ArubaOS 6 topology by performing the following steps:

1. Deploy the MM and perform initial setup

2. Configure licensing on the MM
3. Create a configuration hierarchy on the MM and whitelist the master and local controller MAC addresses
4. Repeat step 1 if a MM is also being installed
5. Configure MM redundancy if a MM has been deployed. The MM VIP will be used moving forward for configuration management
6. Configure clustering between the MCs and enable AP load balancing
7. Create a VIP between the cluster member IPs and optionally create VIPs for RADIUS COA
8. Create an AP group and SSID. GUI**: Managed Network>(select node)>AP Groups**. GUI: **Managed Network>(select node)>Tasks>Create a new WLAN**
9. Whitelist your APs on the MM and map them to the appropriate AP group. GUI: **Managed Network>(select node)>Configuration>Access Points>Whitelist**
10. Back up the existing configuration on the ArubaOS 6 master and local. GUI: **Maintenance>Backup Flash**
11. Upgrade the image on the local to ArubaOS 8 and reboot it. GUI: **Maintenance>Image Management**
12. Provision the local to be managed by the MM via the CLI setup dialog. The local will now become MC1
13. Now repeat steps 11-12 to convert the master to MC2
14. Change **aruba-master** to point to the cluster VIP
15. The APs that were previously terminating on the local will find the cluster VIP, upgrade their images, terminate on MC1 or MC2 (depending on how the cluster leader load balances the APs), and broadcast the configured SSID
16. Connect a wireless client to the SSID and test connectivity
17. Optionally, test seamless client failover by running a voice/video application and disconnecting the user's active controller

Standalone MC with Master Redundancy

Topology

This ArubaOS 8 design consists of a standalone MC backed up by another standalone MC. VRRP is enabled between the two standalone controllers in an active-standby configuration. High Availability (AP Fast Failover) is also configured between the controllers so that APs set up a standby tunnel to the standby standalone controller in addition to terminating their tunnels on the active standalone.

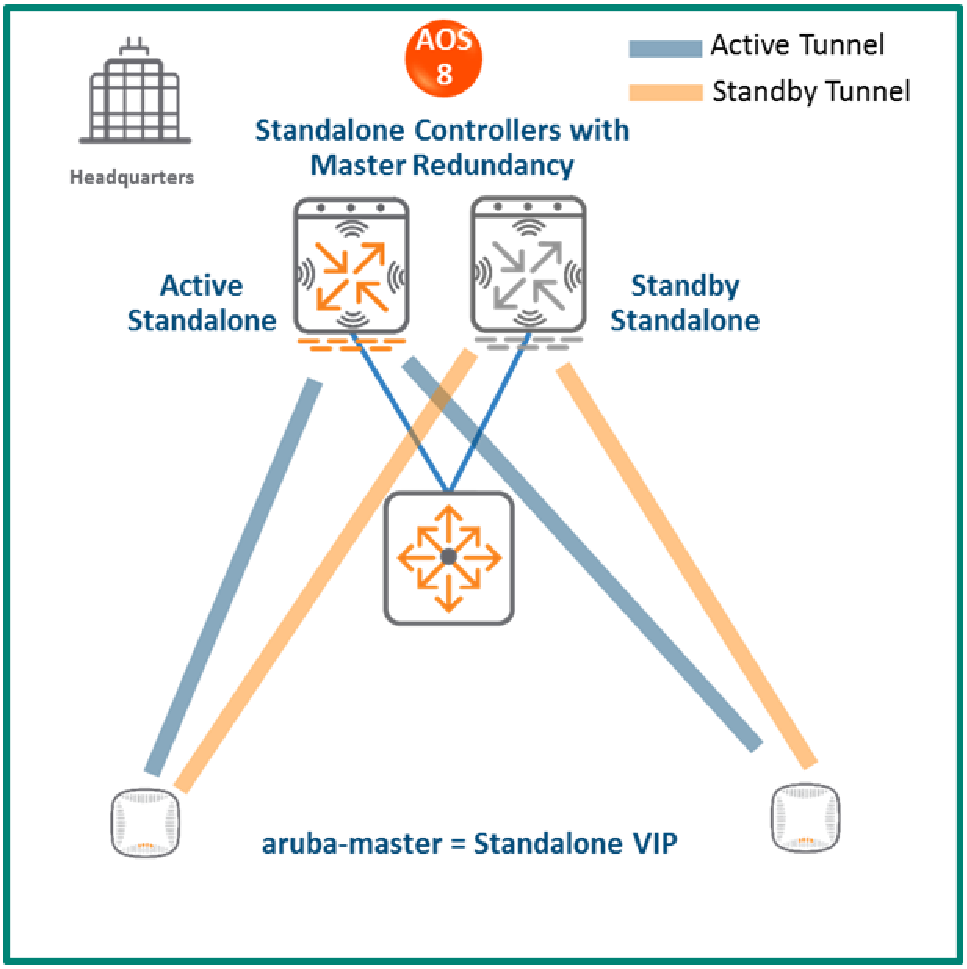

Figure 156 *Standalone MC with Master Redundancy*

The AP Fast Failover detection is not sub-second (i.e. APs will wait for eight missed heartbeats to the master), however the failover itself occurs quickly due to the APs already having standby tunnels to the standby standalone controller. The standalone becomes the new active controller in the event of a failure.

Design Benefits

- No additional hardware is required for migration

- Multi-threaded CLI
- Auto-completion of profile names

Design Caveats

- APs can only terminate on the active standalone controller
- No AP Fast Failover. The master redundancy configuration between the two standalone controllers uses VRRP to detect failover meaning that AP failover will not be sub-second since the failover mechanism is dependent on VRRP latency

Migration Requirements

Licenses such as AP and PEF need to be migrated manually or via the "My Networking Portal"

Migration Options

No migration tool support. Migration can only be performed manually.

Migration Strategy

Existing ArubaOS 6 Deployment

- Master and local
- APs terminating on the local with master as backup

New ArubaOS 8 Deployment

- Active standalone and standby standalone controllers
- APs with the active and standby tunnels terminating on the active and standby controllers respectively

Migration Procedure

Manual migration requires a complete rebuild of the existing ArubaOS 6 topology.

1. Back up the existing configuration on the ArubaOS 6 master and local controller
2. Upgrade the image on the local to ArubaOS 8 and reboot the controller
3. Provision the local as an ArubaOS 8 standalone controller via the CLI setup dialog. The local will become a standalone controller
4. Upgrade the image on the master to ArubaOS 8 and reboot the controller

5. Provision the master as an ArubaOS 8 standalone controller via the CLI setup dialog. The master will become another standalone controller
6. Configure licensing on the active standalone controller. The standby standalone controller will inherit licenses from the active standalone once the database synchronization is configured as part of step 7
7. Configure master redundancy between the two standalone controllers. A VIP will be created as a result of the VRRP configuration. Going forward, configuration management will occur through the VIP
8. Navigate to **/mm** and create an AP group and SSID
9. Configure AP Fast Failover for both standalone controllers
10. Whitelist your APs under **MM>Configuration>Access Points>Whitelist**
11. Change **aruba-master** to point to the standalone VIP
12. The APs will then find the VIP (i.e. active standalone controller), upgrade their images, terminate their tunnels on the VIP, and broadcast the configured SSID
13. Connect a wireless client to the SSID and test connectivity
14. Optionally, test client failover by disconnecting the active standalone controller

Master and Multiple Locals (Single Campus)

In this ArubaOS 6 design, a master (backed up by a standby master) controller is managing a group of local controllers. APs terminate on one of the local controllers with the other locals acting as backups. AP Fast Failover is configured to provide sub-second failover for the APs when connectivity to their primary controller is lost.

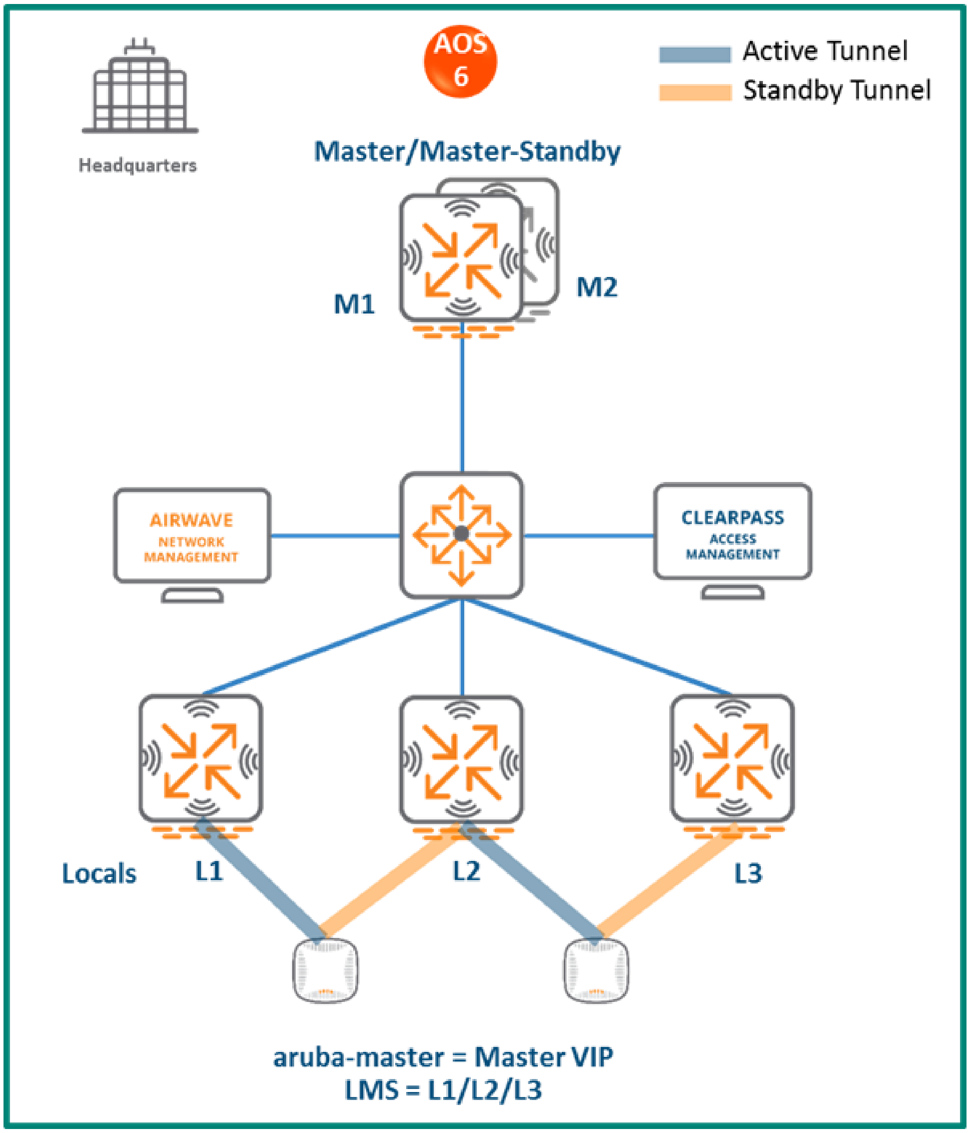

Figure 157 *Master Controller and Multiple Locals*

MM Terminating MCs

Topology

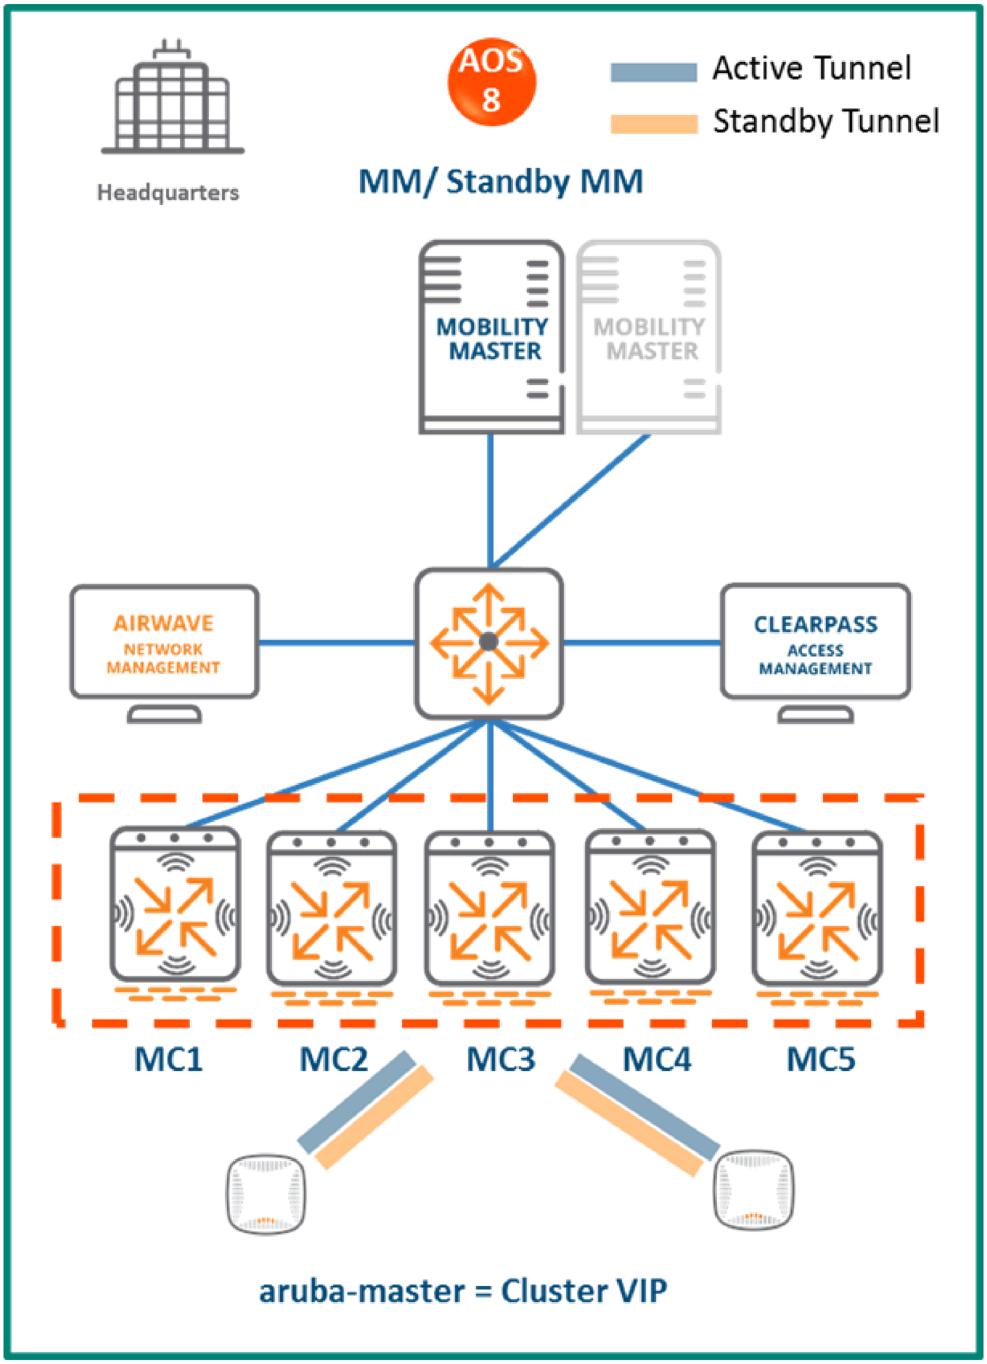

Figure 158 *MM Terminating MCs Topology*

In this ArubaOS 8 design:

- A MM (either virtual or hardware) is deployed and configured along with a backup MM
- Each ArubaOS 6 local controller (L1, L2, and L3) becomes an ArubaOS 8 MC (MC1, MC2, MC3)

- The ArubaOS 6 master (M1) and standby master (M2) become two additional ArubaOS 8 MCs (MC4 and MC5)
- The MCs can be part of a cluster and share the AP and client load
- If the locals were geographically separated from each other, then post migration the APs terminating on L1, L2, and L3 will now terminate on MC1, MC2 and MC3 respectively
- If all the locals were part of a large campus, then the cluster leader will distribute the AP and client load among MC1-MC5

Configuration Hierarchy

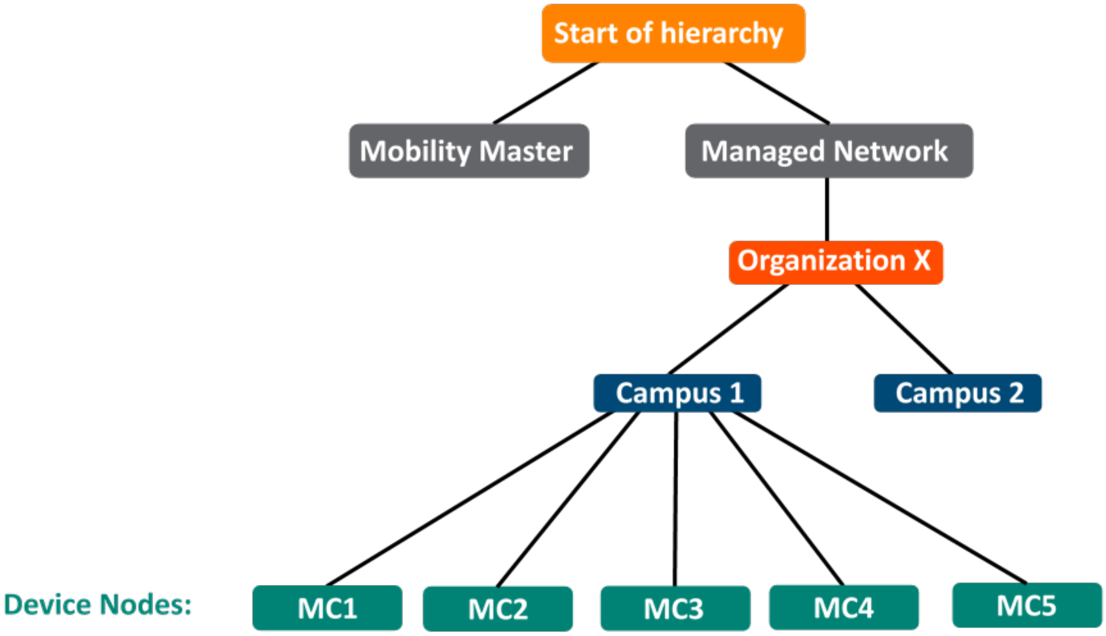

Figure 159 *MM Terminating MCs Configuration Hierarchy*

Design Benefits and Caveats

- **Maximize benefits** - The MM Terminating MCs design is ideal for fully leveraging the capabilities of ArubaOS 8
- **Scalability** - New controllers can be easily added and managed by the MM
- **Ease of migration** - If an existing deployment has multiple topologies they can all be migrated under the MM into their own nodes in the hierarchy
- **Management** - Centralized configuration and management of controllers
- **Hierarchical configuration model** - Configuration is performed on nodes (logical folders such as USA, West Coast, California, Santa Clara, etc.) and each controller is assigned to a particular node depending on location and context
- **Clustering** - Having controllers in a cluster enables seamless client failover between controllers with no impact to user experience. The clustering feature also facilitates client

roaming as well as AP and client load balancing. Clustering is required to support Live Upgrades.

- **Live Upgrade** - Upgrade a cluster of controllers in real time without end users experiencing any loss of connectivity or performance. There is no need for scheduling maintenance cycles for network upgrades
- **AirMatch** - RF intelligence is centralized on the MM which significantly improves the RF management and interference mitigation capabilities of the WLAN
- **REST API support**
- **Multi-version support** - Flexibility of upgrading individual controllers to test out new features or bug fixes. Controllers in a cluster need to run the same ArubaOS software version and may be upgraded together
- **In-service software module upgrades** - Loadable Service Modules (LSMs) such as UCC, AirGroup, and AppRF can be updated during runtime removing the need to schedule any maintenance cycles.

Migration Requirements

- Requires purchase of virtual MM capacity licenses or the purchase of a hardware MM (and optionally a backup hardware MM)
- If you have a backup MM, then the licenses on each MM will be aggregated and synchronized across both MMs
- Other licenses such as AP and PEF need to be migrated manually or via the "My Networking Portal"

Migration Options

- Migration can occur manually or via the Migration Tool
- Manual migration steps are detailed below. To perform migration using the Migration Tool, please refer to the ArubaOS 8 Migration Guide

Migration Strategy

Existing ArubaOS 6 Deployment

- Locals L1, L2, and L3 and masters M1and M2
- 3 AP groups are configured to have groups of APs terminate on each of L1, L2 and L3.

New ArubaOS 8 Deployment

- MM backed up by a standby MM

- MM managing MC1, MC2, MC3, MC4, and MC5
- APs terminating on:
  - MC1, MC2, MC3 in the case of a multi-site campus, with a controller in each site
  - The cluster VIP for a large campus

Migration Procedure

Manual migration requires a complete rebuild of the existing ArubaOS 6 topology by performing the steps listed below. These steps cover termination of APs on a cluster VIP. In the case of a multi-site campus, the APs could terminate on one of three local management switch (LMS) IPs (MC1, MC2, or MC3).

1. Deploy the MM and perform initial setup
2. Configure licensing on the MM
3. Create a configuration hierarchy and whitelist the MAC addresses of M1, M2, and L1-L3 on the MM
4. Repeat step 1 if installing a backup MM
5. Configure MM redundancy if a backup MM is being installed. Going forward, use the MM VIP for configuration management
6. Configure clustering between the MCs and enable AP load balancing
7. Create a VIP between the cluster member IPs and optionally create VIPs for RADIUS COA
8. Create an AP group and SSID. GUI: **Managed Network>(select node)>AP Groups**. GUI: **Managed Network>(select node)>Tasks> Create a new WLAN**
9. Whitelist the APs on the MM and map them to the appropriate AP group. GUI: **Managed Network>(select your node)>Configuration>Access Points>Whitelist**.
10. Back up the existing configuration on the ArubaOS 6 masters and locals. GUI: **Maintenance>Backup Flash**
11. Upgrade the image on local L1 to ArubaOS 8 and reboot it. GUI: **Maintenance>Image Management**
12. Provision local L1 to be managed by the MM via the CLI setup dialog. L1 will become ArubaOS 8 MC1
13. Repeat steps 11-12 for both L2 and L3 to convert them into ArubaOS 8 MC2 and MC3
14. Repeat steps 11-12 to convert M1 and M2 to MC4 and MC5. These controllers can be added to the cluster to share the AP and client load between cluster members
15. Change **aruba-master** to point to the cluster VIP
16. The APs that were terminating on the locals will find the cluster VIP, upgrade their images, terminate on one of MC1-MC5 (depending on how the cluster leader load balances the APs), and broadcast the configured SSID

17. Connect a wireless client to the SSID and test connectivity

18. Optionally, test seamless client failover by running a voice/video application and disconnecting the user's active controller

MC Master Terminating MCs

Topology

This ArubaOS 8 design consists of a hardware controller deployed as a MC Master (optionally backed up by another MC Master) that manages a group of MCs.

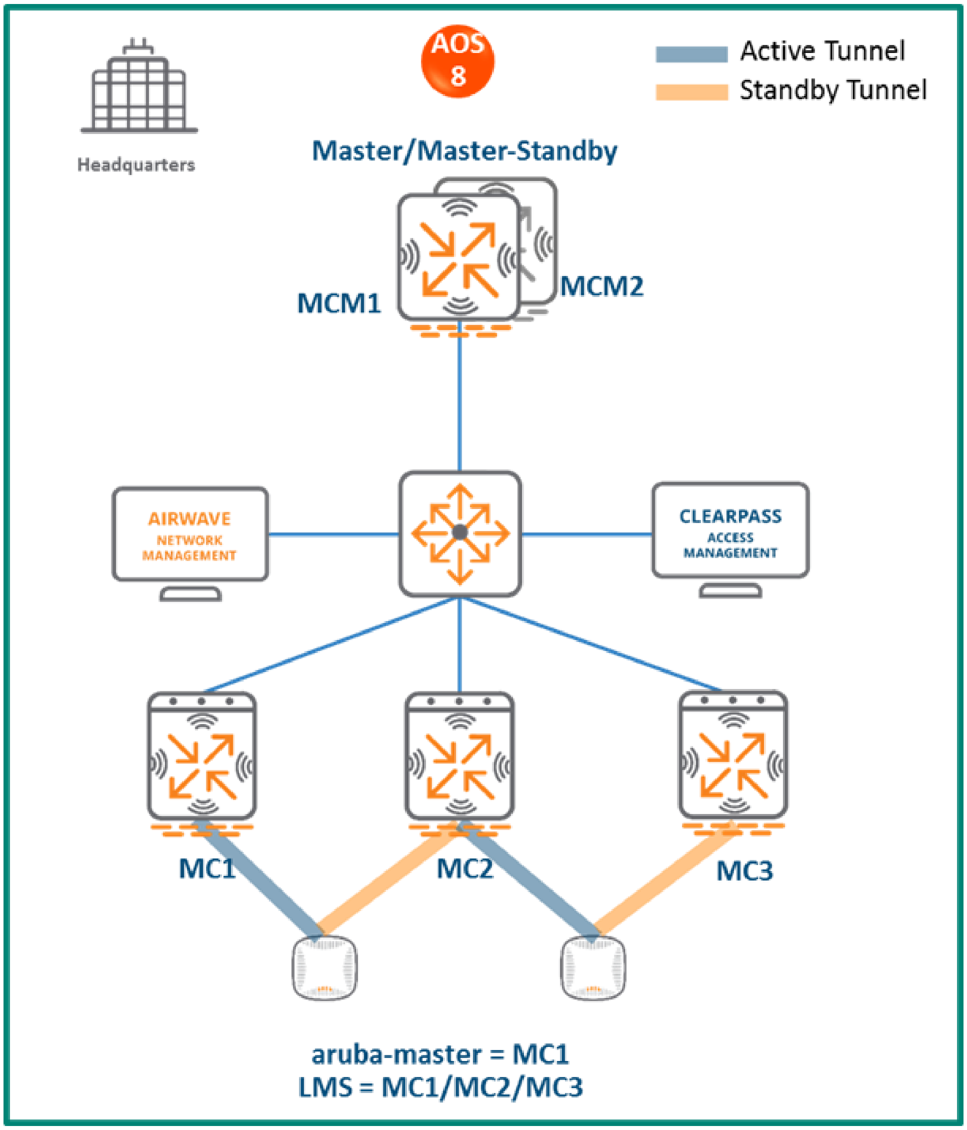

Figure 160 *MC Master Terminating MCs Topology*

This design helps transition deployments to ArubaOS 8 that are unable to deploy a MM. This MC Master topology should eventually be migrated to a MM topology in order to take full advantage of the capabilities offered by ArubaOS 8.

In this design:

- The ArubaOS 6 master (M1) and standby master (M2) become ArubaOS 8 MC Masters (MCM1 and MCM2).
- The ArubaOS 6 local controllers (L1, L2 and L3) become ArubaOS 8 MCs (MC1, MC2 and MC3).
- APs terminating on L1, L2, and L3 will now terminate on MC1, MC2, and MC3 respectively.

Configuration Hierarchy

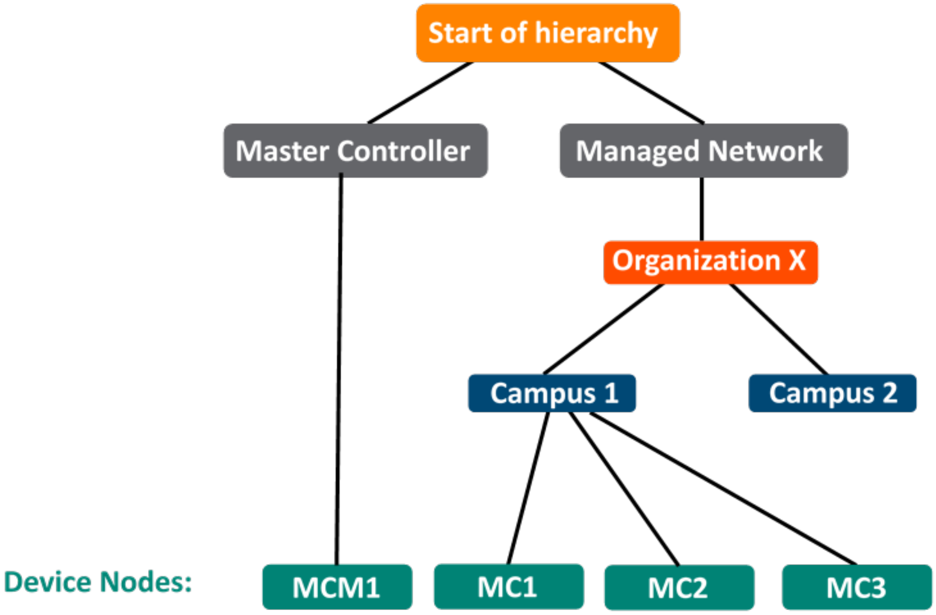

Figure 161 *MC Master Terminating Mobility Configuration Hierarchy*

Design Benefits

- A similar topology is maintained in which the MC Master manages the MCs and no additional hardware is required as long as the MC Master is an Aruba 7030 or larger controller
- The hierarchical configuration model offers fully centralized configuration and management of the WLAN
- Additional controllers could be added later and managed by the MC Master

Design Caveats

- Requires purchase of an Aruba 7030 or larger controller to serve as the MC Master as well as the backup MCM if one is not already present

- AP termination on the MC Master is not supported. This has the following impact on AP termination options:
  - Any APs that are terminating on the master in ArubaOS 6 would need to be redistributed among the locals prior to migration. The locals should have enough capacity to accommodate the additional APs
  - APs can failover between MCs but cannot failover to the MC Master
- The clustering feature is not supported in a MC Master deployment. AP Fast Failover between MCs is the only controller redundancy option
- AirMatch is not supported
- All controllers in the topology must run the same ArubaOS version
- No centralized monitoring

Migration Requirements

- Verify that the ArubaOS 6 master controller meets the MC Master hardware requirements (Aruba 7030 or any Aruba 7200 series controller)
- Ensure that the ArubaOS 6 master is not terminating any APs as an ArubaOS 8 MC Master does not support AP termination
- Ensure that AP, PEF, and all other licenses have been migrated manually or via the "My Networking Portal"

Migration Options

- Migration can occur manually or via the Migration Tool
- Manual migration steps are detailed below. To perform migration using the Migration Tool, please refer to the ArubaOS 8 Migration Guide

Migration Strategy

Existing ArubaOS 6 Deployment

- Locals L1, L2, L3 and masters M1 and M2
- 3 AP groups are configured for termination on L1, L2, and L3

New ArubaOS 8 Deployment

- MCM1 backed up by MCM2
- MCM1 managing MC1, MC2, and MC3

Migration Procedure

Manual migration requires a complete rebuild of the existing ArubaOS 6 topology by performing the steps listed below:

1. Back up the existing configuration on the ArubaOS 6 masters and locals. GUI: **Maintenance>Backup Flash**
2. Upgrade the image on master M1 to ArubaOS 8 and reboot the controller
3. Provision M1 as a MC Master through the CLI setup dialog. M1 will now become MCM1
4. Repeat steps 2 and 3 to convert M2 to MCM2
5. Configure master redundancy between MCM1 and MCM2. The MC Master VIP will be used for configuration management moving forward
6. Configure licensing on the MC Master
7. Create a configuration hierarchy on the MC Master and whitelist the MAC addresses of controllers L1-L3
8. Create three AP groups under **/md** (or a child node), each with the LMS IP of MC1, MC2, and MC3 respectively. GUI: **Managed Network>(select node)>AP Groups**
9. Create an SSID for each AP group. GUI: **Managed Network>(select node)>Tasks> Create a new WLAN**
10. Whitelist the APs on the MC Master. This includes mapping them to their respective AP groups. GUI: **Managed Network>(select node)>Configuration>Access Points > Whitelist**
11. Upgrade the image on local L1 to ArubaOS 8 and reboot it. GUI: **Maintenance>Image Management**
12. Provision local L1 to be managed by the MC Master via the CLI setup dialog. L1 will become MC1
13. Now repeat steps 11-12 for L2 and L3 to convert them to ArubaOS 8 MC2 and MC3
14. Change **aruba-master** to MC1's IP
15. Once MC1 is visible on the MC Master, the APs that were terminating on L1 will find MC1, upgrade their images, download the LMS-IP for MC1, terminate their tunnels on MC1, and broadcast the configured SSID
16. Similarly, the APs on L2 and L3 will show up on MC2 and MC3, respectively
17. Connect a wireless client to the SSID and test connectivity
18. Optionally, configure AP Fast Failover via the MC Master to enable sub-second AP failover between the MCs

Master and Multiple Locals (Multiple Campuses)

In this ArubaOS 6 design, a master controller backed up by a standby master is managing a group of local controllers. APs terminate on one of the local controllers with the other locals acting as backups. AP Fast Failover is configured to provide sub-second failover for the APs when connectivity to their primary controller is lost.

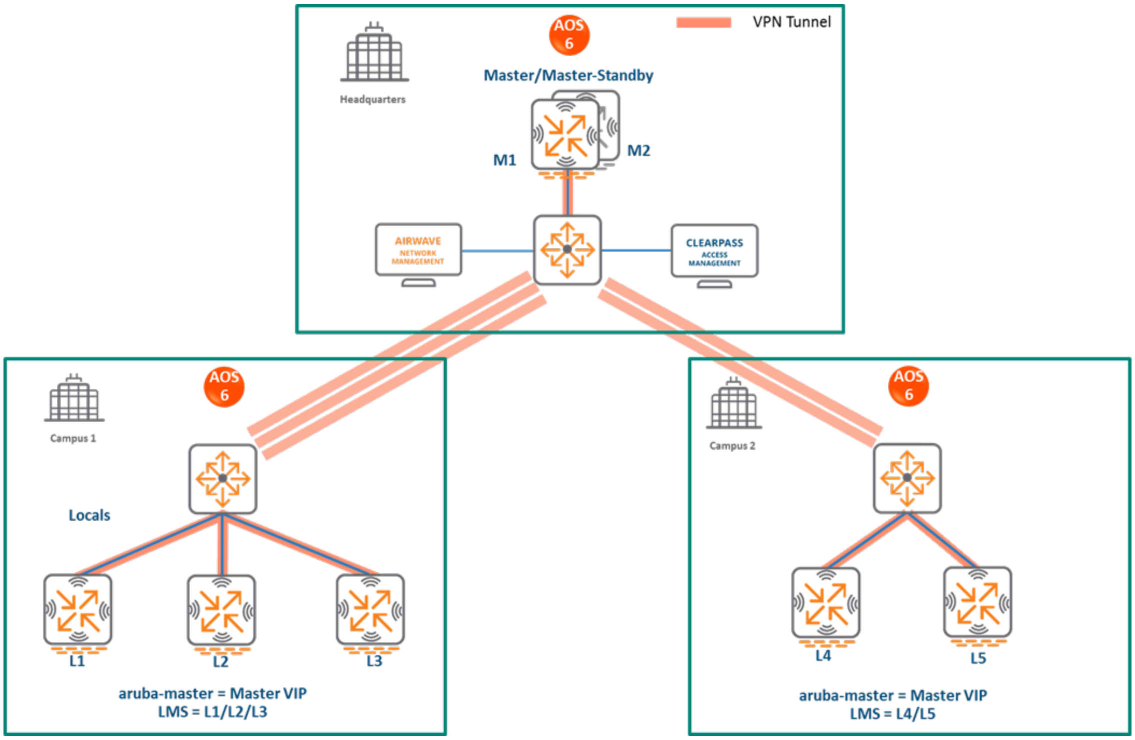

Figure 162 *Master and Multiple Locals (Multiple Campuses)*

MM Terminating MCs

Topology

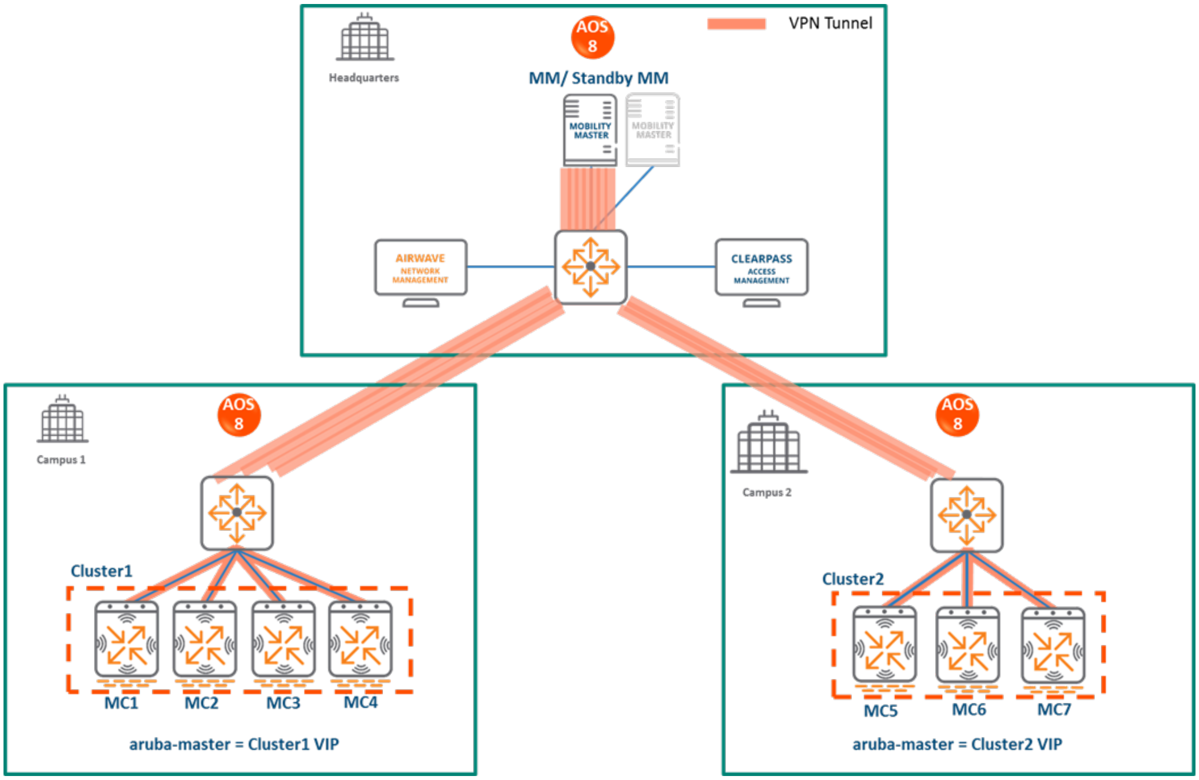

Figure 163 *MM Terminating MCs Topology*

In this ArubaOS 8 design:

- A MM (either virtual or hardware) is deployed and configured along with a backup MM
- In Campus 1, each ArubaOS 6 local controller (L1, L2, and L3) becomes an ArubaOS 8 MC (MC1, MC2, MC3)
- In Campus 2, each ArubaOS 6 local controller (L4 and L5) becomes an ArubaOS 8 MC (MC5 and MC6)
- The MCs in each campus are configured as a cluster and will share the AP and client load
- All MCs terminate their IPsec tunnels on the MM
- If the locals were geographically separated from each other, then the migration is performed so that APs terminating on L1, L2, and L3 will now terminate on MC1, MC2 and MC3 respectively
- If all the locals in each campus are co-located, then post migration the cluster leader will distribute the AP and client load among the cluster members
- The ArubaOS 6 master (M1) and standby master (M2) become two additional ArubaOS 8 MCs (MC4 and MC7) which can be repurposed to become cluster members in each campus
- In the case of remote sites that are separated from the MM via MPLS and/or internet links, if user traffic needs to be routed to access HQ resources then it is recommended to deploy a hardware VPNC at HQ to terminate IPsec connections from the controllers in each site

Configuration Hierarchy

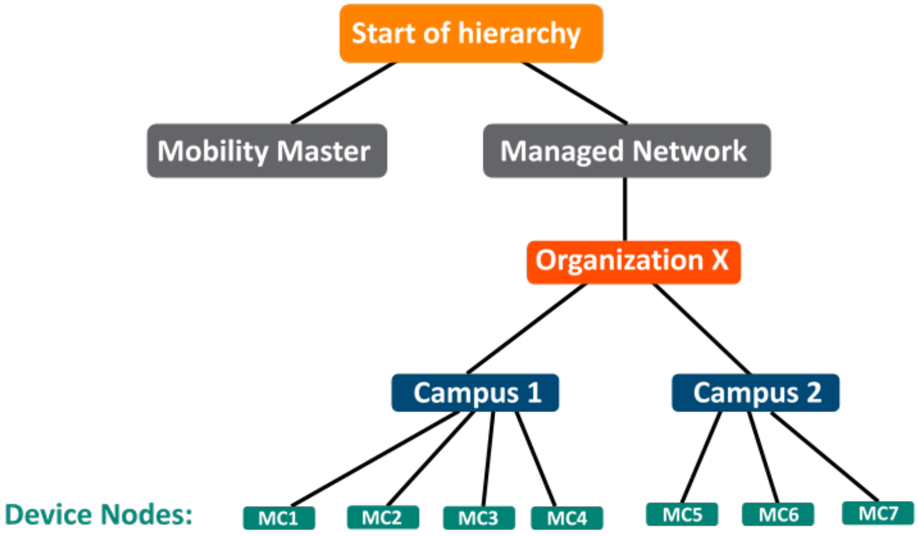

Figure 164 *MM Terminating MCs Configuration Hierarchy*

Design Benefits

- **Maximize benefits** - The MM Terminating MCs design is ideal for fully leveraging the capabilities of ArubaOS 8
- **Scalability** - New controllers can be easily added and managed by the MM
- **Ease of migration** - If an existing deployment has multiple topologies they can all be migrated under the MM into their own nodes in the hierarchy
- **Management** - Centralized configuration and management of controllers
- **Hierarchical configuration model** - Configuration is performed on nodes (logical folders such as USA, West Coast, California, Santa Clara, etc.) and each controller is assigned to a particular node depending on location and context
- **Clustering** - Having controllers in a cluster enables seamless client failover between controllers with no impact to user experience. The clustering feature also facilitates client roaming as well as AP and client load balancing. Clustering is required to support Live Upgrades.
- **Live Upgrade** - Upgrade a cluster of controllers in real time without end users experiencing any loss of connectivity or performance. There is no need for scheduling maintenance cycles for network upgrades
- **AirMatch** - RF intelligence is centralized on the MM which significantly improves the RF management and interference mitigation capabilities of the WLAN
- **REST API support**

- **Multi-version support** - Flexibility of upgrading individual controllers to test out new features or bug fixes. Controllers in a cluster need to run the same ArubaOS software version and may be upgraded together
- **In-service software module upgrades** - Loadable Service Modules (LSMs) such as UCC, AirGroup, AppRF etc. can be updated during runtime removing the need to schedule any maintenance cycles

Design Caveats

- The MM does not terminate APs. APs can only be terminated on a MC
- If the existing ArubaOS 6 deployment has more than 1000 controllers and/or 10,000 APs, then migration to an ArubaOS 8 MM deployment requires the deployment of multiple MMs

Migration Requirements

- Requires purchase of virtual MM capacity licenses or the purchase of a hardware MM
- A backup hardware MM may also be deployed in which case the licenses on each MM will be aggregated and synchronized across both MMs
- Other licenses such as AP and PEF need to be migrated manually or via the "My Networking Portal"

Migration Options

- Migration can occur manually or via the Migration Tool
- Manual migration steps are detailed below. To perform migration using the Migration Tool, please refer to the ArubaOS 8 Migration Guide

Migration Strategy

Existing ArubaOS 6 Deployment

- **HQ:** Master controllers M1 and M2
- **Campus1:** L1, L2, and L3. Three AP groups are configured for termination on each of the local controllers
- **Campus2:** L4 and L5. Two AP groups are configured for termination on each of the local controllers

New ArubaOS 8 Deployment

- MM backed up by a standby MM

- The MM will manage MC1, MC2, MC3, MC4 in Campus1 in addition to MC5, MC6 and MC7 in Campus2
- APs terminate on one of the cluster members in each campus

Migration Procedure

Manual migration requires a complete rebuild of the existing ArubaOS 6 topology by performing the steps listed below. These steps cover termination of APs on a cluster VIP. In the case of a multi-site campus, the APs could terminate on any of the three LMS IPs (MC1, MC2, or MC3)

MM Specific

1. Deploy the MM and perform initial setup
2. Configure licensing on the MM
3. Create a configuration hierarchy and whitelist the MAC addresses of M1, M2, L1-L5 on the MM. Whitelist each device under the following configuration hierarchies:
   - L1, L2, L3, and M1 whitelisted under **Managed Network>Campus1**
   - L4, L5, and M2 whitelisted under **Managed Network>Campus2**
4. Repeat step 1 if a backup MM is being installed
5. Configure MM redundancy if a backup MM is being installed. The MM VIP will be used for configuration management moving forward

Campus1

1. Configure clustering between MC1-MC4. Also enable AP load balancing. GUI: **Managed Network>Campus1>Services>Cluster**
2. Create a VIP (now referred to as "Cluster1 VIP") between the cluster members MC1-MC4. GUI: **Managed Network>Campus1>Services>Redundancy>Virtual Router Table**. Optionally create VIPs for RADIUS COA
3. Create an AP group and SSID. GUI: **Managed Network>Campus1>AP Groups.** GUI: **Managed Network>Campus1>Tasks>Create a new WLAN**
4. Whitelist the Campus1 APs on the MM. This includes mapping them to the appropriate AP group. GUI: **Managed Network>Campus1>Configuration>Access Points>Whitelist**
5. Back up the existing configuration on ArubaOS 6 controllers L1-L3 and M1. GUI: **Maintenance>Backup Flash**
6. Upgrade the image on local L1 to ArubaOS 8 and reboot it. GUI: **Maintenance>Image Management**
7. Provision local L1 to be managed by the MM via the CLI setup dialog. L1 will now become MC1

8. Repeat steps 6-7 to convert L2, L3, and M1 to MC2, MC3, and MC4 respectively
9. In the Campus1 network, point **aruba-master** towards the Cluster1 VIP
10. The APs that were terminating on the L1-L3 will find the cluster VIP, upgrade their images, terminate on one of controllers in the MC1-MC4 range (depending on how the cluster leader load balances the APs), and broadcast the configured SSID for Campus1
11. Connect a wireless client to the SSID and test connectivity
12. Optionally, test seamless client failover by running a voice/video application and disconnecting the user's active controller

Campus2

1. Configure clustering between the MC5, MC6 and MC7 and enable AP load balancing. GUI: **Managed Network>Campus2>Services>Cluster**
2. Create a VIP (now referred to as "Cluster2 VIP") between the cluster members MC5, MC6, and MC7. GUI: **Managed Network>Campus2>Services>Redundancy>Virtual Router Table**. Optionally create VIPs for RADIUS COA
3. Create an AP group and SSID. GUI: **Managed Network>Campus2 > AP Groups.** GUI: **Managed Network>Campus2>Tasks>Create a new WLAN**
4. Whitelist the Campus2 APs on the MM. This includes mapping them to the appropriate AP group. GUI: **Managed Network>Campus1>Configuration>Access Points>Whitelist**
5. Back up the existing configuration on the ArubaOS 6 controllers L4, L5, and M2. GUI: **Maintenance>Backup Flash**
6. Upgrade the image on local L4 to ArubaOS 8 and reboot it. GUI: **Maintenance>Image Management**
7. Provision local L4 to be managed by the MM, via the CLI setup dialog. L4 will now become MC5
8. Repeat steps 6-7 to convert L5 to MC6 and M2 to MC7
9. In the Campus2 network point **aruba-master** towards the Cluster2 VIP
10. The APs that were terminating on the L4 and L5 will find the cluster VIP, upgrade their images, terminate on one of the controllers in the MC5-MC7 range (depending on how the cluster leader load balances the APs), and broadcast the configured SSID for Campus2
11. Connect a wireless client to the SSID and test connectivity
12. Optionally, test seamless client failover by running a voice/video application and disconnecting the user's active controller

MC Master Terminating MCs

Topology

This ArubaOS 8 design consists of a hardware controller deployed as a MC Master (optionally backed up by another MC Master) that manages a group of MCs in different campuses.

This design helps transition deployments to ArubaOS 8 that are unable to deploy a MM. This MC Master topology should eventually be migrated to a MM topology in order to take full advantage of the capabilities offered by ArubaOS 8.

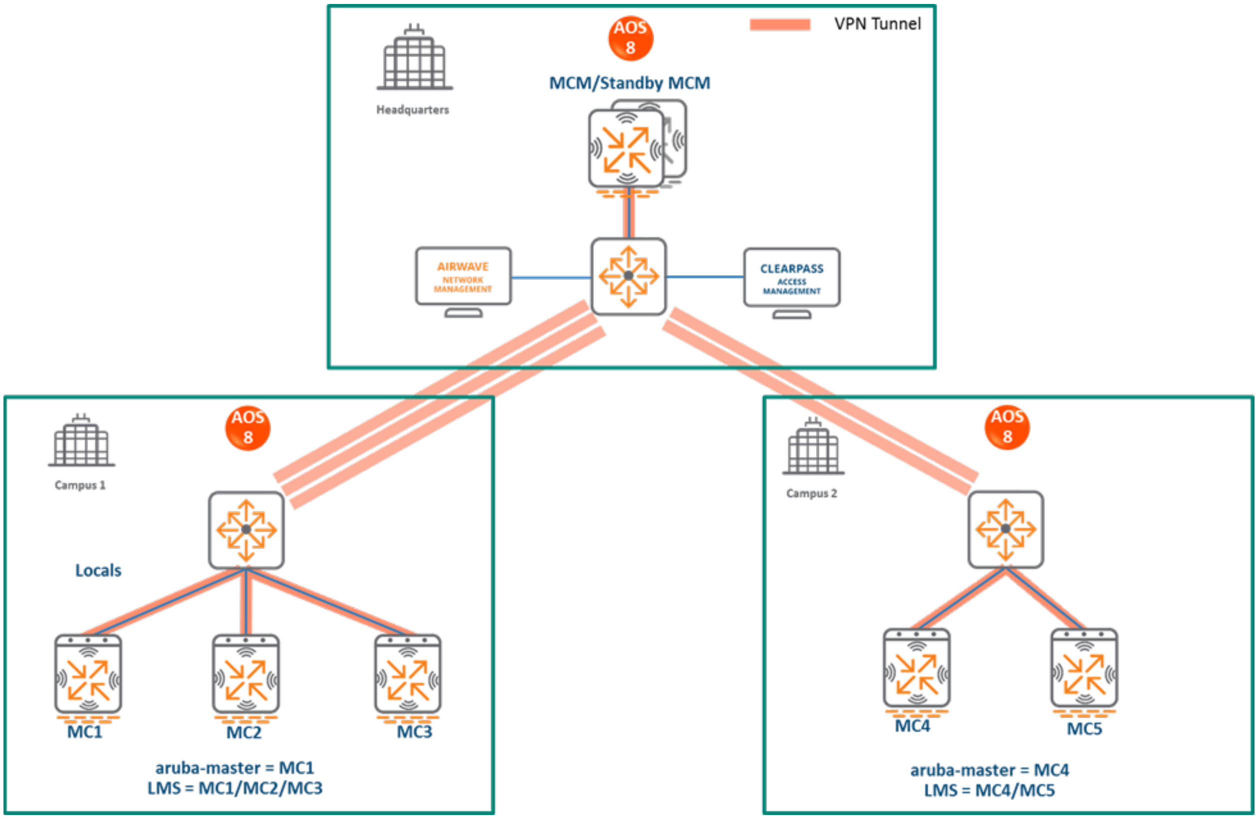

Figure 165 *MC Master Terminating MCs Topology*

In this design:

- The ArubaOS 6 master (M1) and standby master (M2) become ArubaOS 8 MC Masters (MCM1 and MCM2)
- In Campus1, the ArubaOS 6 local controllers (L1, L2, and L3) become ArubaOS 8 MCs (MC1, MC2, and MC3)
- In Campus2, the ArubaOS 6 local controllers (L4 and L5) become ArubaOS 8 MCs (MC4 and MC5)
- All MCs terminate their IPsec tunnels on the MC Master MCM1
- APs terminating on L1, L2, and L3 will now terminate on MC1, MC2, and MC3 respectively
- APs terminating on L4 and L5 will now terminate on MC4 and MC5 respectively

Configuration Hierarchy

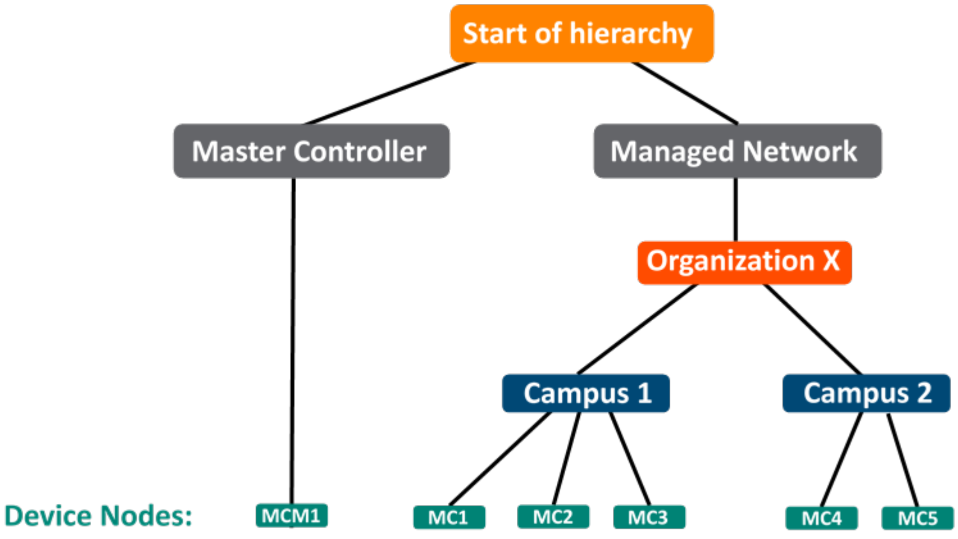

Figure 166 *MC Master Terminating MCs Configuration Hierarchy*

Design Benefits

- A similar topology is maintained in which the MC Master manages the MCs and no additional hardware is required as long as the MC Master is an Aruba 7030 or larger controller
- The hierarchical configuration model offers fully centralized configuration and management of the WLAN
- Additional controllers could be added later and managed by the MC Master

Design Caveats

- Requires purchase of an Aruba 7030 or larger controller to serve as the MC Master as well as the backup MCM if one is not already present
- AP termination on the MC Master is not supported. This has the following impact on AP termination options:
  - ➢ Any APs that are terminating on the master in ArubaOS 6 would need to be redistributed among the locals prior to migration. The locals should have enough capacity to accommodate the additional APs
  - ➢ APs can failover between MCs but cannot failover to the MC Master
- The clustering feature is not supported in a MC Master deployment. AP Fast Failover between MCs is the only controller redundancy option
- AirMatch is not supported

- All controllers in the topology must run the same ArubaOS version
- No centralized monitoring

Migration Requirements

- Verify that the ArubaOS 6 master controller meets the MC Master hardware requirements (Aruba 7030 or any Aruba 7200 series controller)
- Ensure that the ArubaOS 6 master is not terminating any APs as an ArubaOS 8 MC Master does not support AP termination
- Ensure that AP, PEF, and all other licenses have been migrated manually or via the "My Networking Portal"

Migration Options

- Migration can occur manually or via the Migration Tool
- Manual migration steps are detailed below. To perform migration using the Migration Tool, please refer to the ArubaOS 8 Migration Guide

Migration Strategy

Existing ArubaOS 6 Deployment

- **HQ:** M1 and M2
- **Campus1:** L1, L2 and L3
- **Campus2:** L4 and L5
- In Campus1, three AP groups are configured for termination on L1, L2, and L3
- In Campus2, two AP groups are configured for termination onL4 and L5

New ArubaOS 8 Deployment

- MCM1 backed up by MCM2
- MCM1 managing MC1, MC2, and MC3 in Campus1 and MC4 and MC5 in Campus2

Migration Procedure

Manual migration requires a complete rebuild of the existing ArubaOS 6 topology by performing the steps listed below:

MC Master Specific

1. Backup the existing configuration on the ArubaOS 6 masters and locals. GUI: **Maintenance>Backup Flash**
2. Upgrade master M1 to ArubaOS 8 and reboot the controller. GUI: **Maintenance>Image Management**
3. Provision M1 as a MC Master through the CLI setup dialog. M1 will become MCM1
4. Repeat steps 2 and 3 to convert M2 to MCM2
5. Configure master redundancy between MCM1 and MCM2. The MC Master VIP will be used for configuration management moving forward
6. Configure licensing on the MC Master
7. Create a configuration hierarchy on the MC Master and whitelist the MAC addresses of controllers L1-L5. Whitelist each device under the following configuration hierarchies:
   - L1-L3 whitelisted under **Managed Network>Campus1**
   - L4 and L5 whitelisted under **Managed Network>Campus2**

Campus1

1. Create three AP groups, each with the LMS IP of MC1, MC2, and MC3 respectively. GUI: **Managed Network>Campus1>AP Groups**
2. Create a common SSID or one for each AP group. GUI: **Managed Network>Campus1>Tasks>Create a new WLAN**
3. Whitelist the APs on the MC Master. This includes mapping them to the appropriate AP group. GUI: **Managed Network>Campus1>Configuration>Access Points>Whitelist**
4. Upgrade the image on local L1 to ArubaOS 8 and reboot it. GUI: **Maintenance>Image Management**
5. Provision local L1 to be managed by the MC Master via the CLI setup dialog. L1 will become MC1
6. Repeat steps 4-5 for L2 and L3 to convert them to ArubaOS 8 MC2 and MC3
7. Change **aruba-master** to point towards MC1's IP
8. Once MC1 is visible on the MC Master, the APs that were terminating on L1 will find MC1, upgrade their images, download the LMS-IP for MC1, terminate their tunnels on MC1, and broadcast the configured SSID
9. Similarly, the APs on L2 and L3 will be displayed on MC2 and MC3 respectively
10. Connect a wireless client to the SSID and test connectivity
11. Optionally, configure AP Fast Failover via the MC Master to enable AP failover between the MCs

Campus2

1. Create two AP groups, each with LMS IP of MC1 and MC2 respectively. GUI: **Managed Network>Campus2>AP Groups**
2. Create a common SSID or one for each AP group. GUI: **Managed Network>Campus2>Tasks>Create a new WLAN**
3. Whitelist the APs on the MC Master. This includes mapping them to the appropriate AP group. GUI: **Managed Network>Campus2>Configuration>Access Points>Whitelist**
4. Upgrade the image on local L4 to ArubaOS 8 and reboot it. GUI: **Maintenance>Image Management**
5. Provision local L4 to be managed by the MC Master via the CLI setup dialog. L4 will become MC4
6. Repeat steps 4-5 to convert L5 into MC5
7. Change **aruba-master** to point towards MC4's IP
8. The APs that were terminating on L4 will find MC4, upgrade their images, download their LMS IP (i.e. MC4), terminate their tunnels on MC4, and broadcast the configured SSID
9. Similarly, the APs on L5 will be displayed on MC5
10. Connect a wireless client to the SSID and test connectivity
11. Optionally, configure AP Fast Failover via the MC Master to enable AP failover between the MCs

Multiple Master-Locals

This ArubaOS 6 design consists of multiple sites, with the master at each site (typically backed up by a standby master) managing a group of local controllers.

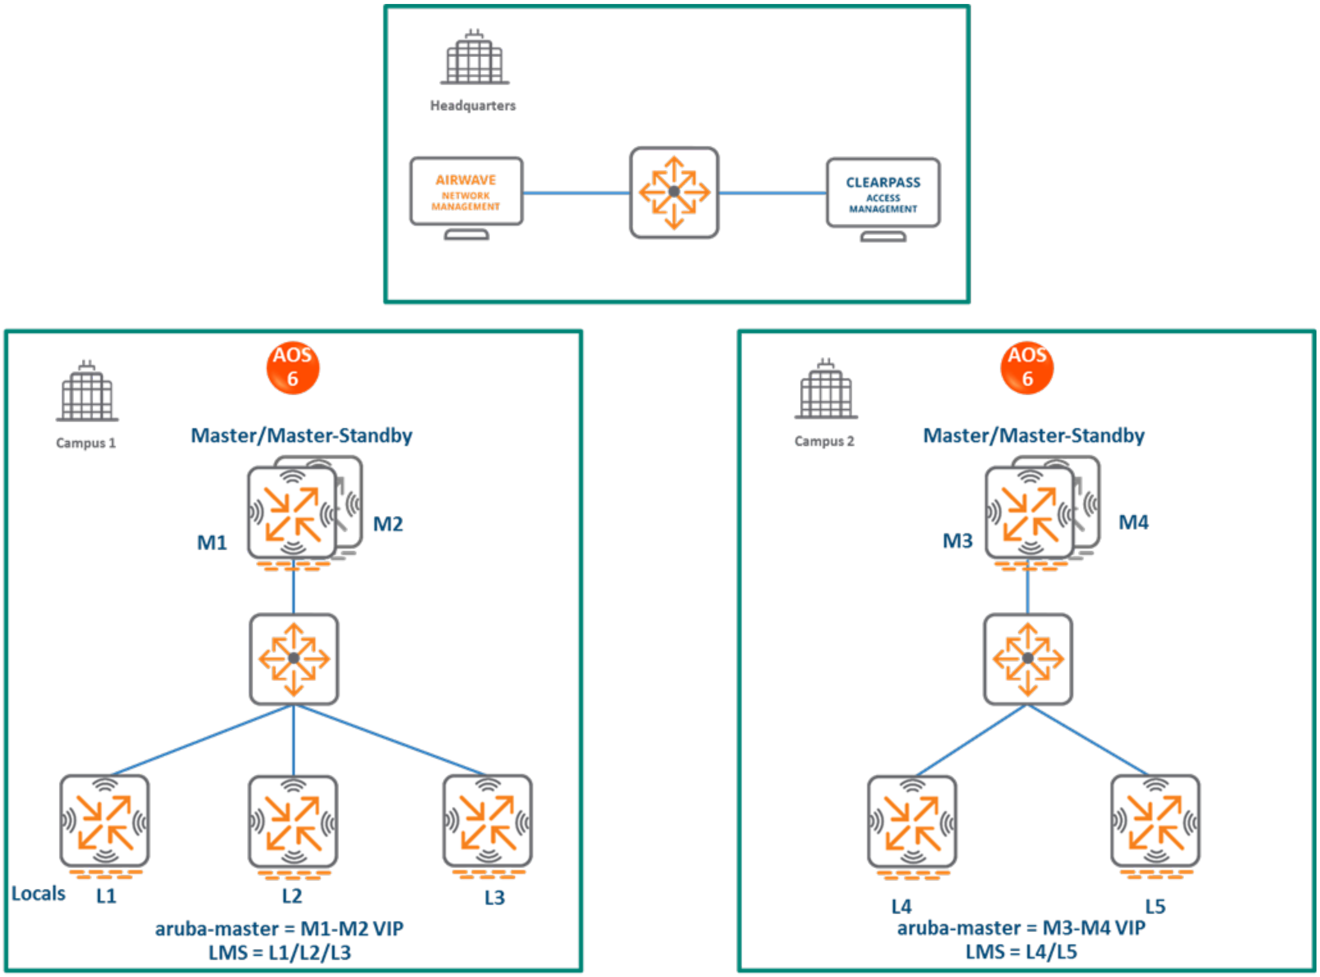

Figure 167 *Multiple Master-Locals*

In this design:

- Each site has its own configuration that is defined on the master and pushed to the respective locals. There is no central point of configuration for multiple sites
- The APs at each site terminate on one of the local controllers with other locals acting as backups. For example, some APs in Campus 1 could terminate on L1, with L2 and L3 providing backup for L1
- AP Fast Failover is configured to provide sub-second failover for the APs when connectivity to their primary controller is lost

MM Terminating MCs

Topology

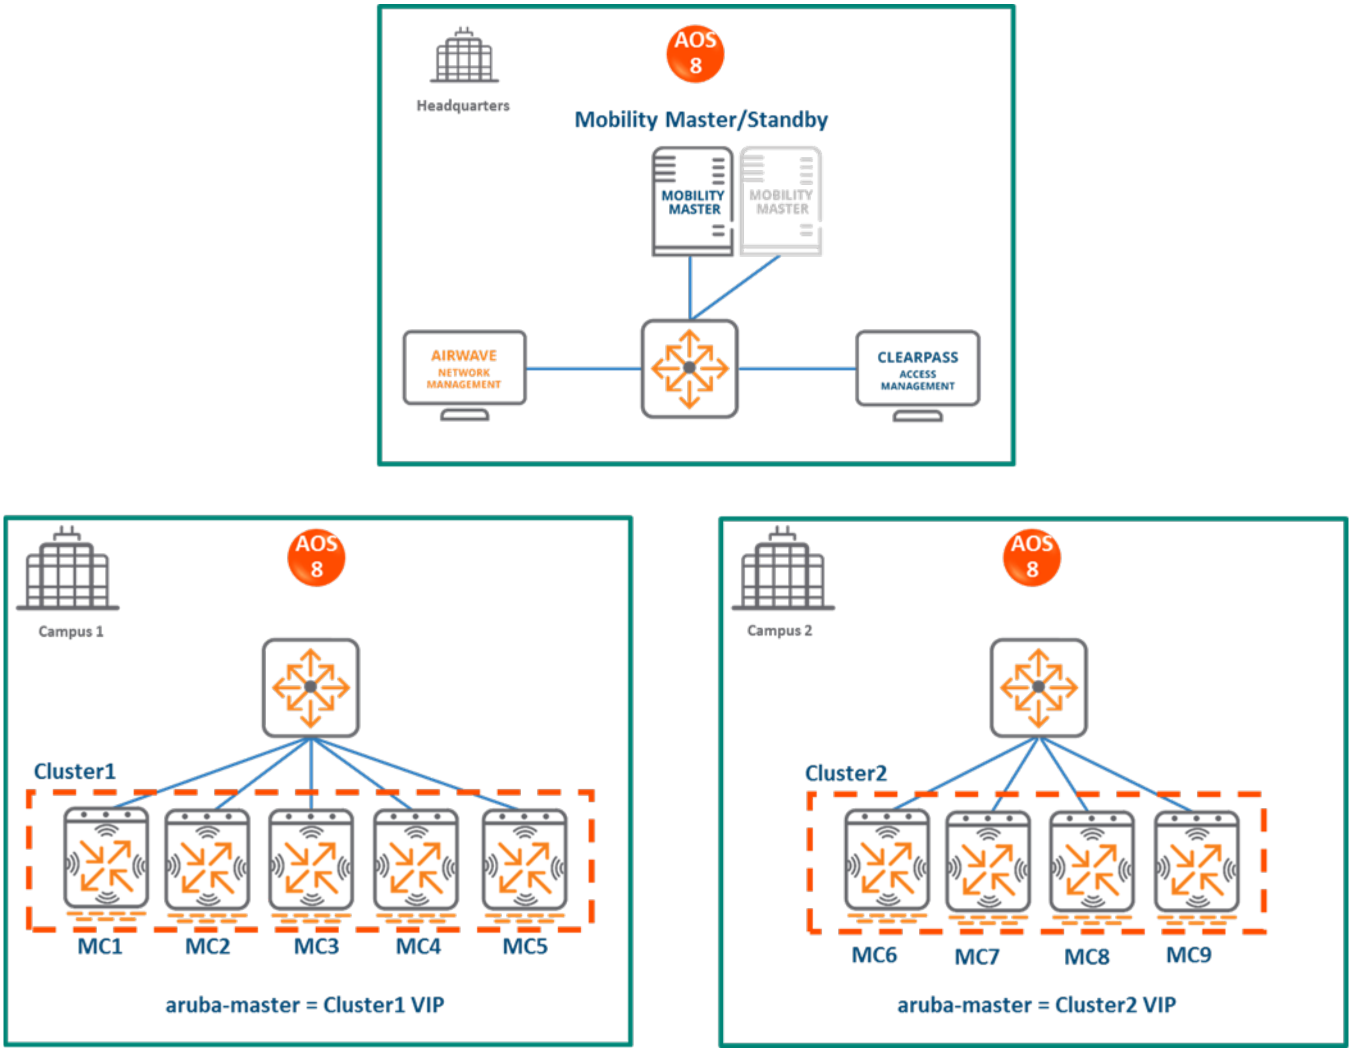

Figure 168 *MM Terminating MCs Topology*

HQ/DC

- A MM (either hardware or virtual) is deployed and configured in the HQ/DC along with a backup MM
- Both campuses are centrally managed by the MM

Campus1

- ArubaOS 6 local controllers L1, L2, and L3 become ArubaOS 8 MC1, MC2, and MC3 respectively
- A cluster is formed between MC1, MC2, and MC3 for controller redundancy, load balancing, and failover of APs and clients

- The ArubaOS 6 masters M1 and M2 become ArubaOS 8 MC4 and MC5
- APs that were terminating on L1, L2 and L3 will now terminate on MC1, MC2, and MC3 respectively
- MC4 and MC5 can be included in the cluster for added redundancy and client and AP load balancing

Campus2

- Similarly, ArubaOS 6 locals L4 and L5 become ArubaOS 8 MC6 and MC7 respectively
- A cluster is formed between MC6 and MC7 for controller redundancy and to load balance and failover APs and clients
- The ArubaOS 6 masters M3 and M4 become ArubaOS 8 MC8 and MC9
- APs that were terminating on L4 and L5 will now terminate on MC6 and MC7 respectively
- MC8 and MC9 can be included in the cluster for added redundancy as well as client and AP load balancing

Configuration Hierarchy

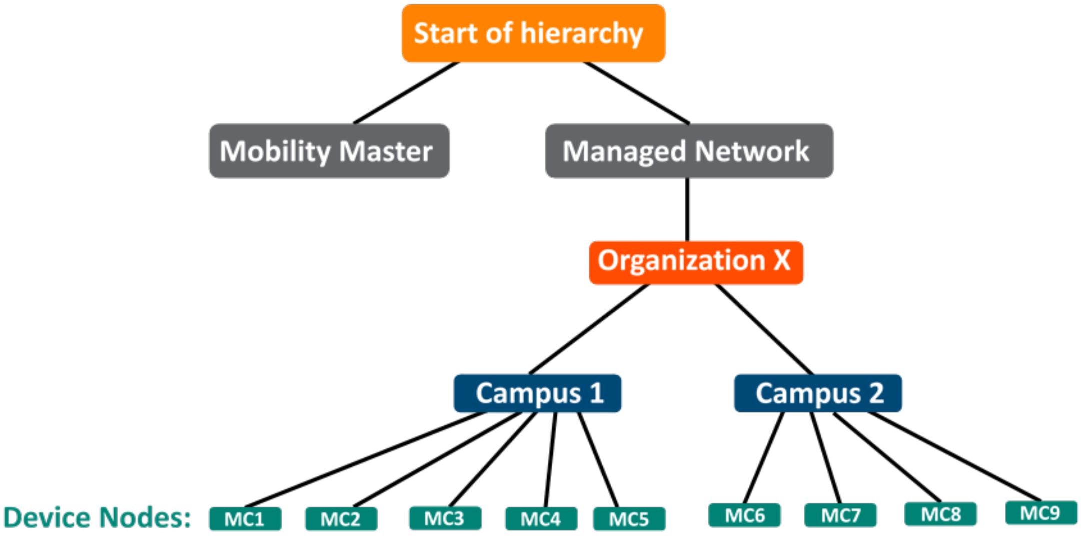

Figure 169 *MM Terminating MCs Configuration Hierarchy*

Design Benefits

- **Maximize benefits** - The MM Terminating MCs design is ideal for fully leveraging the capabilities of ArubaOS 8
- **Scalability** - New controllers can be easily added and managed by the MM
- **Ease of migration** - If an existing deployment has multiple topologies they can all be migrated under the MM into their own nodes in the hierarchy

- **Management** - Centralized configuration and management of controllers
- **Hierarchical configuration model** - Configuration is performed on nodes (logical folders such as USA, West Coast, California, Santa Clara, etc.) and each controller is assigned to a particular node depending on location and context
- **Clustering** - Having controllers in a cluster enables seamless client failover between controllers with no impact to user experience. The clustering feature also facilitates client roaming as well as AP and client load balancing. Clustering is required to support Live Upgrades
- **Live Upgrade** - Upgrade a cluster of controllers in real time without end users experiencing any loss of connectivity or performance. There is no need for scheduling maintenance cycles for network upgrades
- **AirMatch** - RF intelligence is centralized on the MM which significantly improves the RF management and interference mitigation capabilities of the WLAN
- **REST API support**
- **Multi-version support** - Flexibility of upgrading individual controllers to test out new features or bug fixes. Controllers in a cluster need to run the same ArubaOS software version and may be upgraded together
- **In-service software module upgrades** - Loadable Service Modules (LSMs) such as UCC, AirGroup, AppRF etc. can be updated during runtime removing the need to schedule any maintenance cycles

Design Caveats

- The MM does not terminate APs. APs can only be terminated on a MC
- If the existing ArubaOS 6 deployment has more than 1000 controllers and/or 10,000 APs, then migration to an ArubaOS 8 MM deployment requires the deployment of multiple MMs

Migration Requirements

- Requires purchase of virtual MM capacity licenses or the purchase of a hardware MM
- A backup hardware MM may also be deployed in which case the licenses on each MM will be aggregated and synchronized across both MMs
- Other licenses such as AP and PEF need to be migrated manually or via the "My Networking Portal"

Migration Options

- Migration can occur manually or via the Migration Tool
- The Migration Tool is capable of migrating each master-local site to ArubaOS 8 individually. It does not support migration for multiple master-locals at the same time

- Manual migration steps are detailed below. To perform migration using the Migration Tool, please refer to the ArubaOS 8 Migration Guide

Migration Strategy

Existing ArubaOS 6 Deployment

- **Campus1:**
  - Locals L1, L2, L3
  - Masters M1 and M2
  - 3 AP groups are configured to have APs terminate among L1, L2, and L3
- **Campus2**
  - Locals L4 and L5
  - Masters M3 and M4
  - 2 AP groups are configured to have APs terminate among L4 and L5

New ArubaOS 8 Deployment

- MM backed up by a standby MM
- MM managing MC1-MC5 in Campus1 and M6-M9 in Campus2

Migration Procedure

Manual migration requires a complete rebuild of the existing ArubaOS 6 topology by performing the steps listed below:

MM Specific

1. Deploy the MM and perform the initial setup
2. Configure licensing on the MM
3. Create a configuration hierarchy and whitelist the MAC addresses of M1-M4 and L1-L5 on the MM. Whitelist each device under the following configuration hierarchies:
   - M1, M2, L1-L3 whitelisted under **Managed Network>Campus1**
   - M3, M4, L4 and L5 whitelisted under **Managed Network>Campus2**
4. Repeat step 1 if a backup MM is being installed
5. Configure MM redundancy if a backup MM has been installed. The MM VIP will be used for configuration management moving forward

Campus1

6. Configure clustering between MC1-MC5 IPs. Also enable AP load balancing. GUI: **Managed Network>Campus1>Services>Cluster**
7. Create a VIP between the cluster members MC1-MC5. GUI: **Managed Network>Campus1>Services>Redundancy>Virtual Router Table**. Optionally create VIPs for RADIUS COA
8. Create an AP group and SSID. GUI: **Managed Network>Campus1>AP Groups**. GUI: **Managed Network>Campus1>Tasks>Create a new WLAN**
9. Whitelist the Campus1 APs on the MM. This includes mapping them to the appropriate AP group. GUI: **Managed Network>Campus1>Configuration>Access Points>Whitelist**
10. Back up the existing configuration on the ArubaOS 6 masters M1, M2 and locals L1-L3. GUI: **Maintenance>Backup Flash**
11. Upgrade the image on local L1 to ArubaOS 8 and reboot it. GUI: **Maintenance>Image Management**
12. Provision local L1 to be managed by the MM, via the CLI setup dialog. L1 will now become MC1
13. Repeat steps 6-7 to convert L2, L3, M1, and M2 to MC2, MC3, MC4, and MC5 respectively
14. In the Campus1 network point **aruba-master** towards the cluster VIP for MC1-MC5
15. The APs that were terminating on the L1-L3 will find the cluster VIP, upgrade their images, terminate on one of MC1-MC5 (depending on how the cluster leader load balances the APs), and broadcast the configured SSID for Campus1
16. Connect a wireless client to the SSID and test connectivity
17. Optionally, test seamless client failover by running a voice/video application and disconnecting the user's active controller

Campus2

1. Configure clustering between MC6-MC9 IPs and enable AP load balancing. GUI: **Managed Network>Campus2>Services>Cluster**
2. Create a VIP between cluster members MC6-MC9. GUI: **Managed Network>Campus2> Services>Redundancy>Virtual Router Table**. Optionally create VIPs for RADIUS COA
3. Create an AP group and SSID. GUI: **Managed Network>Campus2>AP Groups**. GUI: **Managed Network>Campus2>Tasks>Create a new WLAN**
4. Whitelist the Campus2 APs on the MM. This includes mapping them to the appropriate AP group. GUI: **Managed Network>Campus2>Configuration>Access Points>Whitelist**
5. Back up the existing configuration on the ArubaOS 6 masters M3 and M4 as well as locals L4 and L5. GUI: **Maintenance>Backup Flash**

6. Upgrade the image on local L4 to ArubaOS 8 and reboot the device. GUI: **Maintenance>Image Management**
7. Provision local L4 to be managed by the MM via the CLI setup dialog. L4 will now become MC6
8. Repeat steps 6-7 to convert L5, M3, and M4 to MC7, MC8, and MC9 respectively
9. In the Campus2 network, point **aruba-master** towards the cluster VIP for MC6-MC9
10. The APs that were terminating on the L4 and L5 will find the cluster VIP, upgrade their images, terminate on one of MC6-MC9 (depending on how the cluster leader load balances the APs), and broadcast the configured SSID for Campus2
11. Connect a wireless client to the SSID and test connectivity
12. Optionally, test seamless client failover by running a voice/video application and disconnecting the user's active controller

MC Master Terminating MCs

Topology

In this ArubaOS 8 design, each site consists of a hardware controller deployed as a MC Master (optionally backed up by another MC Master) that manages a group of MCs.

This design helps transition deployments to ArubaOS 8 that are unable to deploy a MM. This MC Master topology should eventually be migrated to a MM topology in order to take full advantage of the capabilities offered by ArubaOS 8.

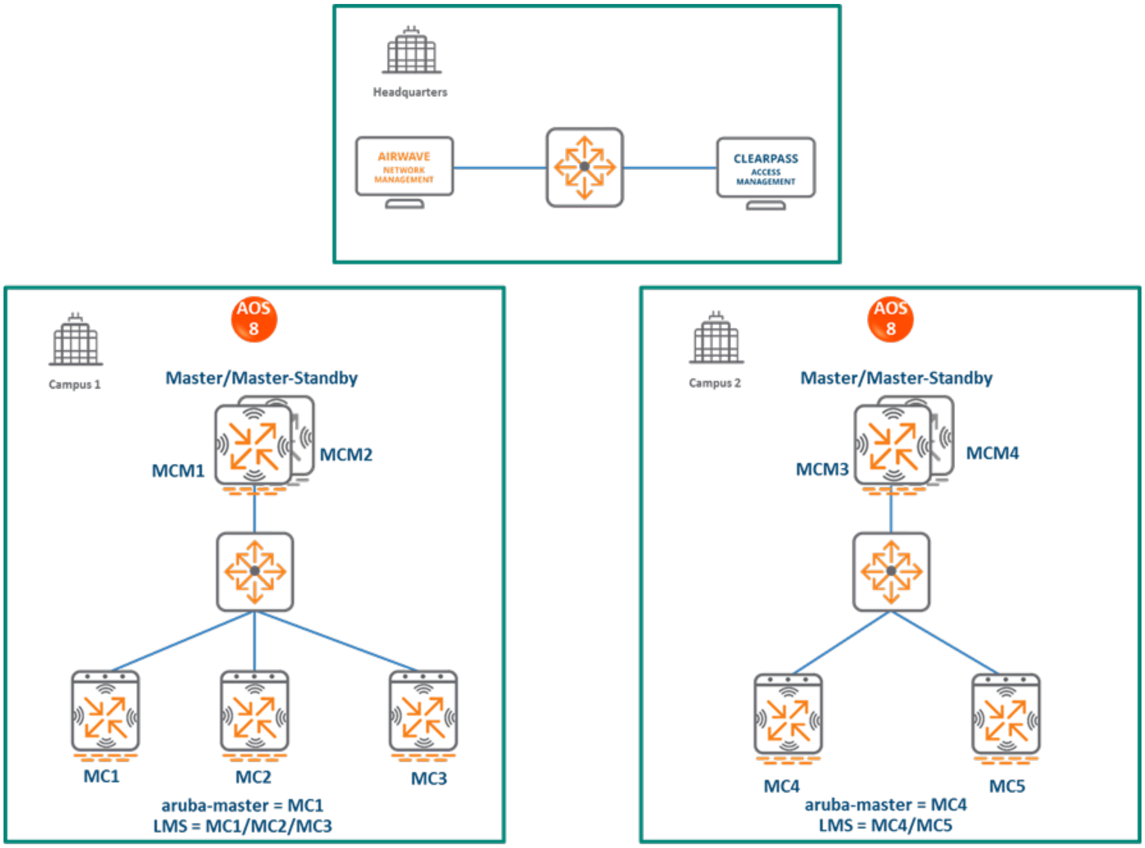

Figure 170 *MC Master Terminating MCs Topology*

In this design, each campus is still managed by its own MC Master.

Campus1

- ArubaOS 6 locals L1, L2, and L3 become ArubaOS 8 MC1, MC2, and MC3 respectively
- The ArubaOS 6 masters M1 and M2 become ArubaOS 8 MCM1 and MCM2
- APs that were terminating on L1, L2, and L3 will now terminate on MC1, MC2 and MC3 respectively

Campus2

- ArubaOS 6 locals L4 and L5 become ArubaOS 8 MC4 and MC5 respectively
- The ArubaOS 6 masters M3 and M4 become ArubaOS 8 MCM3 and MCM4
- APs that were terminating on L4 and L5 will now terminate on MC4 and MC5 respectively

Configuration Hierarchy

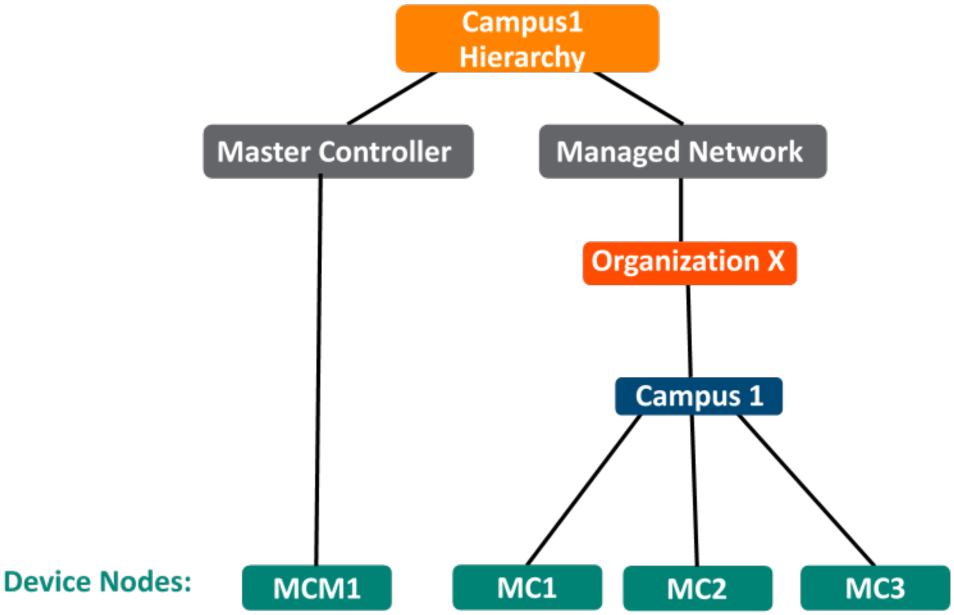

Figure 171 *MC Master Terminating Mobility Configuration Hierarchy Campus 1*

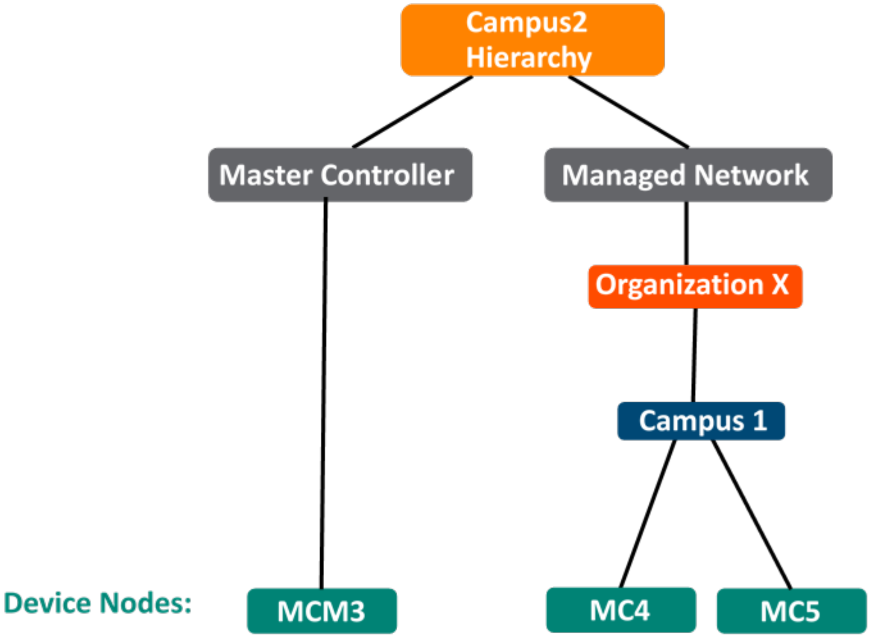

Figure 172 *MC Master Terminating Mobility Configuration Hierarchy Campus 2*

Design Benefits

- A similar topology is maintained in which the MC Master manages the MCs and no additional hardware is required as long as the MC Master is an Aruba 7030 or larger controller

- The hierarchical configuration model offers fully centralized configuration and management of the WLAN
- Additional controllers could be added later and managed by the MC Master

Design Caveats

- Requires purchase of an Aruba 7030 or larger controller to serve as the MC Master as well as the backup MCM if one is not already present
- AP termination on the MC Master is not supported. This has the following impact on AP termination options:
  - Any APs that are terminating on the master in ArubaOS 6 would need to be redistributed among the locals prior to migration. The locals should have enough capacity to accommodate the additional APs
  - APs can failover between MCs but cannot failover to the MC Master
- The clustering feature is not supported in a MC Master deployment. AP Fast Failover between MCs is the only controller redundancy option
- AirMatch is not supported
- All controllers in the topology must run the same ArubaOS version
- No centralized monitoring

Migration Requirements

- Verify that the ArubaOS 6 master controller meets the MC Master hardware requirements (Aruba 7030 or any Aruba 7200 series controller)
- Ensure that the ArubaOS 6 master is not terminating any APs as an ArubaOS 8 MC Master does not support AP termination
- Ensure that AP, PEF, and all other licenses have been migrated manually or via the "My Networking Portal"

Migration Options

- Migration can occur manually or via the Migration Tool
- Manual migration steps are detailed below. To perform migration using the Migration Tool, please refer to the ArubaOS 8 Migration Guide

Migration Strategy

Existing ArubaOS 6 Deployment

- Locals L1, L2, L3
- Masters M1 and M2
- 3 AP groups are configured to have groups of APs terminate among L1, L2, and L3

New ArubaOS 8 Deployment

- **Campus1:**
  - MCM1 backed up by MC2
  - MCM1 managing MC1, MC2, and MC3
- **Campus2:**
  - MCM3 backed up by MCM4
  - MCM3 managing MC4 and MC5

Migration Procedure

Manual migration requires a complete rebuild of the existing ArubaOS 6 topology by performing the steps listed below:

Campus1

1. Backup the existing configuration on the ArubaOS 6 masters and locals. GUI: **Maintenance> Backup Flash**
2. Upgrade master M1 to ArubaOS 8 and reboot the controller
3. Provision M1 as a MC Master through the CLI setup dialog. M1 will now become MCM1
4. Repeat steps 2 and 3 to convert M2 to MCM2
5. Configure master redundancy between MCM1 and MCM2. The MC Master VIP will be used for configuration management going forward
6. Configure licensing on the MC Master
7. Create a configuration hierarchy on the MC Master and whitelist the MAC addresses of controllers L1-L3
   - L1-L3 will be whitelisted under **Managed Network>Campus1**
8. Create three AP groups, each with LMS IP of MC1, MC2, and MC3 respectively. GUI: **Managed Network>Campus1>AP Groups**
9. Create a common SSID or one for each AP group. GUI: **Managed Network>Campus1>Tasks>Create a new WLAN**

10. Whitelist the APs on the MC Master. This includes mapping them to the appropriate AP group. GUI: **Managed Network>Campus1>Configuration>Access Points>Whitelist**
11. Upgrade the image on local L1 to ArubaOS 8 and reboot it. GUI: **Maintenance>Image Management**
12. Provision local L1 to be managed by the MC Master via the CLI setup dialog. L1 now becomes MC1
13. Repeat steps 11-12 for L2 and L3 to convert them to ArubaOS 8 MC2 and MC3
14. Change **aruba-master** to MC1's IP
15. Once MC1 is visible on the MC Master, the APs that were terminating on L1 will find MC1, upgrade their images, download the LMS-IP for MC1, terminate their tunnels on MC1, and broadcast the configured SSID
16. Similarly, the APs on L2 and L3 will be displayed on MC2 and MC3 respectively
17. Connect a wireless client to the SSID and test connectivity
18. Optionally, configure AP Fast Failover via the MC Master to enable sub-second AP failover between the MCs

Campus2

1. Backup the existing configuration on the ArubaOS 6 masters and locals. GUI: **Maintenance>Backup Flash**
2. Upgrade master M3 to ArubaOS 8 and reboot the controller
3. Provision M3 as a MC Master through the CLI setup dialog. M3 will now become MCM3
4. Repeat steps 2 and 3 to convert M4 to MCM4
5. Configure master redundancy between MCM3 and MCM4. Going forward, use the MC Master VIP for configuration management
6. Configure licensing on the MC Master
7. Create a configuration hierarchy on the MC Master and whitelist the MAC addresses of controllers L4 and L5
   - L4, L5 will be whitelisted under **Managed Network>Campus2**
8. Create two AP groups, each with LMS IP of MC4 and MC5 respectively. GUI: **Managed Network>Campus2>AP Groups**
9. Create a common SSID or one for each AP group. GUI: **Managed Network>Campus2>Tasks>Create a new WLAN**
10. Whitelist your APs on the MC Master. This includes mapping them to the appropriate AP group. GUI: **Managed Network>Campus2>Configuration>Access Points>Whitelist**
11. Upgrade the image on local L4 to ArubaOS 8 and reboot it. GUI: **Maintenance>Image Management**

12. Provision local L4 to be managed by the MC Master via the CLI setup dialog. L4 now becomes MC4
13. Repeat steps 11-12 to convert L5 into MC5
14. Change **aruba-master** to MC4's IP
15. The APs that were terminating on L4 will find MC4, upgrade their images, download their LMS IP (i.e. MC4), terminate their tunnels on MC4, and broadcast the configured SSID
16. Similarly, the APs on L5 will show up on MC5 respectively
17. Connect a wireless client to the SSID and test connectivity
18. Optionally, configure AP Fast Failover via the MC Master to enable sub-second AP failover between the MCs

All Masters

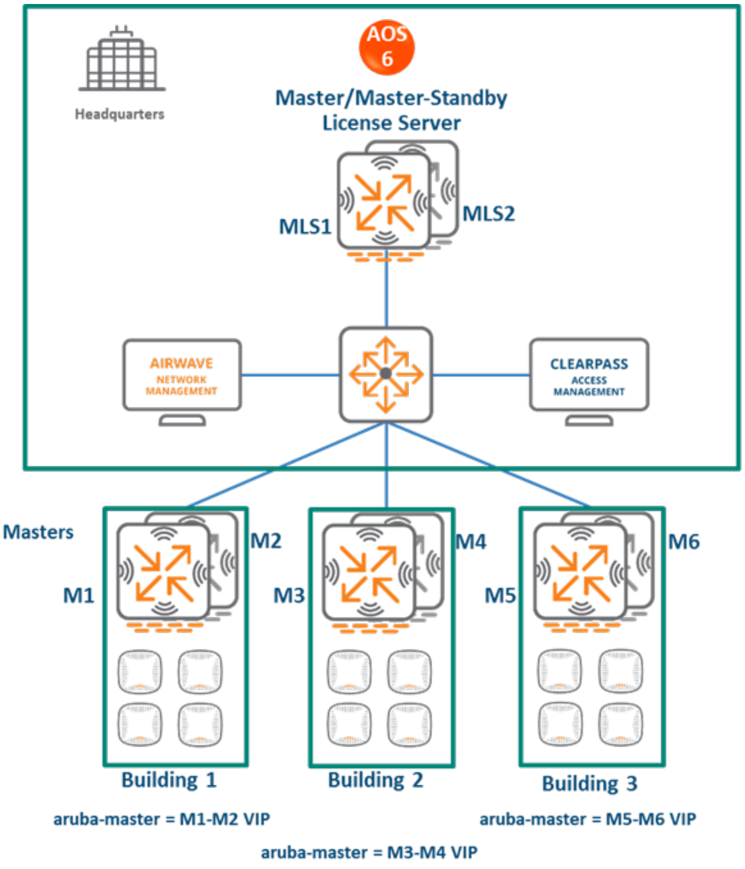

Figure 173 *All Masters Topology*

- In this ArubaOS 6 design, each site is managed by its own master controller, backed up by a standby master
- There is a separate master/standby pair that functions as the license server for all sites.
- All the site masters are centrally managed by AirWave

- The all-master design is typically deployed at sites that need to run different ArubaOS versions (for example, to test new ArubaOS features

MM Terminating MCs

Topology

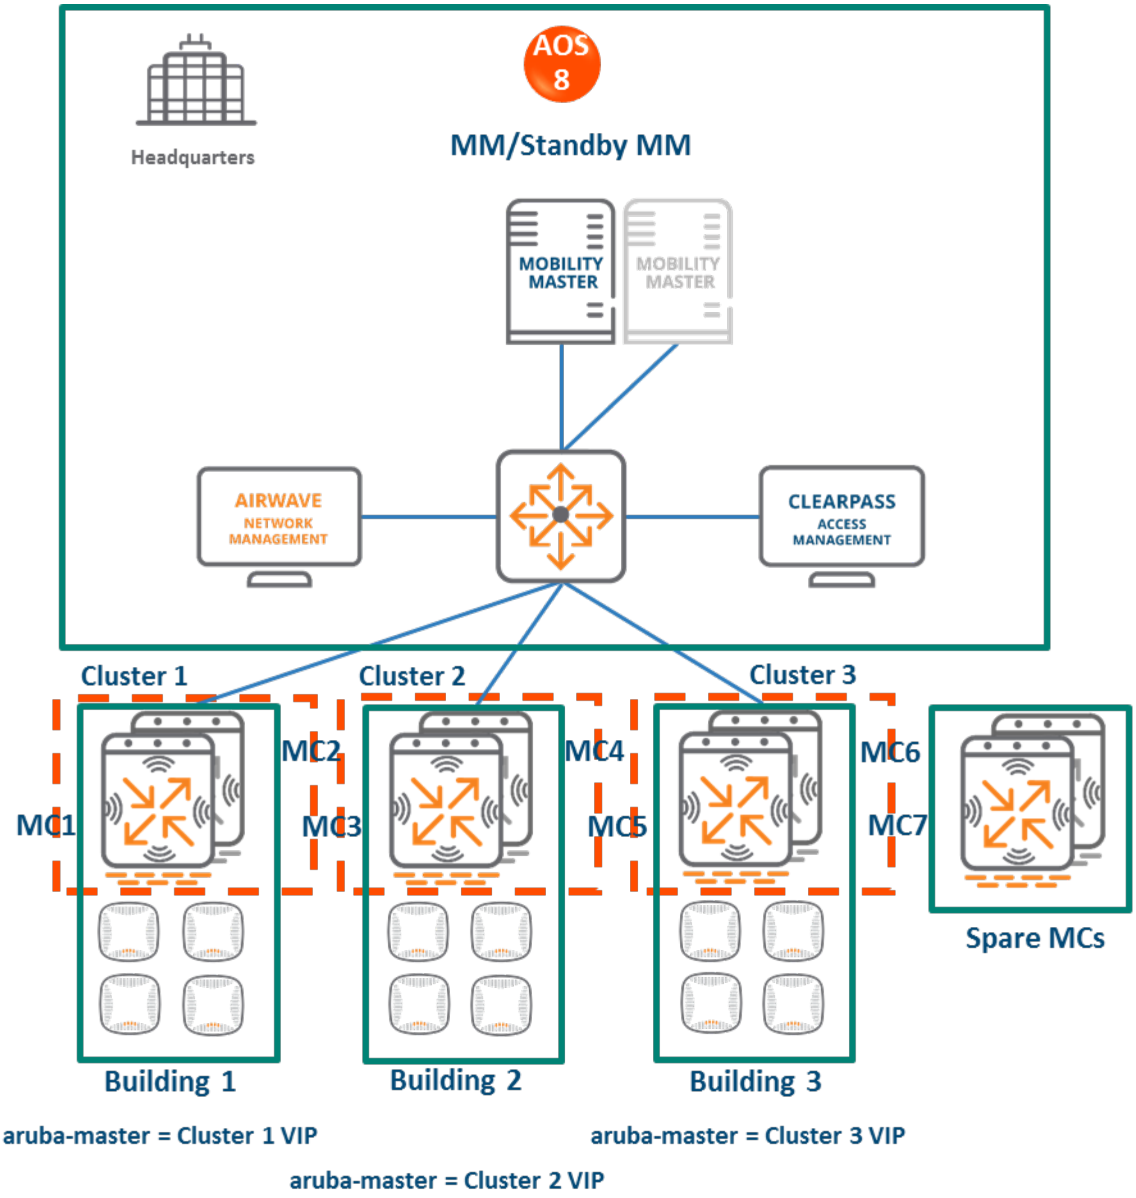

Figure 174 *MM Terminating MCs Topology*

- **HQ/DC:**
  - A MM (either hardware or virtual) is deployed and configured along with a backup MM.
  - All site controllers are centrally managed by the MM
- **Building1:**

  - The ArubaOS 6 master and standby master become ArubaOS 8 MC1 and MC2
  - A cluster can be formed between MC1 and MC2 for controller redundancy as well as client and AP load balancing
- **Building2:** The ArubaOS 6 masters become ArubaOS 8 MC3 and MC4. Both MCs can become cluster members
- **Building3:** The ArubaOS 6 masters become ArubaOS 8 MC5 and MC6. Both MCs can become cluster members
- **License servers:**
  - The ArubaOS 6 master and standby master that were previously being used as licensing servers become MCs managed by the MM
  - These MCs can be repurposed. E.g., they can be used as staging controllers to redirect APs in each site to their LMS controllers, or they can be added to the cluster at any site to provide additional controller redundancy as well as client and AP load balancing

Configuration Hierarchy

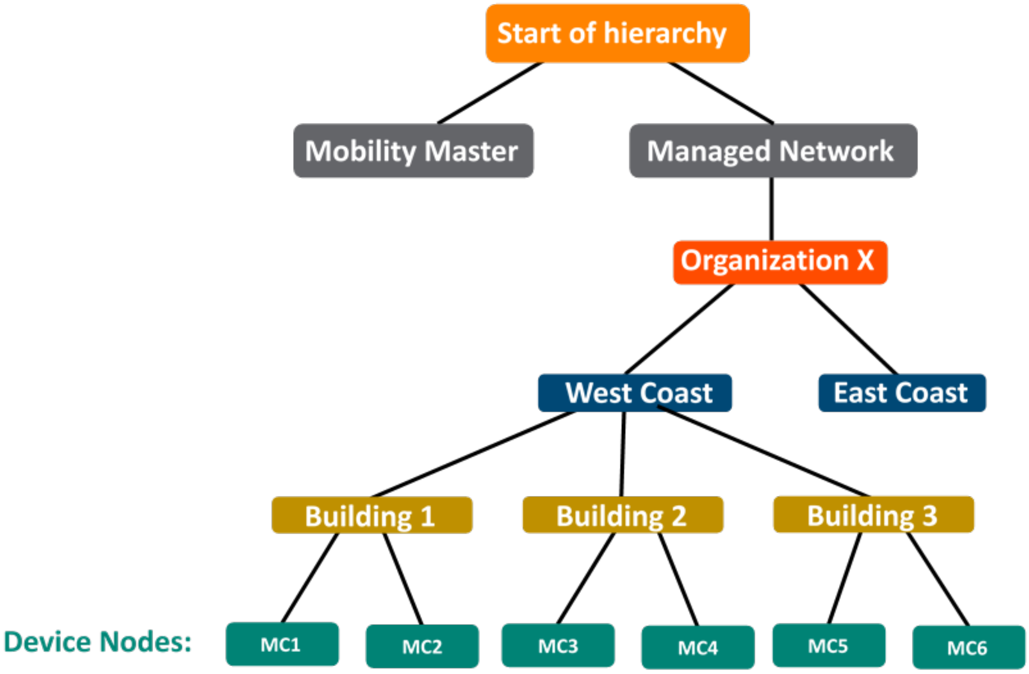

Figure 175 *MM Terminating MCs Configuration Hierarchy*

Design Benefits

- **Maximize benefits** - The MM Terminating MCs design is ideal for fully leveraging the capabilities of ArubaOS 8
- **Scalability** - New controllers can be easily added and managed by the MM
- **Ease of migration** - If an existing deployment has multiple topologies they can all be migrated under the MM into their own nodes in the hierarchy

- **Management** - Centralized configuration and management of controllers
- **Hierarchical configuration model** - Configuration is performed on nodes (logical folders such as USA, West Coast, California, Santa Clara, etc.) and each controller is assigned to a particular node depending on location and context
- **Clustering** - Having controllers in a cluster enables seamless client failover between controllers with no impact to user experience. The clustering feature also facilitates client roaming as well as AP and client load balancing. Clustering is required to support Live Upgrades.
- **Live Upgrade** - Upgrade a cluster of controllers in real time without end users experiencing any loss of connectivity or performance. There is no need for scheduling maintenance cycles for network upgrades
- **AirMatch** - RF intelligence is centralized on the MM which significantly improves the RF management and interference mitigation capabilities of the WLAN
- **REST API support**
- **Multi-version support** - Flexibility of upgrading individual controllers to test out new features or bug fixes. Controllers in a cluster need to run the same ArubaOS software version and may be upgraded together
- **In-service software module upgrades** - Loadable Service Modules (LSMs) such as UCC, AirGroup, AppRF etc. can be updated during runtime removing the need to schedule any maintenance cycles

Design Caveats

- The MM does not terminate APs. APs can only be terminated on a MC
- If the existing ArubaOS 6 deployment has more than 1000 controllers and/or 10,000 APs, then migration to an ArubaOS 8 MM deployment requires the deployment of multiple MMs

Migration Requirements

- Requires purchase of virtual MM capacity licenses or the purchase of a hardware MM
- A backup hardware MM may also be deployed in which case the licenses on each MM will be aggregated and synchronized across both MMs
- Other licenses such as AP and PEF need to be migrated manually or via the "My Networking Portal"

Migration Options

- Migration can be done manually or via the Migration Tool
- Manual migration steps are detailed below. To perform migration using the Migration Tool, refer to the ArubaOS 8 Migration Guide

Migration Strategy

Existing ArubaOS 6 Deployment

- **Building1:** Masters M1-M2
- **Building2:** Masters M3-M4
- **Building3:** Masters M5-M6
- **License servers:** Masters MLS1 and MLS2

New ArubaOS 8 Deployment

- MM managing MC1, MC2, MC3, MC4, MC5, MC6, MC7, and MC8.

Migration Procedure

Manual migration requires a complete rebuild of the existing ArubaOS 6 topology by performing the steps listed below:

MM Specific

1. Deploy the MM and perform initial setup
2. Configure licensing on the MM
3. Create a configuration hierarchy and whitelist the MAC addresses of M1-M6 on the MM. Whitelist each device under the following configuration hierarchies:
   - M1, M2 whitelisted under **Managed Network>Building1**
   - M3, M4 whitelisted under **Managed Network>Building2**
   - M5, M6 whitelisted under **Managed Network>Building3**
4. Repeat step 1 if you are installing a backup MM
5. Configure MM redundancy if a backup MM has been installed. The MM VIP will be used for configuration management moving forward

Building 1

1. Configure clustering between MC1 and MC2 IPs and enable AP load balancing. GUI: **Managed Network>Building1>Services>Cluster**
2. Create a VIP between the cluster members MC1 and MC2. GUI: **Managed Network>Building1>Services>Redundancy>Virtual Router Table**. Optionally create VIPs for RADIUS COA

3. Create an AP group and SSID. GUI: **Managed Network>Building1>AP Groups**. GUI: **Managed Network>Building1>Tasks>Create a new WLAN**
4. Whitelist the Building1 APs on the MM. This includes mapping them to the appropriate AP group. GUI: **Managed Network>Building1>Configuration>Access Points>Whitelist**
5. Back up the existing configuration on the ArubaOS 6 masters. GUI: **Maintenance>Backup Flash**
6. Upgrade the image on local M1 to ArubaOS 8 and reboot it. GUI: **Maintenance>Image Management**
7. Provision local M1 to be managed by the MM via the CLI setup dialog. M1 will now become MC1
8. Repeat steps 6-7 to convert M2 to MC2
9. In Building1, point **aruba-master** towards the cluster VIP for MC1 and MC2
10. The APs that were terminating on M1 will find the cluster VIP, upgrade their images, terminate on either MC1 or MC2 (depending on how the cluster leader load balances the APs), and broadcast the configured SSID for Building1
11. Connect a wireless client to the SSID and test connectivity
12. Optionally, test seamless client failover by running a voice/video application and disconnecting the user's active controller

Building 2

1. Configure clustering between MC3 and MC4 IPs and enable AP load balancing. GUI: **Managed Network>Building2>Services>Cluster**
2. Create a VIP between the cluster members MC3 and MC4. GUI: **Managed Network>Building2>Services>Redundancy>Virtual Router Table**. Optionally create VIPs for RADIUS COA
3. Create an AP group and SSID. GUI: **Managed Network>Building2>AP Groups**. GUI: **Managed Network>Building2>Tasks>Create a new WLAN**
4. Whitelist the Building2 APs on the MM. This includes mapping them to the appropriate AP group. GUI: **Managed Network>Building2>Configuration>Access Points>Whitelist**
5. Back up the existing configuration on the ArubaOS 6 masters. GUI: **Maintenance>Backup Flash**
6. Upgrade the image on master M3 to ArubaOS 8 and reboot it. GUI: **Maintenance>Image Management**
7. Provision local M3 to be managed by the MM via the CLI setup dialog. M3 will now become MC3
8. Repeat steps 6-7 to convert M4 to MC4
9. In Building2 network, point **aruba-master** towards the cluster VIP for MC3 and MC4

10. The APs that were terminating on M3 will find the cluster VIP, upgrade their images, terminate on either MC3 or MC4 (depending on how the cluster leader load balances the APs), and broadcast the configured SSID for Building 2
11. Connect a wireless client to the SSID and test connectivity
12. Optionally, test seamless client failover by running a voice/video application and disconnecting the user's active controller

Follow similar steps for Building3.

Spares MC7 and MC8

- These can be relocated to any of the sites to be repurposed as cluster members for added controller redundancy as well as AP and client load balancing

Master and Branch Controllers

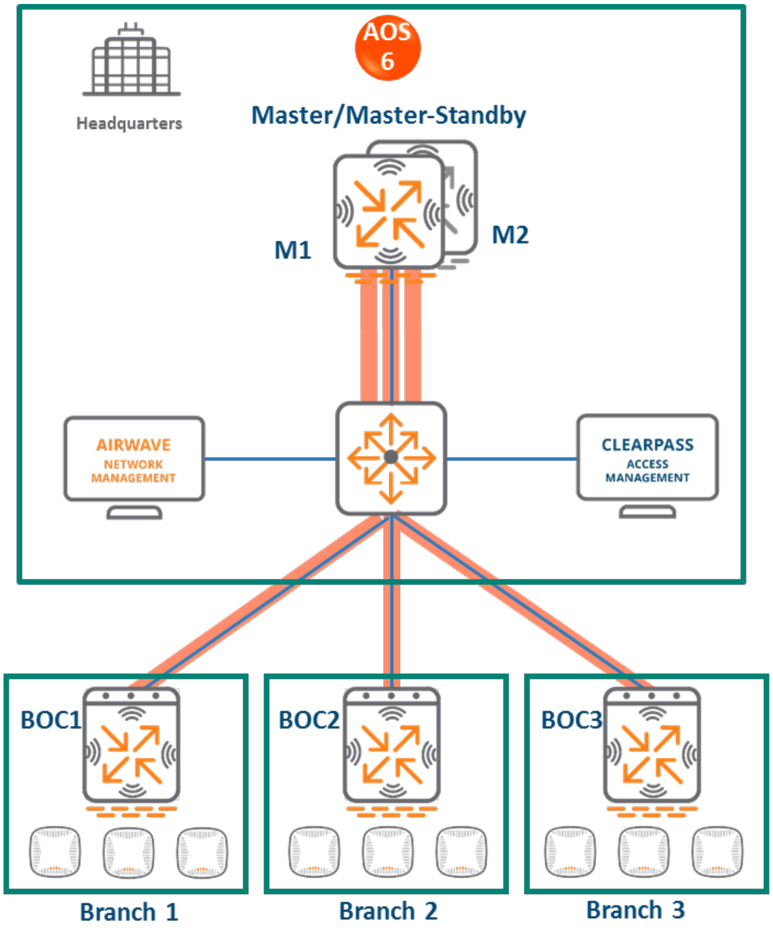

Figure 176 *Master and Branch Controllers Topology*

In this ArubaOS 6 design:

- A master controller manages geographically distributed branches
- The master controller is backed up by a second master controller for redundancy
- Each branch consists of one or more 7000 series controllers i.e. branch controllers/Branch Office Controllers (BOCs)
- Each branch controller uses ZTP to discover and build an IPsec connection with the master controller
- Configuration for the branch controllers is managed on the master

MM Terminating MCs

Topology

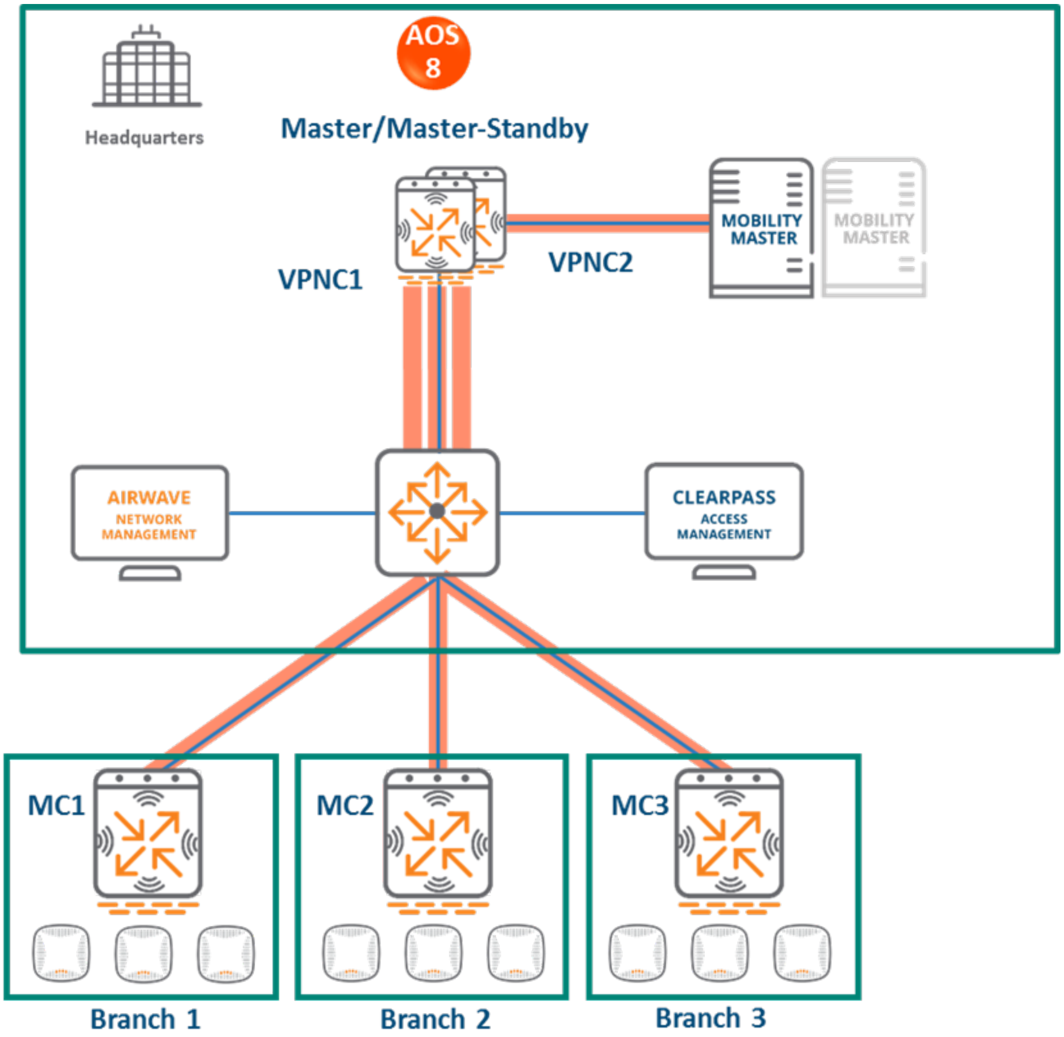

Figure 177 *MM Terminating MCs Topology*

In this design:

- A MM (either hardware or virtual) is deployed and configured, along with a backup MM
- Each ArubaOS 6 BOC (BOC1, BOC2, BOC3) becomes an ArubaOS 8 MC (MC1, MC2, MC3)
- The ArubaOS 6 master (M1) and standby master (M2) become two ArubaOS 8 VPNC MCs (MC4 and MC5)
- Branch MCs are capable of termination on the MM. However, using VPNCs is highly recommended if a deployment consists of distributed branches and user traffic originating from branches needs to reach corporate resources within HQ. User traffic requiring HQ access will be relatively high bandwidth and encryption/decryption is CPU intensive. Using VPNCs helps insulate the MM from the increased load
- APs terminating on L1, L2, and L3 will now terminate on MC1, MC2, and MC3 respectively

Configuration Hierarchy

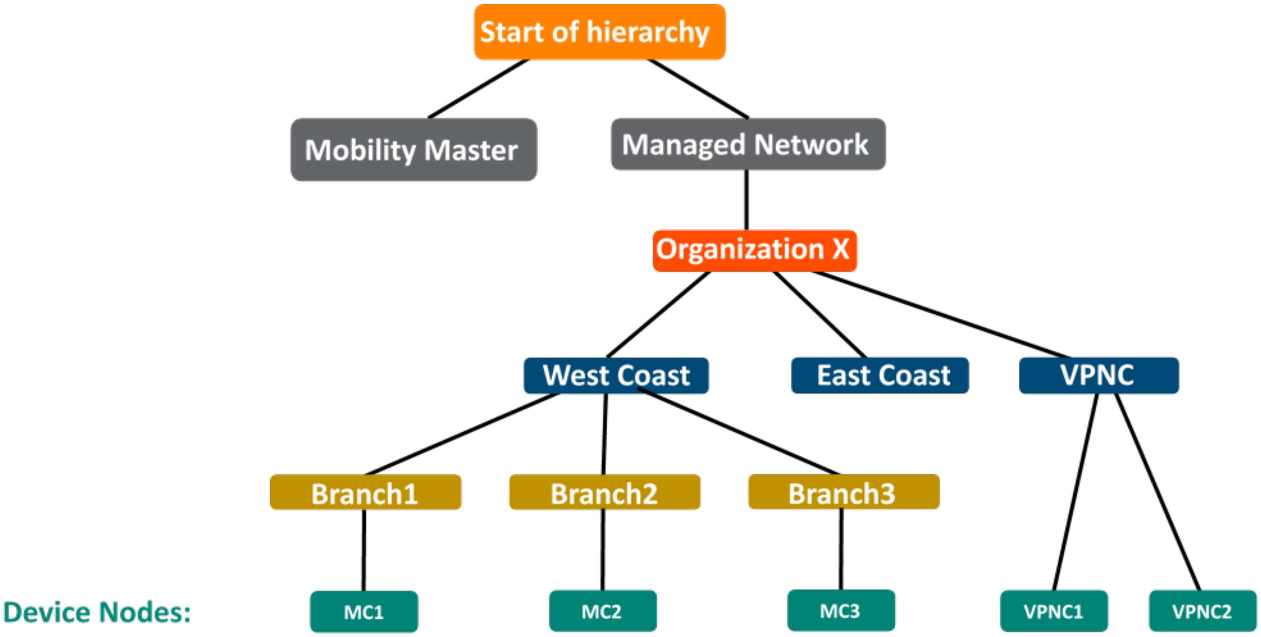

Figure 178 *MM Terminating MCs Configuration Hierarchy*

Design Benefits

- **Maximize benefits** - The MM Terminating MCs design is ideal for fully leveraging the capabilities of ArubaOS 8
- **Scalability** - New controllers can be easily added and managed by the MM
- **Ease of migration** - If an existing deployment has multiple topologies they can all be migrated under the MM into their own nodes in the hierarchy
- **Management** - Centralized configuration and management of controllers

- **Hierarchical configuration model** - Configuration is performed on nodes (logical folders such as USA, West Coast, California, Santa Clara, etc.) and each controller is assigned to a particular node depending on location and context
- **Clustering** - Having controllers in a cluster enables seamless client failover between controllers with no impact to user experience. The clustering feature also facilitates client roaming as well as AP and client load balancing. Clustering is required to support Live Upgrades
- **Live Upgrade** - Upgrade a cluster of controllers in real time without end users experiencing any loss of connectivity or performance. There is no need for scheduling maintenance cycles for network upgrades
- **AirMatch** - RF intelligence is centralized on the MM which significantly improves the RF management and interference mitigation capabilities of the WLAN
- **REST API support**
- **Multi-version support** - Flexibility of upgrading individual controllers to test out new features or bug fixes. Controllers in a cluster need to run the same ArubaOS software version and may be upgraded together
- **In-service software module upgrades** - Loadable Service Modules (LSMs) such as UCC, AirGroup, and AppRF can be updated during runtime removing the need to schedule any maintenance cycles.

Design Caveats

- The MM does not terminate APs. APs can only be terminated on a MC
- If the existing ArubaOS 6 deployment has more than 1000 controllers and/or 10,000 APs, then migration to an ArubaOS 8 MM deployment requires the deployment of multiple MMs

Migration Requirements

- Requires purchase of virtual MM capacity licenses or the purchase of a hardware MM
- A backup hardware MM may also be deployed in which case the licenses on each MM will be aggregated and synchronized across both MMs
- Other licenses such as AP and PEF need to be migrated manually or via the "My Networking Portal"

Migration Options

- Migration can occur manually or via the Migration Tool
- The Migration Tool is capable of migrating each master-local site to ArubaOS 8 individually. It does not support migration for multiple master-locals at the same time

- Manual migration steps are detailed below. To perform migration using the Migration Tool, please refer to the ArubaOS 8 Migration Guide

Migration Strategy

Existing ArubaOS 6 Deployment

- **Branch1:** BOC1
- **Branch2:** BOC2
- **Branch3:** BOC3
- **HQ:** M1, M2

New ArubaOS 8 Deployment

- MM managing MC1, MC2, MC3 and VPNC1, VPNC2.

Migration Procedure

Manual migration requires a complete rebuild of the existing ArubaOS 6 topology. Use the following steps to perform manual migration of a branch network. The ArubaOS8 Branch Network ASE recipe may also be used to understand MM/VPNC/branch controller configuration in ArubaOS 8.

MM

1. Deploy the MM and perform initial setup
2. Configure licensing on the MM
3. Repeat step 1 if a backup MM is being installed
4. Configure MM redundancy if a backup MM has been installed. The MM VIP will be used for configuration management moving forward
5. Configure Activate, a configuration hierarchy, VPN peers, and whitelist the MAC addresses of M1, M2, BOC1, BOC2 and BOC3 on the MM. Whitelist each device under the following configuration hierarchies:
    - M1, M2 whitelisted under **Managed Network>VPNC**
    - BOC2 whitelisted under **Managed Network>Branch2**
    - BOC3 whitelisted under **Managed Network>Branch3**
6. Configure interfaces and VLANs and VPNC VIP
7. Branch MC basic configuration - Configure interfaces, VLANs, DHCP pools for APs and users, IP VLAN pool of controller IPs for Branch MCs

8. Uplink Configuration of Branch MCs - Add uplinks, load balancing, and policy based routing in Branch MCs
9. Advertise routes of Branch MC to VPNCs
10. Routing Configuration of VPNCs - Static routes and OSPF configuration in the VPNCs
11. Create an AP group and SSID for Branch1. GUI: **Managed Network>Branch1>AP Groups**. GUI: **Managed Network>Branch1>Tasks>Create a new WLAN**
12. Whitelist the Branch1 APs on the MM. This includes mapping them to the appropriate AP group. GUI: **Managed Network>Branch1>Configuration>Access Points>Whitelist**
13. Create an AP group and SSID for Branch2. GUI: **Managed Network>Branch2>AP Groups.** GUI: **Managed Network>Branch2>Tasks>Create a new WLAN**
14. Whitelist the Branch2 APs on the MM. This includes mapping them to the appropriate AP group. GUI: **Managed Network>Branch2>Configuration>Access Points>Whitelist**

Activate

1. Set up provisioning rules in Activate to whitelist the branch controllers and redirect them to the MM
2. Optionally, if VPNCs are being used set up provisioning rules to whitelist them and redirect them to the MM as well

VPNC

1. Back up the existing configuration on the ArubaOS 6 masters. GUI: **Maintenance>Backup Flash**
2. Upgrade the image on master M1 to ArubaOS 8 and reboot it. GUI: **Maintenance>Image Management**
3. Provision M1 to be a VPNC managed by the MM via the CLI setup dialog. M1 will now become VPNC1
4. Repeat steps 2-3 to convert M2 into VPNC2

Branch 1

1. Back up the existing configuration on the ArubaOS 6 BOC1. GUI: **Maintenance>Backup Flash**
2. Upgrade the image on BOC1 to ArubaOS 8 and reboot it. GUI: **Maintenance>Image Management**
3. If provisioning rules have been created on Activate, the controller will perform ZTP, establish communication with the MM, and download its configuration

4. Optionally, the controller may be manually configured. Provision BOC1 to be a MC managed by the MM via the CLI setup dialog. BOC1 will now become MC1

Branch 2

1. Back up the existing configuration on the ArubaOS 6 BOC2. GUI: **Maintenance>Backup Flash**
2. Upgrade the image on BOC2 to ArubaOS 8 and reboot it. GUI: **Maintenance>Image Management**
3. If provisioning rules have been created on Activate, the controller will perform ZTP, establish communication with the MM, and download its configuration
4. Optionally, the controller may be manually configured. Provision BOC2 to be a MC managed by the MM via the CLI setup dialog. BOC2 will now become MC2

Follow similar steps for Branch 3.