

# Campus WLAN Redundancy

**Author:**  
**Makarios Moussa**

**Contributors:**  
**Justin Noonan**  
**Michael W. Wong**  
**David Balding**



Validated Reference Design

## **Copyright Information**

© Copyright 2017 Hewlett Packard Enterprise Development LP.

## **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

---

<b>Contents</b>	<b>3</b>
<b>Figures</b>	<b>7</b>
<b>Symbols</b>	<b>10</b>
<b>About this Guide</b>	<b>12</b>
Acronyms	12
Scope	14
Reference Material	15
Introduction to WLAN Campus Redundancy	16
<b>AP Boot Up and Master Discovery</b>	<b>18</b>
Controller Roles	18
Master Controller	18
Standalone Controller	18
Local Controller	18
Branch Controller	19
Controller Functions	19
Standby Master Controller	19
AP Master Controller	19
LMS Controller	19
Backup-LMS Controller	19
Campus AP Boot Process	20
Static Discovery	20
Dynamic Discovery Methods	21
DHCP Option 43	22
Aruba Discovery Protocol	23
DNS Query	23
LMS Controller	26
Control Plane Security	26
Boot Process with CPSec	27
CPSec Tunnels	28
CPSec Cluster Root for Multiple Masters	29

---

<b>Redundancy Overview</b>	<b>30</b>
Legacy Redundancy	30
VRRP	31
Master Redundancy	32
LMS/Backup-LMS	34
Legacy Redundancy Key Considerations	34
Redundancy with AP Fast Failover	35
High Availability Controller Roles	35
HA and Legacy Redundancy Comparison	35
Deployment Models	36
Active and Standby Tunnel Details	37
Flow	39
HA Enhancement Features	40
Inter-Controller Heartbeat	40
Introduction	40
Failover Scenario	41
Inter-Controller Heartbeat Flow	43
Client State Synchronization	44
Introduction	44
Key Considerations	44
Failover Scenario	45
Capacity Extension	47
Introduction	47
Standby AP Over-Subscription	48
Standby AP Over-Subscription Example	49
N+1 Over-Subscription	50
Master Redundancy	51
HA Support	51
HA Constraints	51
HA Failover	52
HA Support for Bridge Mode	53
HA Supported AP Forwarding Modes	53
<b>Centralized Licensing</b>	<b>54</b>
Supported Topologies	55
Additional Supported Topology in ArubaOS 6.5	56
License Server Redundancy Failover	57
Single License Server Failure - No Redundancy	58

---

Frequently Asked Questions .....	59
<b>Redundancy Migration .....</b>	<b>60</b>
HA AP Fast Failover Redundancy Guidelines .....	60
Migration from VRRP to HA AP Fast Failover .....	61
Migration from LMS and Backup LMS to HA AP Fast Failover .....	63
Migration for Master-Local Deployment .....	65
<b>Performance Comparison .....</b>	<b>66</b>
Legacy VRRP Failover .....	66
HA AP Fast Failover .....	66
HA AP Fast Failover vs. Legacy .....	67
<b>Wired Campus Redundancy .....</b>	<b>68</b>
Component Level Redundancy .....	68
Management Module Redundancy .....	68
Power Supplies Redundancy .....	69
Device Level Redundancy .....	71
Stacking Overview .....	71
Virtual Switching Framework .....	73
Backplane Stacking .....	77
Ring Topology .....	78
Mesh Topology .....	79
Advantages .....	81
Key Considerations .....	81
Link Level Redundancy .....	82
Overview .....	82
Virtual Router Redundancy Protocol .....	82
Link Aggregation Control Protocol .....	83
Distributed Trunking .....	85
Spanning Tree Protocol .....	87
Advantages .....	88
VRRP Advantages .....	88
LACP Advantages .....	88
STP Advantages .....	88
Key Considerations .....	88
VRRP Key Considerations .....	88
LACP Key Considerations .....	88

---

STP Key Considerations .....	88
Campus Network Redundancy - Implementation Guidance .....	89
<b>HA Deployment Models .....</b>	<b>91</b>
Master / Standby Master with HA Active-Active Locals .....	91
Introduction .....	91
Configuration Methodology .....	92
Failover Scenario .....	93
Benefits .....	94
Key Considerations .....	94
One Master - One Local .....	95
Master-Local (HA Active-Active) .....	96
Master-Local (HA Active-Standby) .....	97
Failover Scenario .....	98
Benefits .....	99
Key Considerations .....	99
Independent Masters / All Standalone Masters .....	100
Prerequisites .....	100
Configuration Methodology .....	101
Failover Scenario .....	102
Benefits .....	103
Key Considerations .....	103
Master Redundancy (Master / Standby Master) .....	104
Introduction .....	104
Configuration Methodology .....	105
Failover Scenario .....	106
Benefits .....	107
Key Considerations .....	107
N+1 (Over-Subscription) .....	108
Configuration Methodology .....	108
Failover Scenario .....	110
Benefits .....	111
Key Considerations .....	111

Figure 1 VRD Core Technologies .....	14
Figure 2 Campus Global Network Architecture .....	17
Figure 3 Campus AP Boot Process .....	20
Figure 4 Dynamic Controller Discovery Process .....	21
Figure 5 Master Discovery Packet Capture .....	22
Figure 6 Controller Discovery DHCP Option 43 .....	22
Figure 7 Controller Discovery DHCP Option 43 Steps .....	23
Figure 8 Controller Discovery DNS Entry .....	24
Figure 9 Controller Discovery DNS Entry Steps .....	25
Figure 10 LMS Controller .....	26
Figure 11 Campus AP Boot Process with CPSec .....	27
Figure 12 AP Tunnels Setup with CPSec .....	28
Figure 13 CPSec Cluster Root for Multiple Masters .....	29
Figure 14 VRRP between Controllers .....	31
Figure 15 Master Redundancy .....	32
Figure 16 Active Master .....	33
Figure 17 LMS/Backup-LMS .....	34
Figure 18 HA-AP Fast Failover Deployment Models .....	36
Figure 19 AP Failover to Pre-Established Standby Tunnels .....	37
Figure 20 Active Tunnel .....	37
Figure 21 Standby Tunnel Becomes Active .....	38
Figure 22 AP Fast Failover Flow .....	39
Figure 23 Inter-Controller Heartbeat Failover Scenario 1 .....	41
Figure 24 Inter-Controller Heartbeat Failover Scenario 2 .....	42
Figure 25 Inter-Controller Heartbeat Failover Scenario 3 .....	42
Figure 26 Inter-Controller Heartbeat Flow .....	43
Figure 27 Client State Synchronization Failover Scenario Step 1 .....	45
Figure 28 Client State Synchronization Failover Scenario Step 2 .....	45
Figure 29 Client State Synchronization Failover Scenario Step 3 .....	46
Figure 30 Client State Synchronization Failover Scenario Step 4 .....	46
Figure 31 Over-Subscription 1 .....	50
Figure 32 Over-Subscription 2 .....	50
Figure 33 Over-Subscription 3 .....	50
Figure 34 Master Redundancy .....	52
Figure 35 Master Failover .....	52
Figure 36 Supported Topologies .....	55

Figure 37 Additional Supported Topology .....	56
Figure 38 Master License Server .....	57
Figure 39 Standby License Server .....	57
Figure 40 Single License Server .....	58
Figure 41 Migration from VRRP to HA-APFF .....	62
Figure 42 Migration from LMS to HA .....	64
Figure 43 Migration Master-Local Deployment .....	65
Figure 44 Management Module Redundancy – 5400R .....	68
Figure 45 Management Modules (MM) – Aruba 5400R .....	69
Figure 46 Power Supply Redundancy – 5400R/3810 .....	69
Figure 47 Redundant Power Supplies .....	70
Figure 48 Virtual Switching Framework (VSF) .....	71
Figure 49 Backplane Stacking .....	72
Figure 50 VSF on 5400R and 2930F .....	73
Figure 51 VSF 1 .....	74
Figure 52 VSF 2 .....	75
Figure 53 VSF 3 .....	76
Figure 54 Backplane Stacking on 2920 and 3810M .....	77
Figure 55 Backplane Stacking Ring Topology 1 .....	78
Figure 56 Backplane Stacking Ring Topology 2 .....	78
Figure 57 Backplane Stacking Ring Topology 3 .....	79
Figure 58 Backplane Stacking Mesh Topology 1 .....	79
Figure 59 Backplane Stacking Mesh Topology 2 .....	80
Figure 60 Backplane Stacking Mesh Topology 3 .....	80
Figure 61 Link Level Redundancy VRRP Example .....	82
Figure 62 Link Level Redundancy LACP Example 1 .....	83
Figure 63 Link Level Redundancy LACP Example 2 .....	83
Figure 64 Link Level Redundancy LACP Example 3 .....	84
Figure 65 Link Level Redundancy DT Example 1 .....	85
Figure 66 Link Level Redundancy DT Example 2 and 3 .....	85
Figure 67 Link Level Redundancy DT Example 4 .....	86
Figure 68 Link Level Redundancy STP Example .....	87
Figure 69 Implementation Guidance 1 .....	89
Figure 70 Implementation Guidance 2 .....	89
Figure 71 Implementation Guidance 3 .....	90
Figure 72 Master / Standby Master with HA Active-Active Locals .....	92
Figure 73 Active-Active Locals with AP Fast Failover Operation .....	93
Figure 74 Master-Local Deployment Models .....	95
Figure 75 Master-Local (HA Active-Active) .....	96
Figure 76 Master-Local (HA Active-Standby) .....	97








---



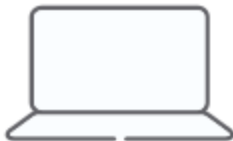

<i>Figure 77 Master-Local (HA Active-Standby) AP Fast Failover Operation</i> .....	98
<i>Figure 78 Independent Masters (HA Active-Active)</i> .....	101
<i>Figure 79 Independent Masters HA Fast Failover Operation</i> .....	102
<i>Figure 80 Master Redundancy (HA Master-Standby)</i> .....	105
<i>Figure 81 Master Redundancy HA Failover</i> .....	106
<i>Figure 82 HA with 2+1 Deployment</i> .....	108
<i>Figure 83 HA with 2+1 Deployment Failover</i> .....	110

The table below describes the symbols used in the figures in this guide.

**Table 1:** *Symbols*

Description	Symbol
Wireless Controller	
Access Point	
Layer 2 Switch	
Layer 3 Switch	
Router	

**Table 1:** *Symbols*

Description	Symbol
Servers/PBX	
Wired Client - Desktop Computer	
Wireless Client - Laptop	
Wireless Client - Smart Phone	

This chapter includes the following topics:

- [Acronyms on page 12](#)
- [Scope on page 14](#)
- [Reference Material on page 15](#)
- [Introduction to WLAN Campus Redundancy on page 16](#)

## Acronyms

[Table 2](#) describes the acronyms used in this guide.

**Table 2:** *Acronyms*

Acronym	Definition
ACR	Advanced Cryptography
AP	Access Point
BSSID	Basic Service Set Identifier
CAP	Campus Access Point
CPSec	Control Plane Security
DHCP	Dynamic Host Configuration Protocol
DNS	Domain Name System
DT	Distributed Trunking
EAP	Extensible Authentication Protocol
FTP	File Transfer Protocol
GRE	Generic Routing Encapsulation
HA	High Availability
ISC	Interswitch Connect
IP	Internet Protocol
IPsec	Internet Protocol Security
LACP	Link Aggregation Control Protocol
LAG	Link Aggregation Group
LAN	Local Area Network

**Table 2: Acronyms**

Acronym	Definition
LMS	Local Management Switch
MAC	Media Access Control
MSTP	Multiple Spanning Tree Protocol
OSPF	Open Shortest Path First
PAPI	Proprietary Access Protocol Interface
PBR	Policy Based Routing
PEF	Policy Enforcement Firewall
PEFNG	Policy Enforcement Firewall Next Generation
PEFV	Policy Enforcement Firewall Virtual Private Network
PMK	Pairwise Master Key
PMKID	Pairwise Master Key Identification
PMKSA	Pairwise Master Key Security Association
RAP	Remote Access Point
RF	Radio Frequency
RPVST	Rapid Per-VLAN Spanning Tree
RSTP	Rapid Spanning Tree Protocol
SSID	Service Set Identifier
STP	Spanning Tree Protocol
TFTP	Trivial File Transfer Protocol
VAP	Virtual AP
VLAN	Virtual Local Area Network
VR	Virtual Router
VRD	Validated Reference Design
VRRP	Virtual Router Redundancy Protocol
WebCC	Web Content Classification
WLAN	Wireless Local Area Network
WMS	WLAN Management System
xSec	Extreme Security
ZTP	Zero Touch Provisioning

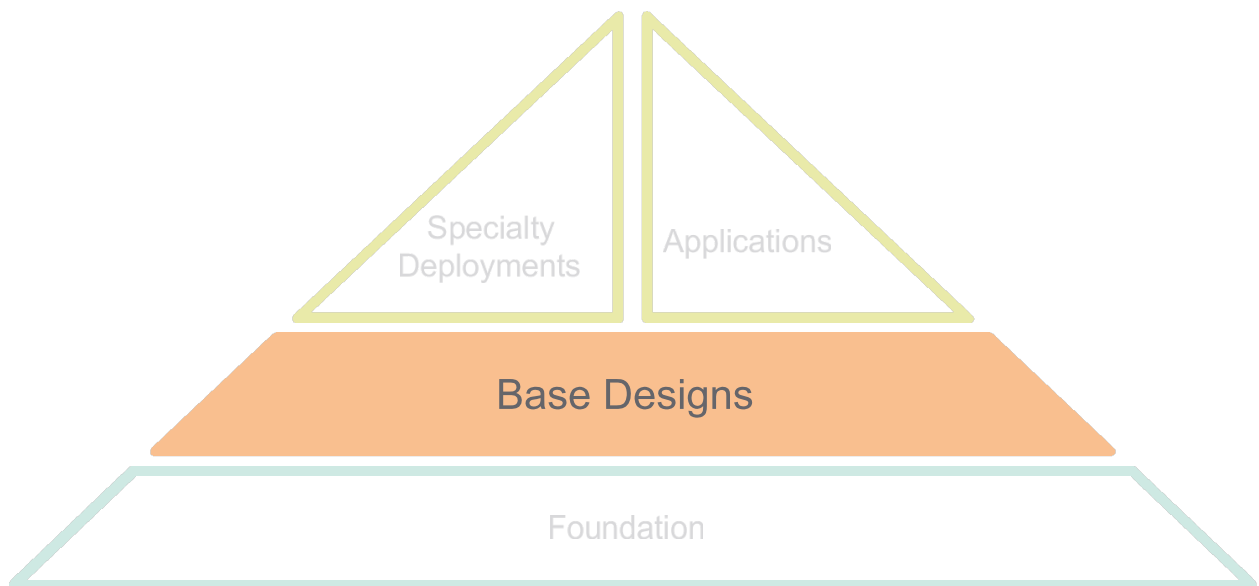
## Scope

The Aruba Validated Reference Design (VRD) is a series of technology deployment guides that include descriptions of Aruba technology, recommendations for product selection, network design decisions, configuration steps, and best practices. Together these guides comprise a reference model for understanding Aruba technology and design from common customer deployment scenarios.

The VRD series has four types of guides:

- **Foundation:** These guides explain the core technologies of an Aruba WLAN. The guides also describe different aspects of planning, operation, and troubleshooting deployments.
- **Base Designs:** These guides describe the most common deployment models, recommendations, and configurations.
- **Application:** These guides build on the base designs. These guides deliver specific information that is relevant to deploying particular applications such as voice, video or outdoor campus extension.
- **Specialty deployments:** These guides involve deployments in conditions that differ significantly from the common base design deployment models, such as high-density WLAN deployments.

**Figure 1** VRD Core Technologies



This Campus WLAN Redundancy design guide is part of “Base Designs” guides within the VRD core technology series.

- It is designed for Aruba Mobility Controllers running ArubaOS 6.4.3.4 and later.
- It does not cover the fundamental concepts of wireless networks. This guide assumes that the reader has a working knowledge of Aruba WLAN architecture.
- This design guide focuses on a large campus deployment model where multiple buildings with contiguous radio frequency (RF) are part of the campus.

## Reference Material

This is a base designs guide, and therefore it will not cover the fundamental wireless concepts. Readers should have a good understanding of wireless concepts and the Aruba technology explained in the foundation-level guides.

- For information on Aruba Mobility Controllers and deployment models, see the Aruba Mobility Controllers and Deployment Models Validated Reference Design, available on the Aruba website at <http://www.arubanetworks.com/vrd>
- The complete suite of Aruba technical documentation is available for download from the Aruba support site. These documents present complete, detailed feature and functionality explanations beyond the scope of the VRD series. The Aruba support site is located at: <https://support.arubanetworks.com/>
- For more training on Aruba products, or to learn about Aruba certifications, visit the Aruba training and certification page on our website. This page contains links to class descriptions, calendars, and test descriptions: <http://www.arubanetworks.com/support-services/training-services/>
- Aruba hosts a user forum site and user meetings called Airheads. The forum contains discussions of deployments, products, and troubleshooting tips. Airheads Online is an invaluable resource that allows network administrators to interact with each other and Aruba experts. Please visit: <http://community.arubanetworks.com/>

## Introduction to WLAN Campus Redundancy

Redundancy in the networking industry has been an inherent part of every sound system design, and wireless networking is no exception. The evolution of the wireless networking technology in reliability and speed went hand in hand with its fast adoption by consumers and businesses alike. This explosive growth mandated the need for wireless networks to stay up and available all the time.

This Validated Reference Design (VRD) document was written to provide Aruba WLAN networks designers with the necessary technical knowledge and assistance to properly set up fully redundant Wi-Fi solutions taking advantage of state of the art ArubaOS 6.5 High Availability features.

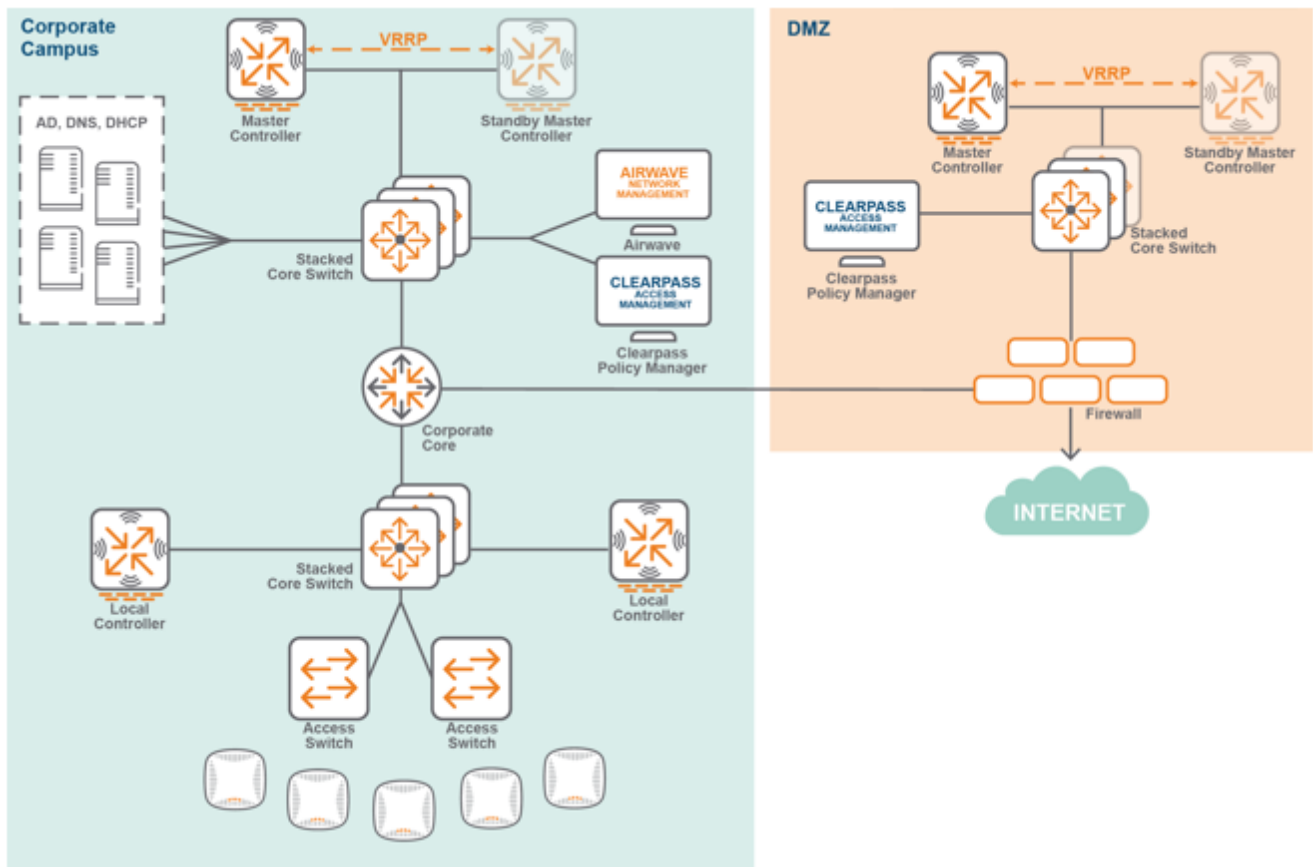
The document adds seven chapters to this introduction:

- Chapter two covers the components of the Aruba WLAN network and defines the Mobility Controllers roles and functions, the Campus Access Point (CAP) boot up process and its master discovery along with the discovery techniques, then concludes with the CPSec security scheme.
- Chapter three gives an overview of legacy redundancy techniques before launching into the more sophisticated High Availability (HA) AP Fast Failover feature. A brief description of the different deployment models and the enhanced features that provided both scalability with security in a sub-second failover time.
- Chapter four introduces Centralized Licensing to provide ease of use license planning and management with redundant license servers.
- Chapter five provides the necessary migration guidance to assist administrators with their migration from legacy to HA AP Fast Failover redundancy.
- Chapter six adds a failover performance comparison between legacy and HA using actual numbers provided by Aruba Test Engineering.
- Chapter seven provides the redundancy guidance of the indispensable Aruba wired switches that brings the WLAN network components together and connects Wi-Fi clients to the Data center and Internet.
- Chapter eight brings the information learned in the previous chapters together to give the guidelines of multiple HA deployments models.

A full redundancy of Master Mobility Controllers, Core and Distribution switches, and Local Mobility Controllers makes up the Aruba's recommended best practices deployment model. A campus global network architecture based on such redundancy is displayed below in [Figure 2](#).



Figure 2 Campus Global Network Architecture



This chapter includes the following topics:

- [Controller Roles on page 18](#)
- [Controller Functions on page 19](#)
- [Campus AP Boot Process on page 20](#)
- [LMS Controller on page 26](#)
- [Control Plane Security on page 26](#)

## Controller Roles

This section describes the different roles the Aruba Mobility Controllers can be assigned in the Campus WLAN Redundancy model. Controller roles are first assigned through Initial Setup.

Controller roles include the following:

- [Master Controller on page 18](#)
- [Standalone Controller on page 18](#)
- [Local Controller on page 18](#)
- [Branch Controller on page 19](#)

### Master Controller

The master controller is the anchor controller of the Wireless Local Area Network (WLAN) domain; it is responsible for all global configuration. This includes the access point (AP) groups, all the WLAN profiles, and the firewall roles and policies.

Several centralized features include: WLAN configuration, management and monitoring, whitelist database, and a licensing server. The master controller exchanges management traffic and pushes the global configuration to all other controllers via persistent internet protocol security (IPsec) tunnels. The master controller hosts several databases including the control plane security (CPSec) and remote access point (RAP) whitelist, and the WLAN management system (WMS) database storing RF monitoring information. Furthermore, it is capable of terminating APs and passing Wi-Fi user traffic.

### Standalone Controller

The standalone controller is a role selected through the controller initial-setup. It fulfills the same functions as the master controller role. It remains as an option in the Initial Setup to retain backward compatibility with older controller hardware platforms.

### Local Controller

You can deploy controllers in a local role, which initially takes local configuration (such as ports, VLANs, IP addresses, name, and master controller IP), and then receives the global configuration from the master through a persistent IPsec tunnel.

The local controller is the preferred controller to act as the APs tunnel termination point. All Wi-Fi user traffic flowing through the generic routing encapsulation (GRE) tunnels is decrypted and encrypted, firewalled, and switched and routed in and out of the local controller after a packet header conversion: 802.11 <-> 802.3.

## Branch Controller

The branch controller is a role reserved only for the 70xx hardware platforms and designed to be deployed in small branch offices (with a maximum of 64 APs for the 7030) with minimal IT configuration through the zero touch provisioning (ZTP) feature. Additional features include WAN Compression and Bandwidth Contracts, Policy Based Routing (PBR), and WAN Health Check.

## Controller Functions

This section describes the controller functions, including the following:

- [Standby Master Controller on page 19](#)
- [AP Master Controller on page 19](#)
- [LMS Controller on page 19](#)
- [Backup-LMS Controller on page 19](#)

### Standby Master Controller

The standby master controller runs the master-redundancy feature. It is a backup master that remains in sync with the active master as far as the global configuration and the various databases. Virtual Router Redundancy Protocol (VRRP) is the redundancy protocol used, while a persistent IPsec tunnel is used for data synchronization.

The standby master remains passive (it has no active AP tunnels or users), until a failure of the active master takes place. Then the standby master steps in to take over as an active master. Databases such as WMS, AP whitelist, and Local-userdb sync with the active master. Note that the backup master can still run services like Dynamic Host Configuration Protocol (DHCP), Open Shortest Path First (OSPF), and Spanning Tree Protocol (STP).

### AP Master Controller

Before discussing the ap boot up process and master discovery, it is important to understand the meaning of 'AP master' and not confuse it with the master controller role. An AP master could be any Aruba Mobility Controller whether it is a master or a local. This is a controller the AP needs to discover its IP and establish a Proprietary Access Protocol Interface (PAPI) (UDP 8211) communication with, in order to check and update its firmware and receive its configuration on a per ap-group basis.

### LMS Controller

The Local Management Switch (LMS) controller terminates one or multiple groups of APs and handles their users' traffic. The final tunnel termination point is determined by the LMS IP option that is part of the configuration downloaded by the AP. Therefore, the LMS IP option plays a major role in distributing the AP tunnels and users load among multiple controllers in the WLAN domain.

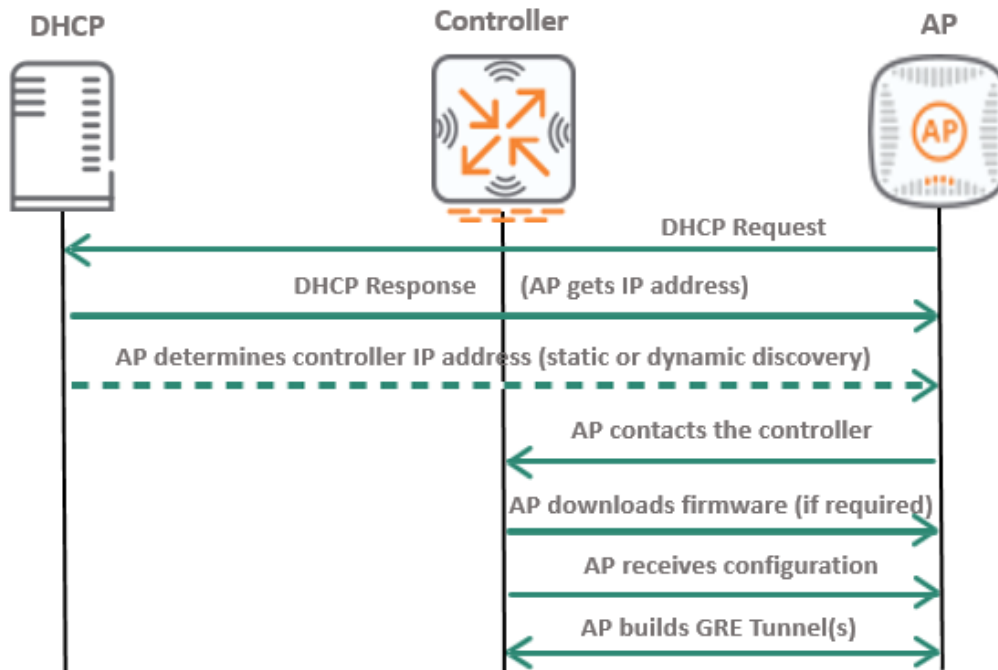
### Backup-LMS Controller

The Backup-LMS IP is another option in the AP system profile used for redundancy purposes, in case the LMS controller is no longer reachable.

## Campus AP Boot Process

During the campus AP boot process, the AP must discover and connect to a provisioning mobility controller. [Figure 3](#) illustrates the steps in the process.

**Figure 3** Campus AP Boot Process



The process includes the following steps:

1. The AP sends a DHCP Request.
2. The AP receives an IP address in the DHCP Response.
3. The AP discovers the AP master controller IP address; this can be either a static or dynamic process. See [Static Discovery on page 20](#) and [Dynamic Discovery Methods on page 21](#) below.
4. The AP contacts the discovered controller.
5. If required, the AP downloads its firmware from the discovered controller in case of a version mismatch.
6. The AP receives the configuration from the controller.
7. The AP creates the GRE tunnel(s) through which the user traffic passes.

### Static Discovery

In the static method, the master controller IP address is provisioned and saved in AP Flash. You can set other parameters like the AP IP address, subnet mask, default gateway, and so on. However, such static provisioning may not scale when deploying hundreds of APs. A more scalable approach is to use one of the other dynamic methods.

## Dynamic Discovery Methods

When an AP first boots up, it acquires an IP address that is typically given by a DHCP server, along with other IP stack information (subnet mask, default gateway, domain name system (DNS) IP, and domain name). The next step is to locate a controller (AP master).

If the AP is running a different firmware than the controller it discovered, it will use File Transfer Protocol (FTP) /Trivial File Transfer Protocol (TFTP) to download the matching firmware version. Once the AP reboots with the new firmware, it will repeat the same process starting from step 1 (that is, getting an IP address). The controller uses PAPI (an Aruba proprietary communication protocol) to send the AP its configuration. This configuration provides the AP with all RF and WLAN settings needed to operate. The AP will create a GRE (protocol 47) tunnel to pass the user traffic to the controller.

In summary, the AP <-> controller communication will consist of a control channel (PAPI) and one or more GRE tunnels for Wi-Fi user data traffic.

There are three types of dynamic discovery methods:

- [DHCP Option 43 on page 22](#)
- [Aruba Discovery Protocol on page 23](#)
- [DNS Query on page 23](#)

[Figure 4](#) illustrates the different dynamic discovery methods in the discovery process.

**Figure 4** Dynamic Controller Discovery Process

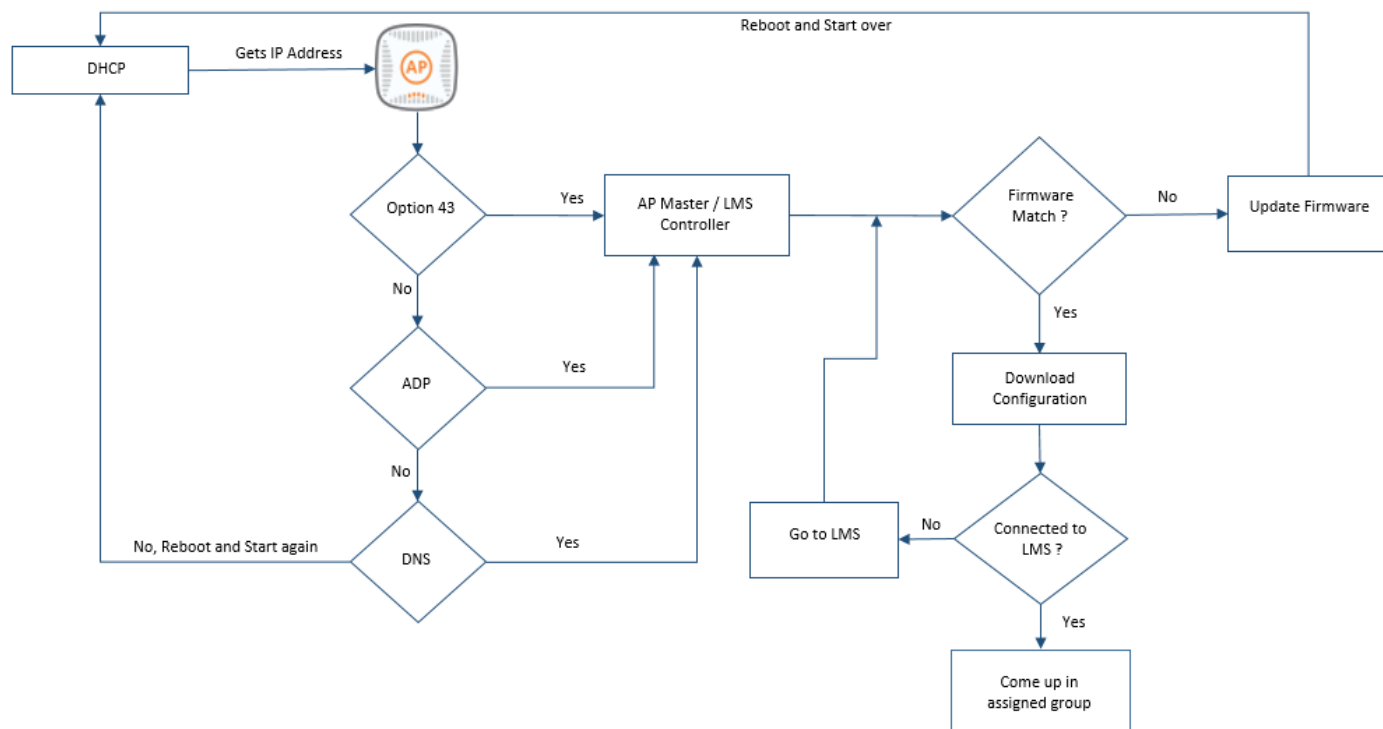
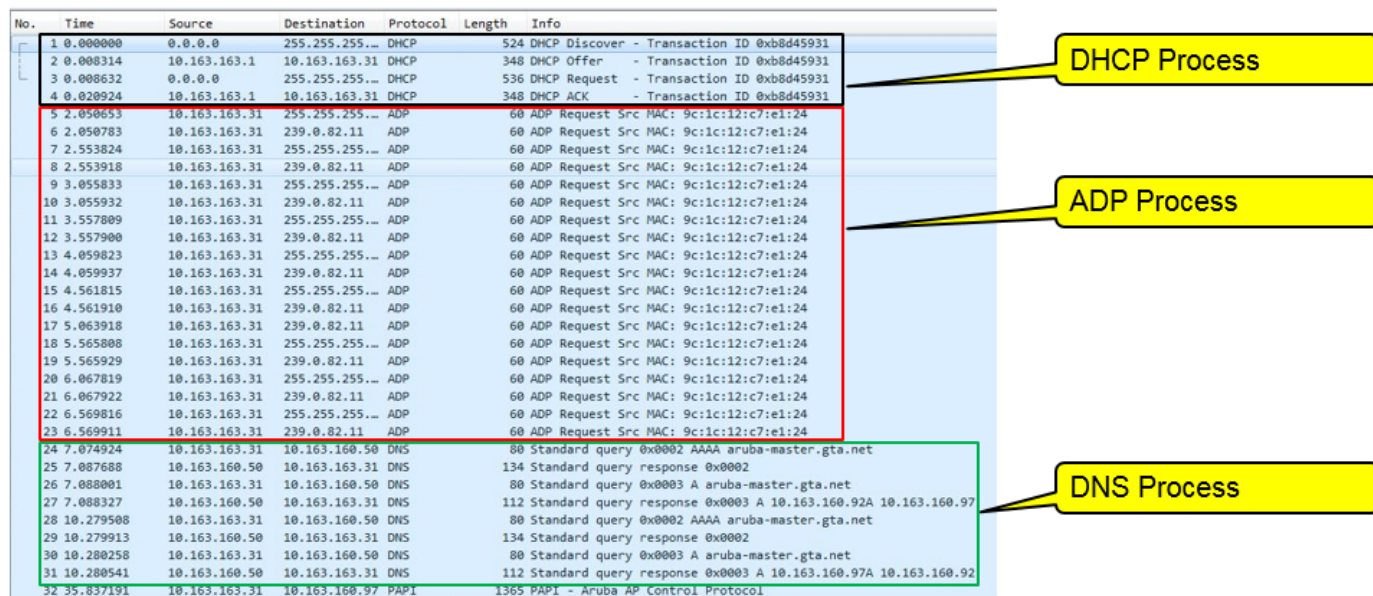


Figure 5 illustrates a packet capture showing the AP master discovery process.

**Figure 5** Master Discovery Packet Capture



## DHCP Option 43

Figure 6 illustrates the campus scenario with centralized DHCP option 43 for discovering the controller's IP address. Both Option 60 and Option 43 are required. Configure option 60 with the string ArubaAP, and option 43 with the controller's IP address as a string. For master redundancy, set the option 43 value to the VRRP IP address. That way, if one master controller is down, the AP can always discover an available master controller from which to download its configuration.

**Figure 6** Controller Discovery DHCP Option 43

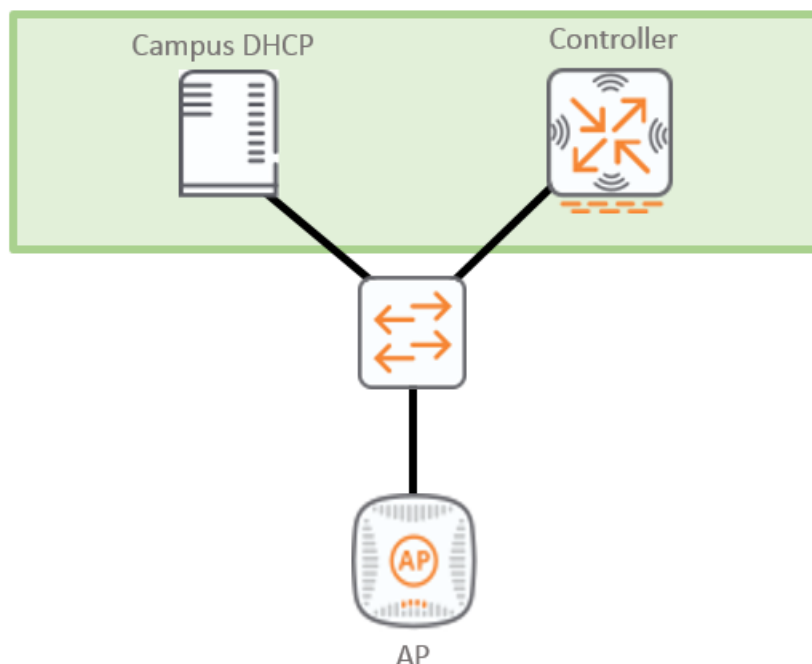
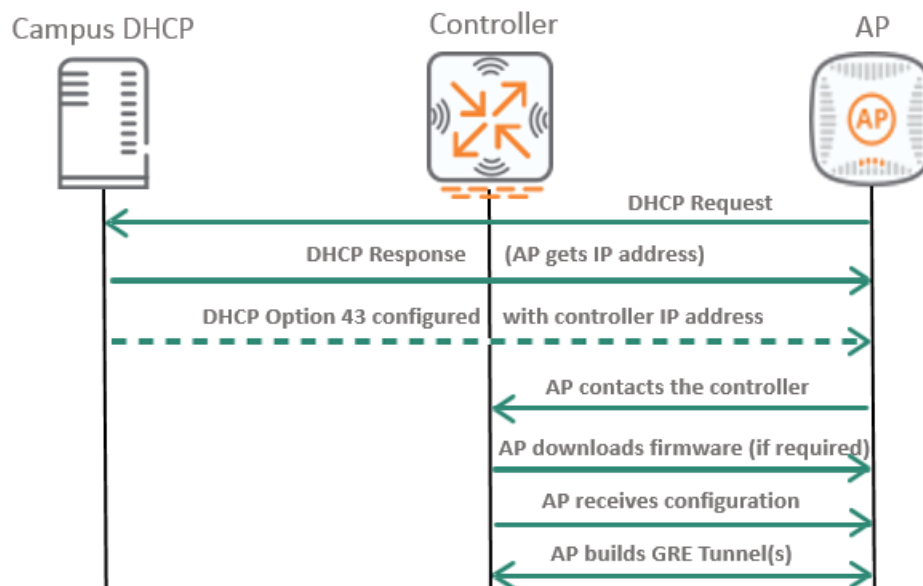


Figure 7 illustrates the steps in the process.

**Figure 7** Controller Discovery DHCP Option 43 Steps



The process includes the following steps:

1. The AP sends a DHCP Request.
2. The AP receives an IP address in the DHCP Response.
3. The AP discovers the AP master controller IP address using the dynamic process with the DHCP option 43 configured.
4. The AP contacts the discovered controller.
5. If required, the AP downloads its firmware from the discovered controller in case of a version mismatch.
6. The AP receives the configuration from the controller.
7. The AP creates the GRE tunnel(s) through which the user traffic passes.

### Aruba Discovery Protocol

ADP (Aruba Discovery Protocol) is enabled by default on all Aruba APs and controllers. If DHCP Option 43 or DNS are not configured, the AP sends out periodic L2/L3 broadcast and multicast (239.0.82.11) queries at 0.5 second intervals to locate a controller.

If the AP and controller are on the same L2 network, the broadcast packet will elicit a response from the controller with its IP address, while the multicast packet does require an upstream multicast routing in an L3 environment to reach the controller.

If multiple controllers are available in the same Layer 2 network, the AP updates its firmware as needed and downloads its configuration from the first controller that responds to the ADP message, to come up in its assigned ap-group or in the "default" group if not yet provisioned.

### DNS Query

If one of the previous methods has not discovered a controller, then the AP sends a DNS query for '**aruba-master.domain.com**', where 'domain.com' is the domain received through DHCP. The 'DNS server' is supplied by DHCP. If the AP receives a DNS response with more than one IP address, it caches the additional addresses. The AP uses the first IP received to contact the controller. If no response is received, the AP continues through the list of cached IP addresses until it establishes a connection with the controller.

APs are factory configured to use the host name “aruba-master” for the master controller discovery. Therefore an A or CNAME record needs to be added to the internal DNS server for the name “aruba-master”.

Aruba recommends the DNS Query method for AP Master discovery because it involves minimal changes to the network and provides the greatest flexibility in the placement of APs.

[Figure 8](#) illustrates the campus scenario with the DNS query for discovering the controller’s IP address.

**Figure 8** *Controller Discovery DNS Entry*


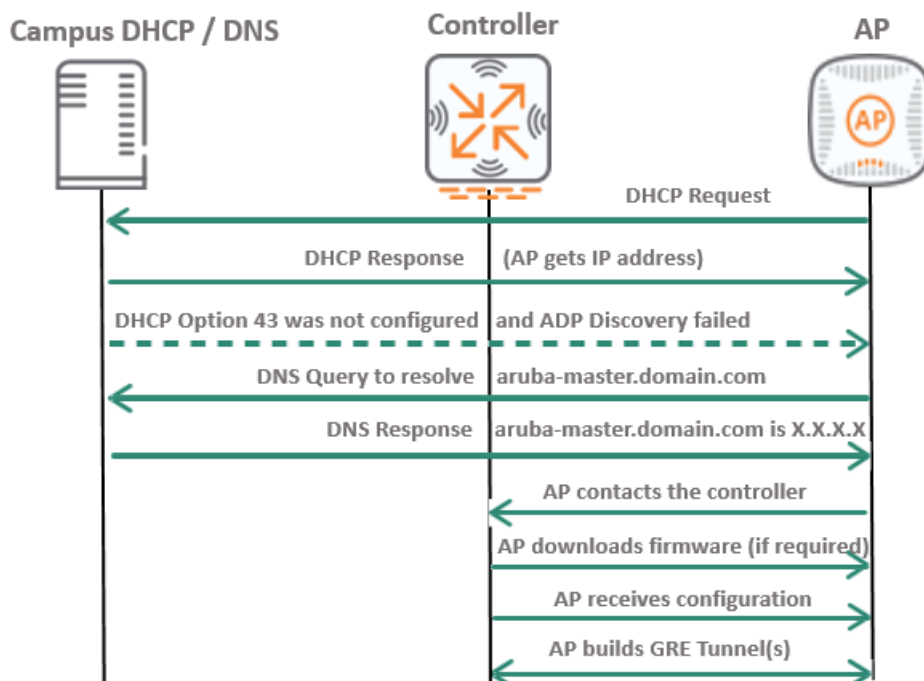




Figure 9 illustrates the steps in the process.

**Figure 9** *Controller Discovery DNS Entry Steps*



The process includes the following steps:

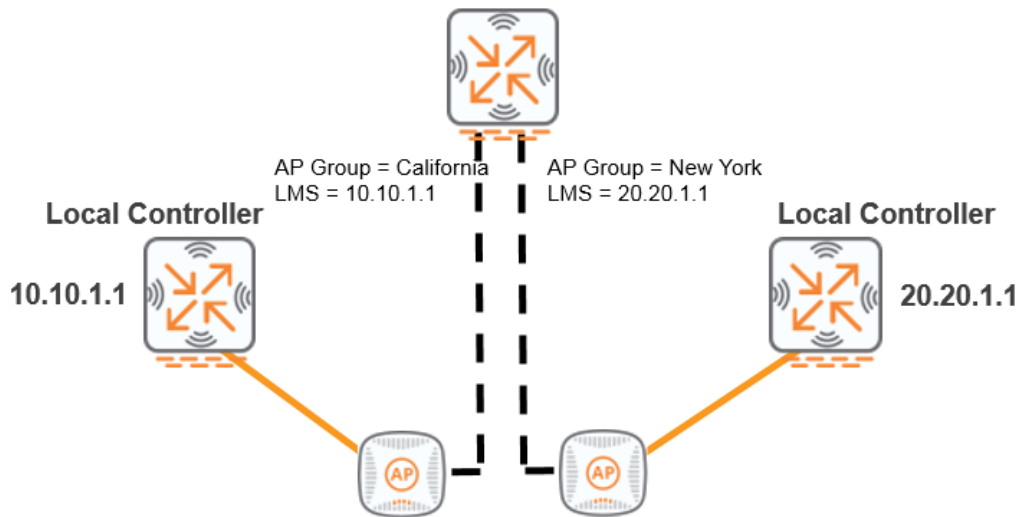
1. The AP sends a DHCP Request.
2. The AP receives an IP address in the DHCP response.
3. The DHCP Option 43 was not configured and ADP discovery has failed.
4. The AP sends a DNS query to resolve **aruba-master.domain.com**.
5. The DNS response includes the master controller IP address.
6. The AP contacts the controller.
7. If required, the AP downloads firmware from the discovered AP master in case of a version mismatch.
8. The AP receives the configuration from the controller.
9. The AP creates the GRE tunnel(s) through which the user traffic passes.

## LMS Controller

An AP configured to be in the “California” AP group (see [Figure 10](#)) receives an LMS IP address as 10.10.1.1 and builds its tunnels to terminate on that LMS controller.

Similarly, if an AP is provisioned to be in the “New York” AP group (see [Figure 10](#)), it receives an LMS IP address as 20.20.1.1 and builds its tunnels to terminate on that LMS controller.

**Figure 10** LMS Controller



## Control Plane Security

The Control Plane Security (CPSec) feature has two main goals:

1. Secure the control channel (PAPI) between Aruba Mobility Controllers and their attached APs.
2. Allow only authorized APs to connect to controllers and join the Aruba WLAN network.

The above goals are achieved in the following manner:

- The control traffic (PAPI) is secured by a certificate based 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 Control Plane Security (CPSec) is enabled by default, upon boot up, the master controller certifies its local controllers using its generated factory certificate. In turn, local controllers certify their APs (sign their factory default certificates).

Once the APs are authorized through the CPsec whitelist (in 'certified-factory-cert' state), they initiate secure PAPI (UDP 8209 inside ipsec) communication with the controller, sync their firmware, and download their configuration.



Only PAPI communication is encrypted, while GRE tunnels remain in the clear (their Wi-Fi data would be encrypted in tunnel mode if the service set identifier (SSID) opmode is not 'Open').

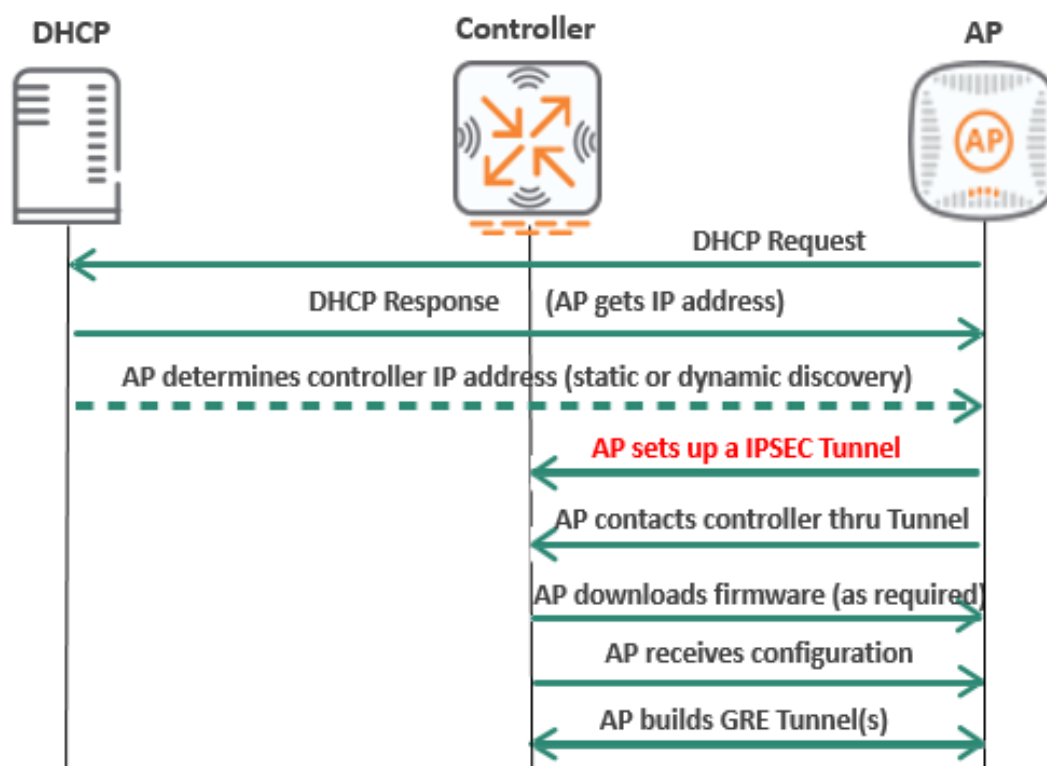
This section includes the following topics:

- [Boot Process with CPsec on page 27](#)
- [CPsec Tunnels on page 28](#)
- [CPsec Cluster Root for Multiple Masters on page 29](#)

## Boot Process with CPsec

[Figure 11](#) illustrates the steps in the campus AP boot process with CPsec.

**Figure 11** Campus AP Boot Process with CPsec



The process includes the following steps:

1. The AP sends a DHCP Request.
2. The AP receives an IP address in the DHCP Response.
3. The AP determines the master controller IP address; this can be either a static or dynamic process. See [Static Discovery on page 20](#) and [Dynamic Discovery Methods on page 21](#).
4. The AP establishes an IPsec tunnel with the controller.
5. The AP exchanges PAPI (UDP 8209) over the IPsec tunnel with the controller.
6. If required, the AP downloads firmware from the discovered AP master in case of a version mismatch.
7. The AP receives the configuration from the controller
8. The AP creates the GRE tunnel through which the user traffic flows.

## CPSec Tunnels

[Figure 12](#) illustrates the AP tunnels setup with CPSec.

**Figure 12** *AP Tunnels Setup with CPSec*



The tunnels setup process includes the following steps:

1. The AP sets up the IPsec tunnel with the controller.
2. The PAPI Control Traffic flows inside IPsec.
3. The GRE tunnels are set up in the clear.

## CPSec Cluster Root for Multiple Masters


Keeping the CPsec certificate trust hierarchy in mind, it is important in multi-master deployments to group independent master controllers in a CPsec cluster with one master acting as the CPsec cluster root as illustrated in [Figure 13](#). The advantages are the following:

- Independent masters' CPsec whitelists remain in synchronization.
- The CPsec cluster root acts as the certificate trust anchor for all controllers and APs.
- All other masters are certified by that CPsec cluster root self-signed certificate.
- APs could move from one master controller domain to another without re-certification.



We highly recommend that the CPsec cluster root have a backup master controller to safeguard the trust anchor certificate and keys in case of failure of the trust anchor.

**Figure 13** CPsec Cluster Root for Multiple Masters



This chapter includes the following topics:

- [Legacy Redundancy on page 30](#)
- [Redundancy with AP Fast Failover on page 35](#)
- [HA Support for Bridge Mode on page 53](#)
- [HA Supported AP Forwarding Modes on page 53](#)

## Legacy Redundancy

This section includes the following topics:

- [VRRP on page 31](#)
- [LMS/Backup-LMS on page 34](#)

Two mechanisms are used to achieve legacy redundancy (prior to ArubaOS 6.3):

## VRRP

VRRP is a layer 2 protocol that provides redundancy at the controller level as illustrated in [Figure 14](#).

VRRP eliminates controller single point of failure when an election process between two controllers yields a VRRP Active controller that owns the virtual IP address (VIP) of the VRRP instance. A VRRP Backup controller steps in to take ownership of the VIP when the VRRP Active controller becomes unavailable.

Two commonly used VRRP scenarios are considered:

1. VRRP between two Master controllers:

Redundancy of the Master controllers is achieved through the master-master redundancy feature based on a single VRRP instance between the 2 master controllers.

2. VRRP between two Local controllers:

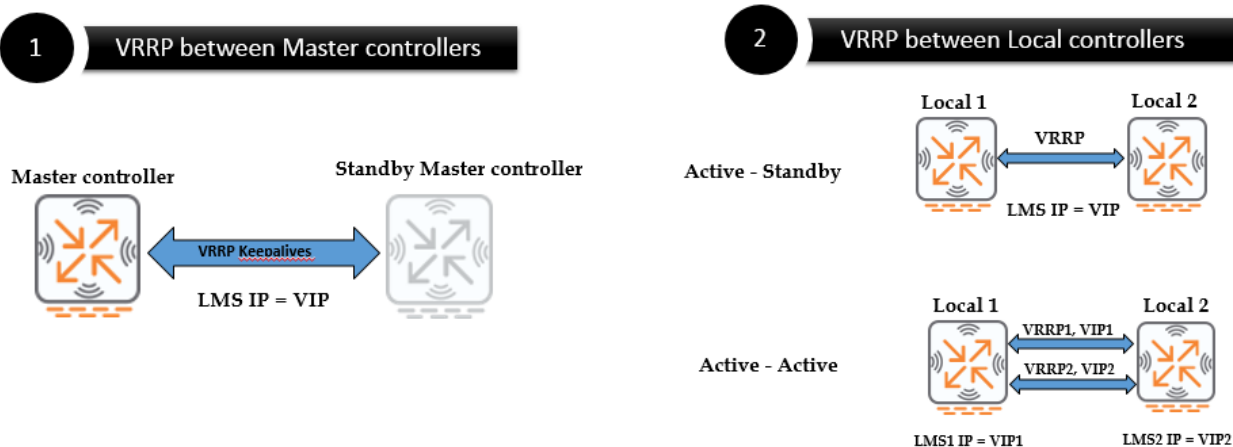
Two scenarios for VRRP redundancy between two local controllers are reviewed: Active-Standby and Active-Active.

a. Active-Standby runs a single instance of VRRP and the LMS IP = VIP. As a result, all APs terminate on the active local that handles all users traffic, while the standby local is idle.

b. Active-Active runs two instances of VRRP: Local1 would be active for instance VRRP1, while Local2 would be Active for instance VRRP2.

The AP load is divided between the 2 locals on a per ap-group basis where VIP1 is used as LMS IP for the first group, and VIP2 is the LMS IP for the second group.

**Figure 14** VRRP between Controllers

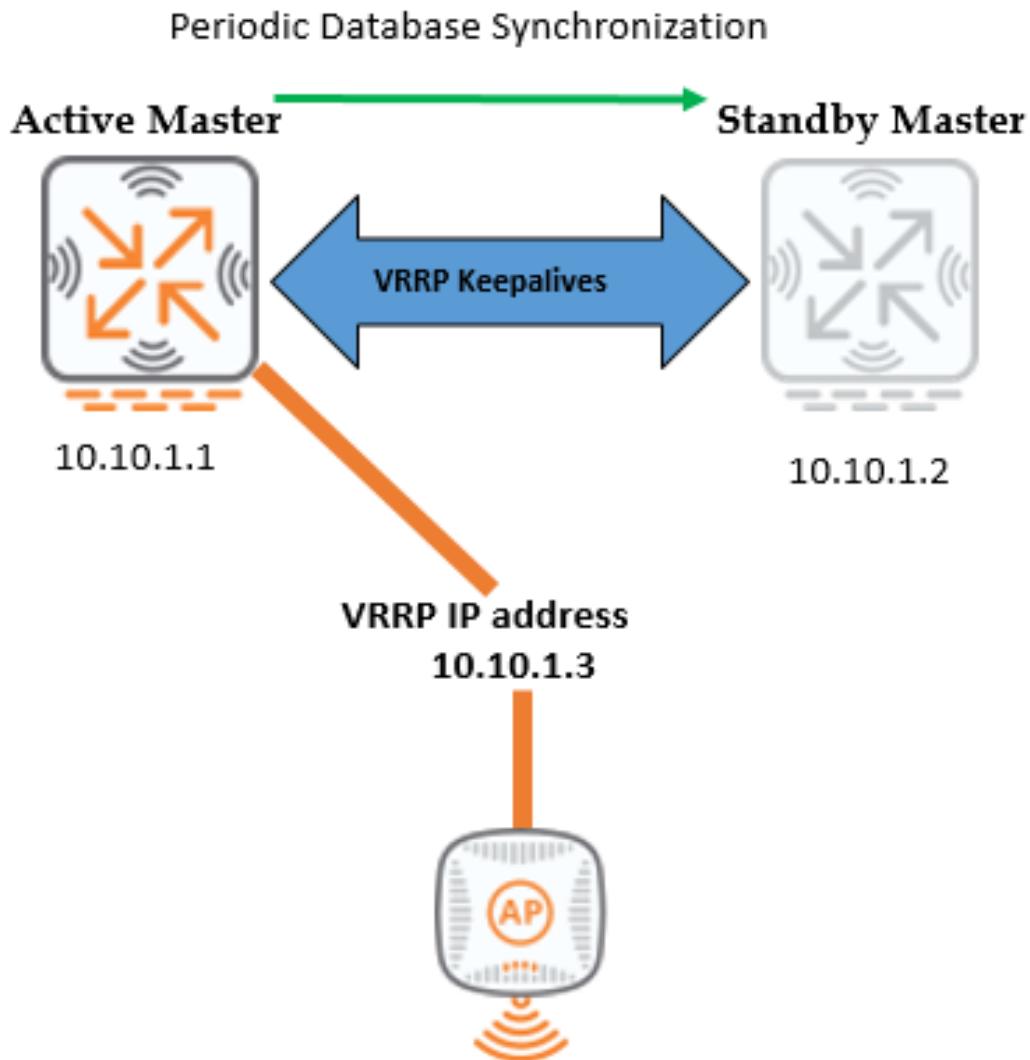


## Master Redundancy

As illustrated in [Figure 15](#), the Master redundancy is deployed in the following manner:

- VRRP between the two master controllers.
- Master-redundancy enabled.
- Periodic database synchronization enabled.
- LMS-IP configured to be the VRRP VIP for AP termination.
- If the master controller fails, then the standby becomes the active master.

**Figure 15** *Master Redundancy*

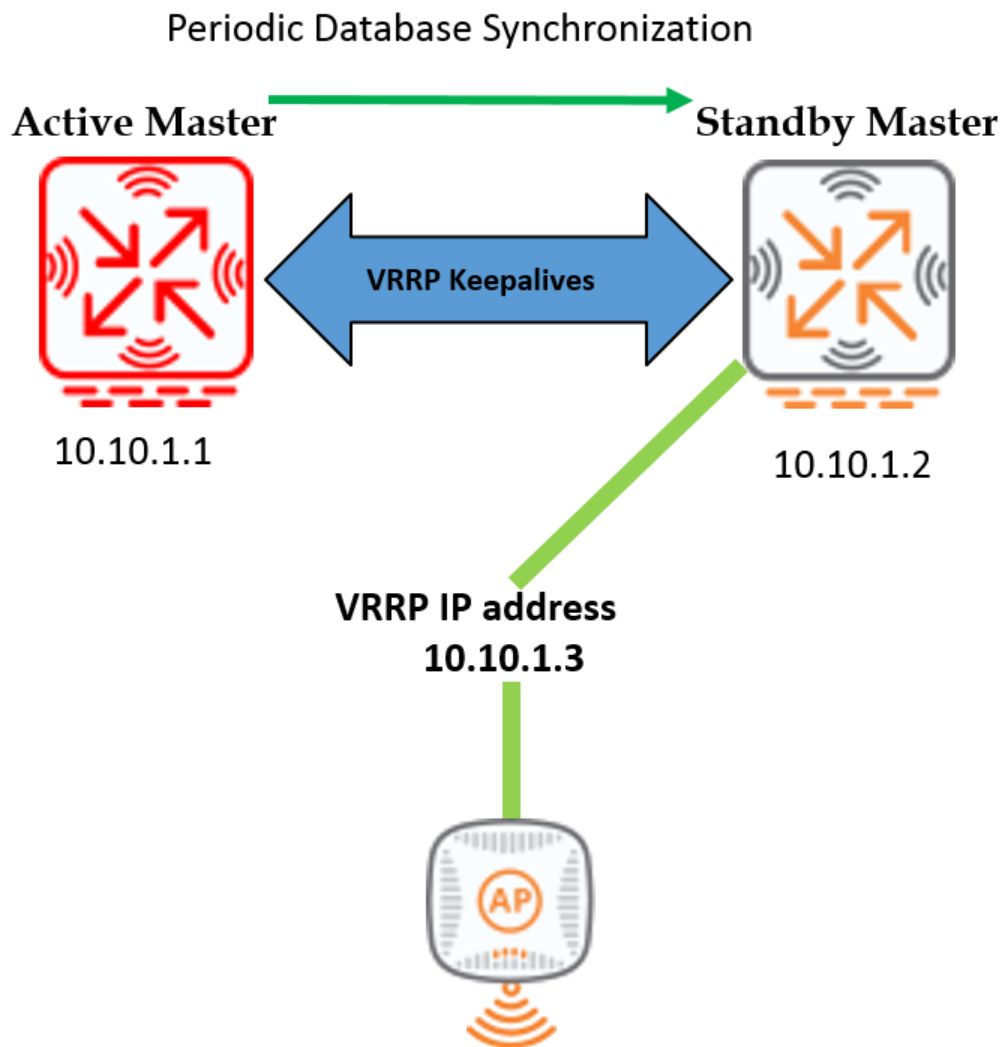




A Master failover illustrated in [Figure 16](#) goes through the following sequence of events:

1. The Active Master 10.10.1.1 fails.
2. The Standby Master detects the failure after three (3) consecutive VRRP keep-alive timeouts.
3. The Standby Master takes ownership of the VIP and becomes the new Active Master.
4. The AP bootstraps and rebuilds its tunnels with the new Active Master.

**Figure 16** *Active Master*



## LMS/Backup-LMS


LMS/Backup-LMS is primarily a mechanism at the AP level that relies on heartbeat timeouts with the LMS controller to fail over to a pre-configured Backup-LMS controller. Since APs are typically one or more hops away from their controllers, the mechanism is inherently a layer 3 mechanism.

When controllers are in separate L3 networks, VRRP cannot be used for redundancy. In such case, the LMS/Backup-LMS should be used for redundancy as illustrated in [Figure 17](#).

In this scenario:

- The AP terminates on the LMS controller.
- If the LMS controller fails.
- Eight missed heartbeats trigger an AP failover.
- The AP comes up on the Backup-LMS.

**Figure 17** LMS/Backup-LMS



## Legacy Redundancy Key Considerations

- Controller unavailability detected after 3 seconds (VRRP) and 8 heartbeat misses - 8 seconds (LMS/Backup-LMS).
- APs re-bootstrap when failing over that introduces additional delays while virtual APs (VAPs) are initialized and radios are reset.
- SSIDs are not available during AP failover.
- The setup of tunnels for a large number of APs introduces slower failover performance that affects the scalability of legacy redundancy solutions.

## Redundancy with AP Fast Failover

The High Availability (HA) - AP Fast failover was introduced in 6.3 onward to significantly enhance and reduce AP failover time to a standby controller.

The enhanced AP failover is achieved through the following mechanisms:

- Pre-established CPsec and standby GRE tunnels to a designated HA Standby controller.
- The AP does not bootstrap upon failover.
- The AP does not turn off its SSIDs and radios upon failover.
- Sub-second controller failover detection in 6.4 onward thanks to inter-controller heartbeat feature.

The AP Fast Failover feature can operate in L2 and L3 networks within the same campus.

This section includes the following topics:

- [High Availability Controller Roles on page 35](#)
- [HA and Legacy Redundancy Comparison on page 35](#)
- [Deployment Models on page 36](#)
- [Flow on page 39](#)
- [HA Enhancement Features on page 40](#)

### High Availability Controller Roles

You can configure a controller with one of the following three HA roles:

- **Active** – The LMS controller that terminates active AP tunnels and hosts user sessions. In such a role, it cannot act as an HA standby controller.
- **Standby** – The controller acts as a failover Standby controller only, and cannot be an LMS controller for APs to establish their initial connections and download their configuration.
- **Dual** – A dual controller can support both roles (acting as active controller for one set of APs, and a standby controller for another set of APs). This is the recommended role except for N+1 deployments.

We recommend configuring HA controllers as Dual except for N+1 deployments, where you would configure your standby controller as "Standby".

### HA and Legacy Redundancy Comparison

The following table compares HA - AP fast failover with the legacy redundancy (VRRP, LMS/Backup-LMS).

**Table 3:** *HA and Legacy Redundancy Comparison*

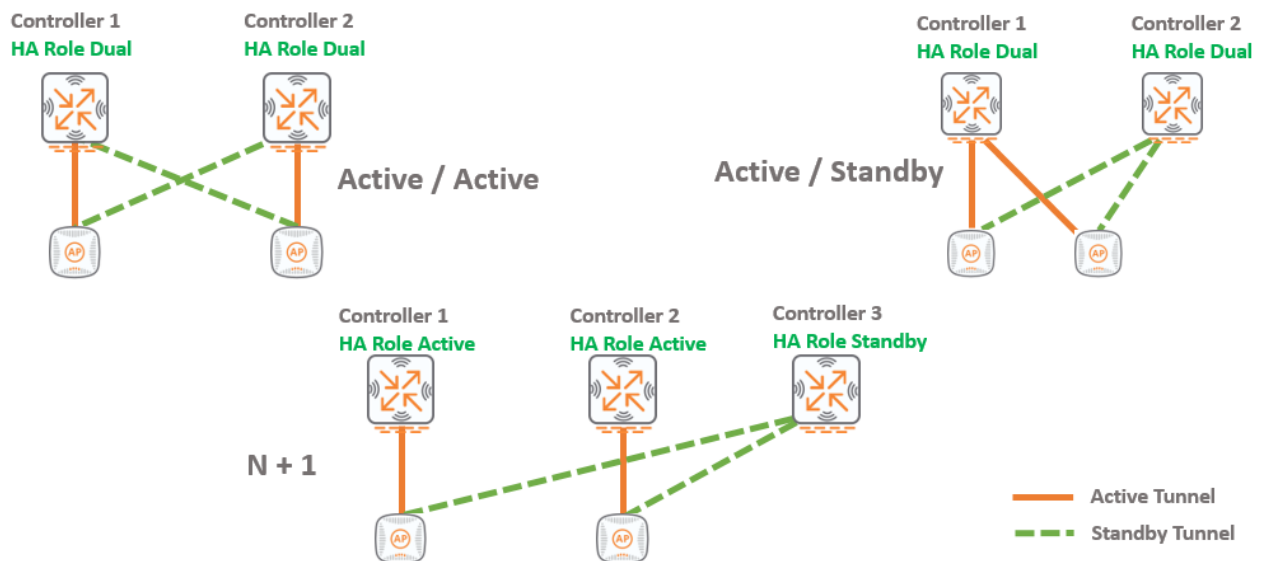
Feature	Legacy Redundancy	HA-AP Fast Failover
Failover	Slow failover to standby controller due to AP config re-download (re-bootstrap) and tunnel build up.	Much faster failover with use of standby tunnel and no configuration to re-download.
Messages	Approximately 10 messages exchanged between AP and controller.	Only one message exchanged during failover.
Tunnels	Additional IKE/IPsec overhead in case of CPsec.	Standby IPsec tunnel already established.

## Deployment Models

In HA-AP fast failover, you can deploy controllers in one of the following three models as illustrated in [Figure 18](#).

- **Active / Active** – Both controllers are deployed in HA Dual role and are actively terminating APs and processing user traffic. In this model, they are both acting as LMS controllers (HA Active) for 40% of their maximum AP capacity and backup (HA Standby) for each other. If one of the controllers becomes unreachable, the other controller will carry the total AP load.
- **Active / Standby** – Although both controllers are deployed in HA Dual role, only one controller acts as the LMS that terminates all APs.  
The other controller acts as HA Standby that terminates all standby tunnels from all deployed APs. If the Active controller becomes unreachable, all APs fail over to the Standby controller. See [Active and Standby Tunnel Details on page 37](#) for more information.
- **N+1** – In this model, multiple controllers, deployed in an HA Active role, are sharing the termination of all APs, while a single controller deployed in HA Standby role, acts as the backup for all the active controllers. This model requires that the AP capacity of the standby controller is able to support the total number of APs distributed across all active controllers in the HA group.


**Figure 18** HA-AP Fast Failover Deployment Models



## Active and Standby Tunnel Details


The Active / Standby model is illustrated in [Figure 19](#).

**Figure 19** AP Failover to Pre-Established Standby Tunnels




The process includes the following steps:

**Figure 20** Active Tunnel



1. The active controller goes down.
2. The AP detects failure after eight missed heartbeats.
3. The AP tears down its GRE tunnel.

**Figure 21** *Standby Tunnel Becomes Active*

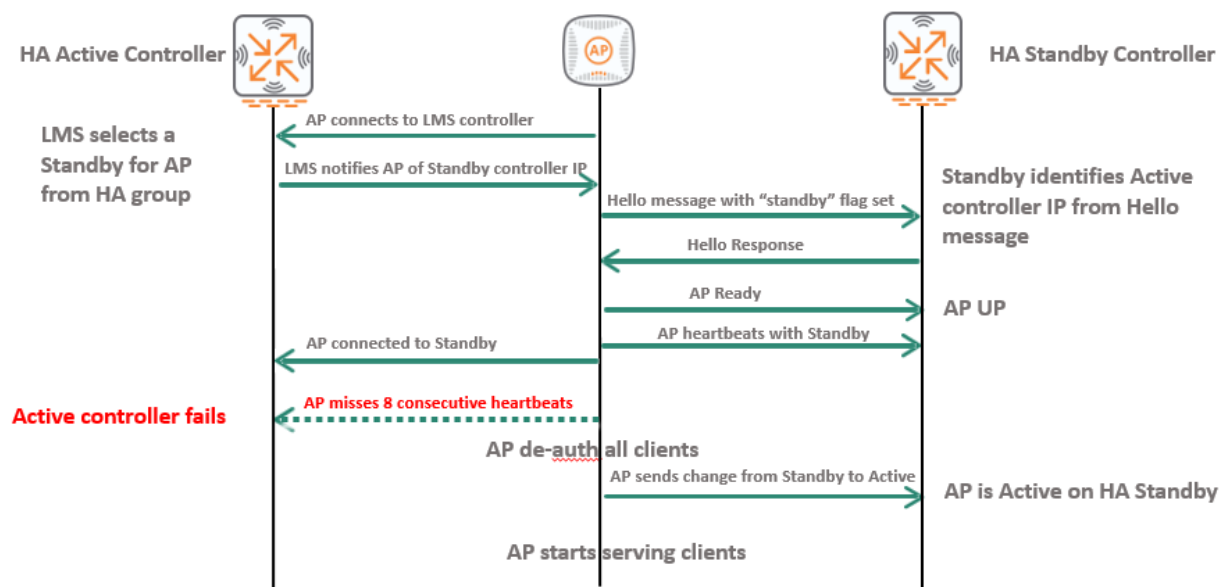


4. The GRE standby tunnel becomes active for user traffic.

## Flow

Figure 22 illustrates the HA-AP fast failover flow in the network.

**Figure 22** AP Fast Failover Flow



1. A provisioned AP ends up connected to its LMS controller.
2. From an HA perspective, that LMS controller is considered the HA Active controller for that AP.
3. Other controllers in either Dual or Standby roles within the same HA group become potential HA Standby controllers for that AP.
4. Once the AP connects with its HA Active controller, that controller notifies the AP of its assigned HA Standby controller IP.
5. The AP contacts its assigned Standby controller using a 'standby' flag.
6. Once the handshake is complete, the HA Standby controller acknowledges the AP connection as a Standby connection.
7. The AP informs its HA Active controller of its successful connection with its assigned HA Standby controller.
8. The AP exchanges heartbeats with the HA Standby controller in the same manner it does with the HA Active controller.
9. If CPSec was enabled, the AP would have also established an IPsec tunnel with the HA Standby controller to secure the control channel communication (PAPI).

An AP standby connection does not consume an additional AP license. However, such connection is counted against the AP platform capacity on that controller.

Example: A 7005 controller can accommodate 16 combined active and standby APs.

## HA Enhancement Features

ArubaOS 6.4 introduced several major enhancements to the AP Fast Failover introduced in 6.3 to further improve controller failure detection and 802.1X client reconnection, as well as introducing a platform over-subscription feature to better support the N+1 model. Another feature added was to bring HA support to the Master-Master redundancy topology.

Here is a brief description of each enhancement:

- **Inter-Controller Heartbeat** – Rather than waiting for the AP eight consecutive heartbeats to detect a controller that has become unreachable, a heartbeat sent by the HA Standby controller and acknowledged by the HA Active controller was implemented to quickly detect an Active controller failure (~1 second). See [Inter-Controller Heartbeat on page 40](#) for more information.
- **Client State Synchronization** – This feature reduces the time taken by dot1x clients to reconnect after a controller failover by synchronizing the PMK entries between Active and Standby controllers. As a result, only a 4-way handshake takes place rather a full dot1x authentication. See [Client State Synchronization on page 44](#) for more information.
- **Capacity Extension** – This feature is also known as **N+1 Over-Subscription**. This feature supports the N+1 deployments where one controller is dedicated to backup N LMS controllers. It overcomes the platform AP capacity restriction to allow more 'standby' APs to connect with standby tunnels beyond the platform capacity limitation. See [Capacity Extension on page 47](#) for more information.
- **Master-Master Redundancy** – In ArubaOS 6.4, a standby master using the master-master redundancy feature is permitted to terminate 'standby' tunnels from APs, thus supporting the HA-AP Fast Failover feature. See [Master Redundancy on page 51](#) for more information.

### Inter-Controller Heartbeat

The inter-controller heartbeat enhancement was introduced in ArubaOS 6.4 to provide faster AP failover in case the LMS controller is unreachable.



---

Inter-controller heartbeat works independently from the AP mechanism that sends heartbeats from the AP to the controller and it supersedes the AP's heartbeat to its controller, since it is able to detect an unreachable active controller much faster than AP heartbeats.

---

This section includes the following topics:

- [Introduction on page 40](#)
- [Failover Scenario on page 41](#)
- [Inter-Controller Heartbeat Flow on page 43](#)

### Introduction

The inter-controller heartbeat includes the following features:

- Faster detection of active controller failure.
- AP failover time less than 1 second.
- All controller platforms except 650/620 are supported.
- The Active/Active, Active/Standby, and N+1 topologies are supported.
- A single standby can simultaneously heartbeat up to seven (7) active controllers.
- Independent of the traditional AP to controller heartbeat mechanism.

Once the inter-controller heartbeat feature is enabled, the HA standby controller sends heartbeats to the HA active controller every 100 msec (default), and if five (5) consecutive heartbeats are missed (not acknowledged), the standby controller instructs the APs to failover.





Inter-controller heartbeat interval and threshold values may need to be changed from their default values to account for link latency and network congestion between the HA active and standby controllers.

It is to be noted that an AP may still fail over to its HA Standby controller, independently of the Inter-Controller Heartbeat feature, if it misses 8 consecutive heartbeats along its path to its HA Active controller.

The Inter-Controller Heartbeat feature is not supported by the Master Redundancy topology due to the dependency of this topology on VRRP heartbeat timeouts for its failover.


The Heartbeat frame is a PAPI message that originates from the control-plane of the HA Standby controller-ip to the HA Active controller-ip, and awaits acknowledgment by the data-plane of the HA Active controller.

### Failover Scenario

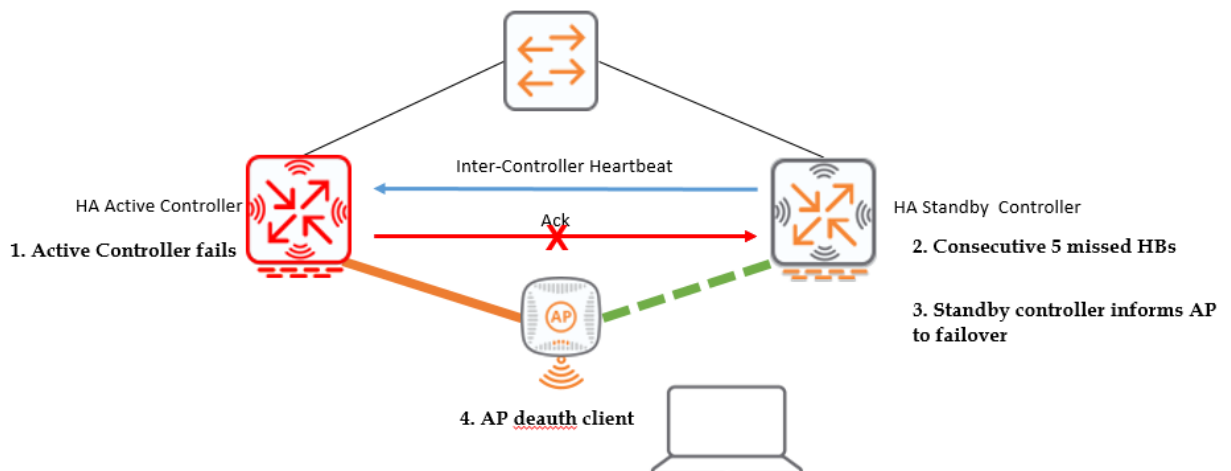
An Inter-Controller Heartbeat failover scenario is illustrated in [Figure 23](#), [Figure 24](#) and [Figure 25](#).

1. Active controller fails.
2. Standby controller misses 5 consecutive heartbeats after 500 msec.
3. Standby controller instructs AP to fail over.
4. AP sends 802.11 de-authentication frames to its associated clients.
5. AP tears down its tunnel to failed Active controller.
6. AP existing Standby tunnel becomes Active to new Active controller.

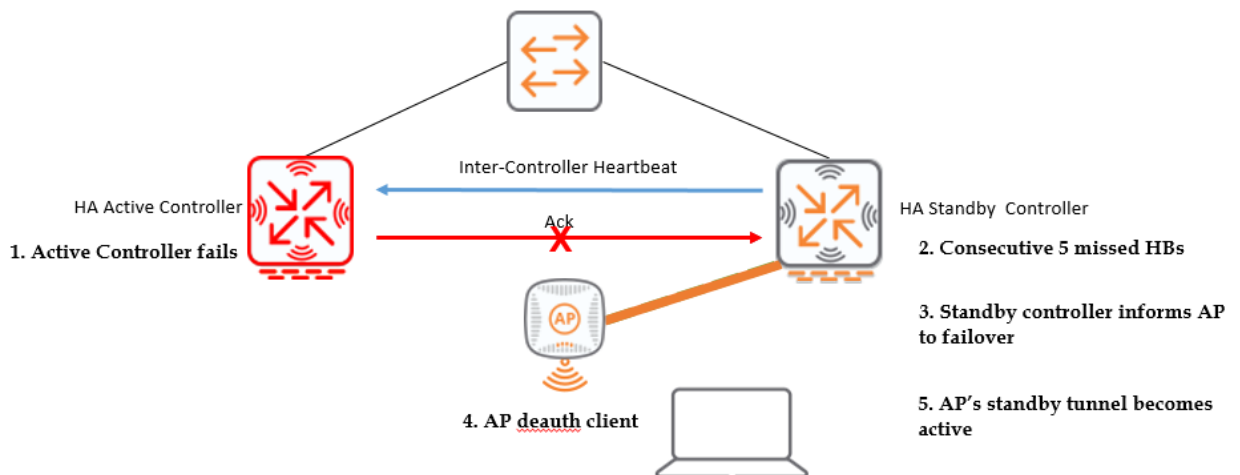
**Figure 23** *Inter-Controller Heartbeat Failover Scenario 1*



**Figure 24** *Inter-Controller Heartbeat Failover Scenario 2*



**Figure 25** *Inter-Controller Heartbeat Failover Scenario 3*

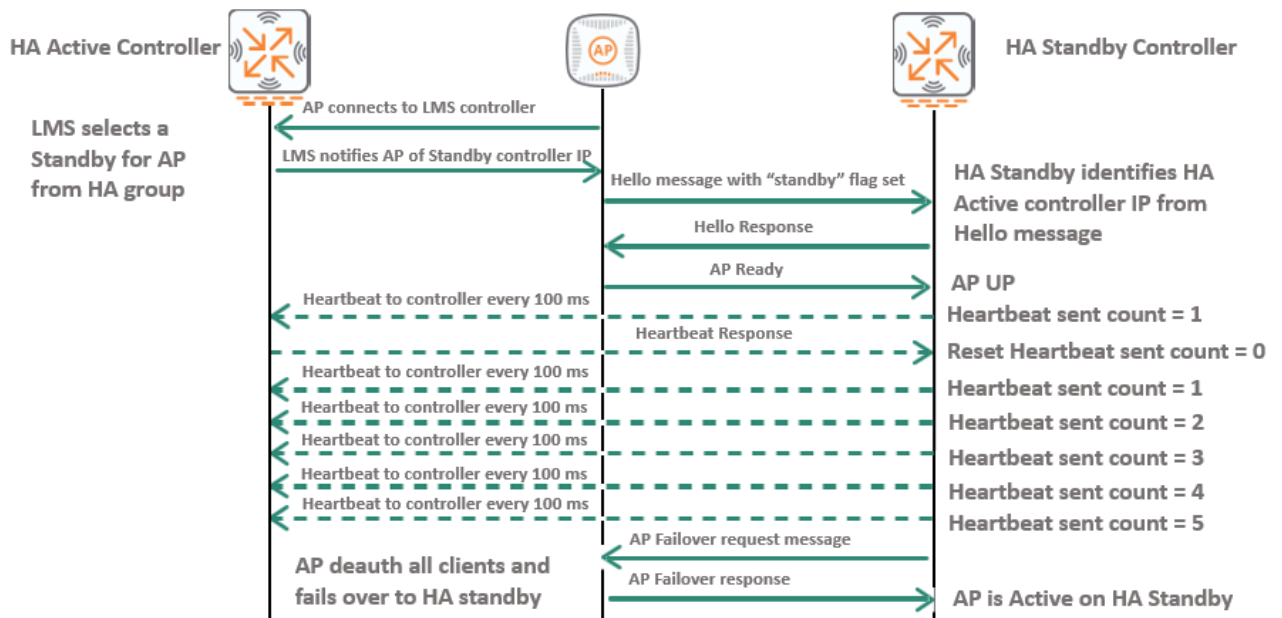


## Inter-Controller Heartbeat Flow

Before the heartbeats begin, the AP connects to the active controller. The active controller notifies the AP of the standby controller IP address. The AP contacts the standby controller with a Hello message with the "standby" flag set, to ensure that it has connectivity to the standby controller. The standby controller responds to the Hello message sent by the AP. The AP is in the UP status on the standby controller with the standby tunnel set. The AP informs its active controller of its successful standby status on the designated standby controller.

Figure 26 illustrates the inter-controller heartbeat flow in the AP fast failover.

**Figure 26** Inter-Controller Heartbeat Flow



The flow includes the following steps:

1. Once the first AP to establish a standby tunnel with its HA standby controller is UP, that controller learns the IP address of the active controller and sends its first heartbeat to that controller, and increments its heartbeat counter by 1.
2. If an acknowledgment heartbeat is received from the active controller, the standby heartbeat counter is reset to zero.
3. In the event of the active controller failure, the standby heartbeat counter keeps increasing with every unacknowledged heartbeat sent up to the configured missed heartbeat threshold (5 by default).
4. Once that threshold is reached, the standby controller sends a failover request message to the AP.
5. Upon receiving that failover request from its standby, the AP de-authenticates all its associated clients, drops its existing tunnel(s) with the active controller, and fails over to its HA standby controller.
6. The existing tunnel(s) to the standby changes state from standby to active and becomes available to handle clients' traffic.

## Client State Synchronization

This section includes the following topics:

- [Introduction on page 44](#)
- [Key Considerations on page 44](#)
- [Failover Scenario on page 45](#)

### Introduction

The client state synchronization feature provides the following enhancements:

- Shortens the time the 802.1X client connects after an active controller failover.
- Synchronizes 802.1X pairwise master key (PMK) entries between the HA active and standby controllers.
- Full 802.1X does not occur after a failover when the 802.1X client attempts to reconnect to the 802.1X SSID.
- When 802.1X clients reconnect to the network, they only have to perform the 4-way key handshake exchange, without the full extensible authentication protocol (EAP) exchange.
- There is no communication with the back-end RADIUS authentication server as full authentication is not required.

To further speed up client reconnection after failover, we recommend maintaining the same user virtual local area network (VLAN) and subnet in the 802.1X SSID configuration and between the HA active and standby controllers so that clients can retain their IP addresses.



---

When the client state synchronization feature is enabled, although there is no authentication request to the RADIUS server upon reconnecting, the failover and reconnection of the client does trigger a new accounting session with the RADIUS accounting server.

---

### Key Considerations

Client state synchronization key considerations include the following:

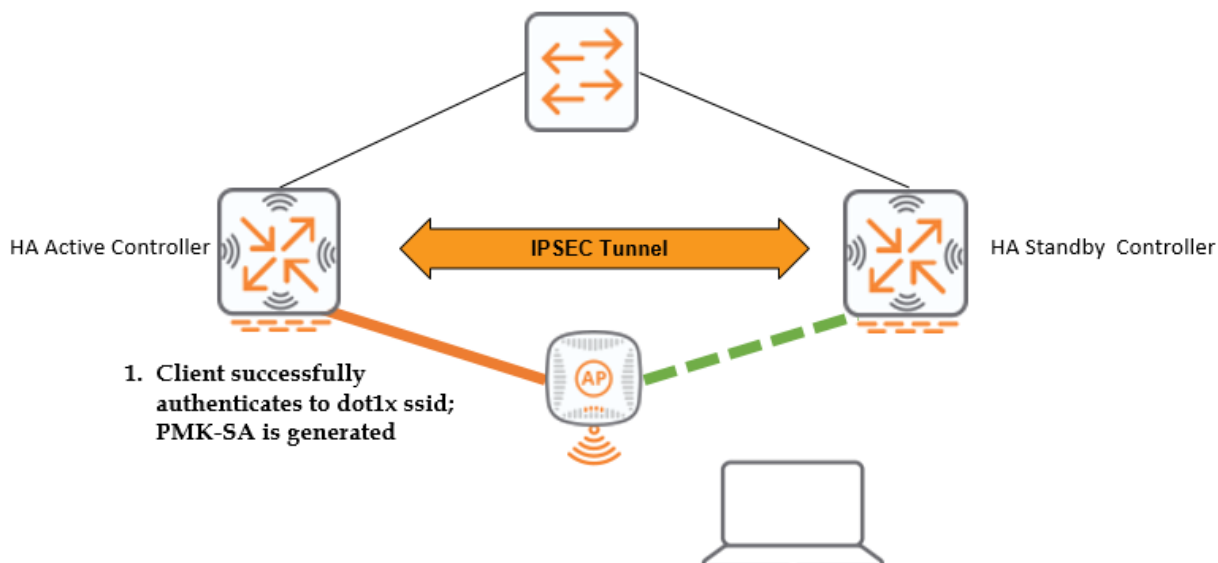
- Client state data synchronization between HA active and standby controllers occurs securely through an IPsec tunnel. Data passing between the two controllers cannot be compromised.
- Supported only on the following HA deployment models between two controllers:
  - Active - Active (1:1)
  - Active - Standby (1+1)
- Not supported on the 600 and 3200/3400 platforms. Supported on the M3, 3600, 70xx, and 72xx platforms.
- Mutually exclusive with the N+1 over-subscription feature. You cannot enable the features together; use either one of the two.

## Failover Scenario

Enabling the client state synchronization feature when there is an HA active controller and standby controller establishes an IPsec tunnel between the two controllers. PMK data passes securely between the two controllers.

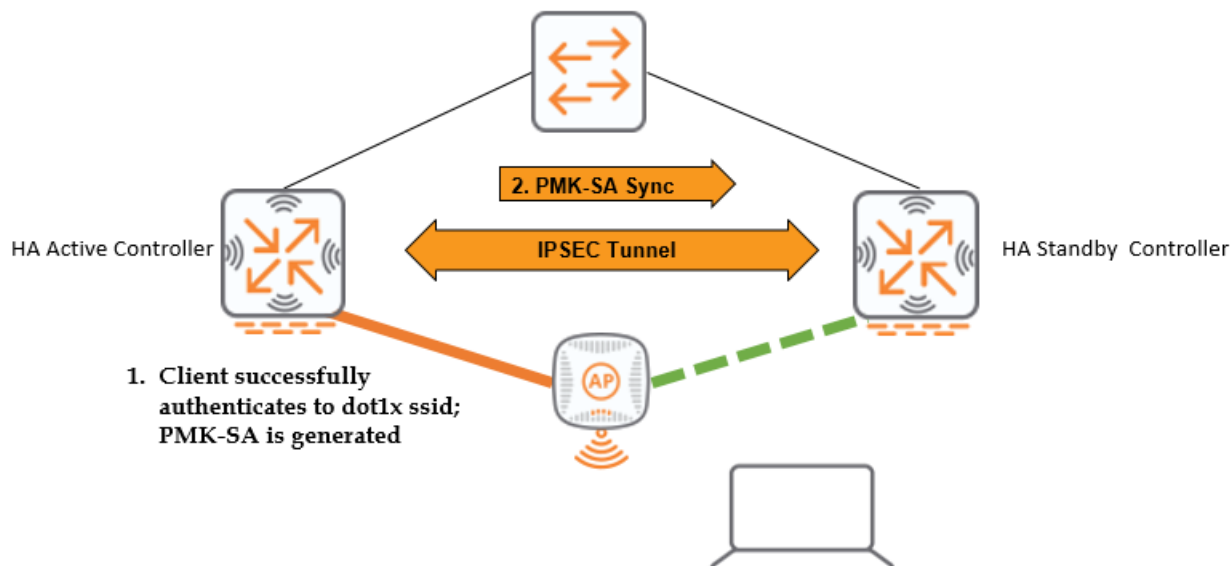
A client state synchronization failover scenario is illustrated in [Figure 27](#), [Figure 28](#), [Figure 29](#), and [Figure 30](#).

**Figure 27** Client State Synchronization Failover Scenario Step 1



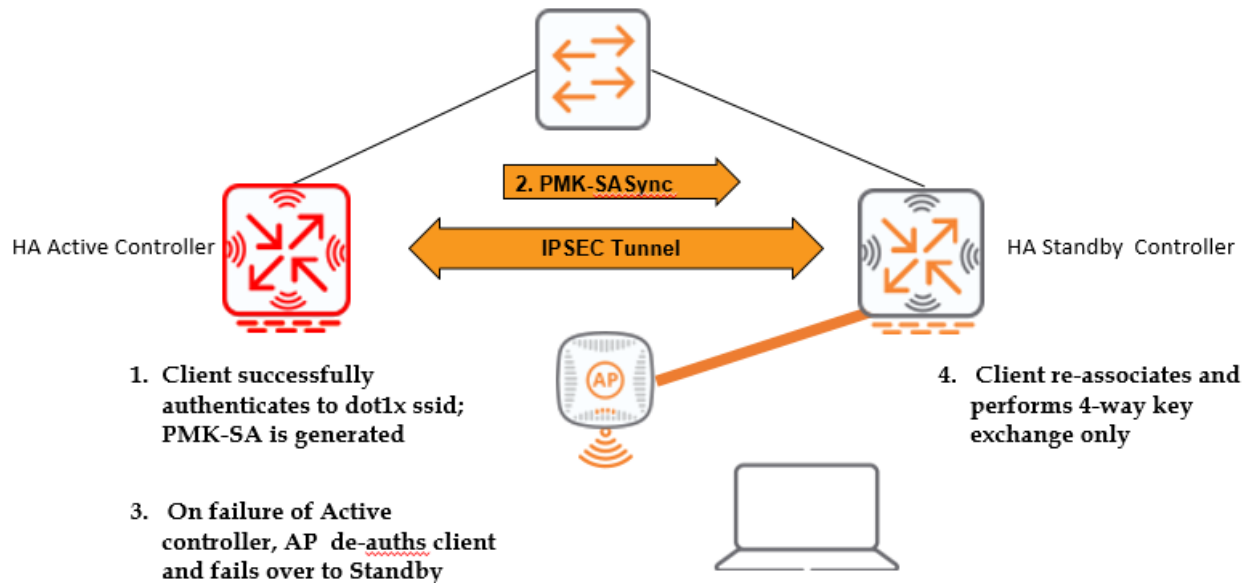
1. Client successfully connects to the 802.1X SSID. A PMK-SA is generated.

**Figure 28** Client State Synchronization Failover Scenario Step 2



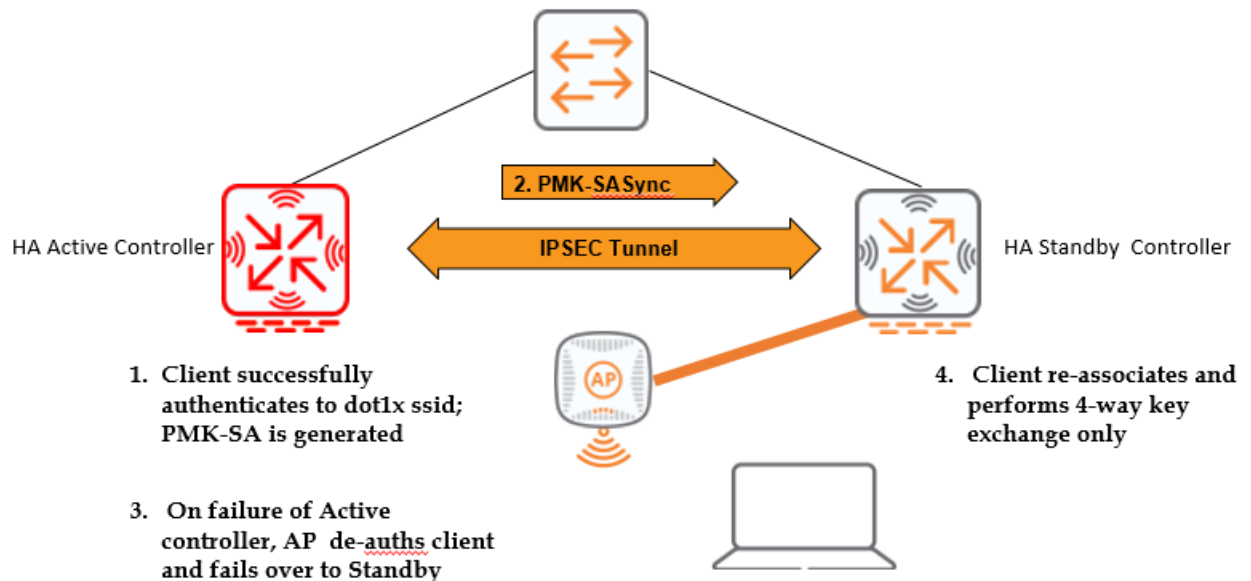
2. PMK-SA data is synchronized from the active controller to the standby controller inside the IPsec tunnel.

**Figure 29** Client State Synchronization Failover Scenario Step 3



3. If the active controller fails, the AP de-authenticates its clients and fails over to the standby controller.

**Figure 30** Client State Synchronization Failover Scenario Step 4



4. As soon as the AP fails over, the client re-associates again, and the 802.1X client performs a 4-way key exchange only.

## Capacity Extension

This section includes the following topics:

- [Introduction on page 47](#)
- [Standby AP Over-Subscription on page 48](#)
- [Standby AP Over-Subscription Example on page 49](#)
- [N+1 Over-Subscription on page 50](#)

### Introduction

The capacity extension (also known as N+1 over-subscription) feature includes the following:

- An HA standby controller will be able to terminate standby AP tunnels above its AP platform limit, thus overcoming the controller's original rated AP capacity in ArubaOS 6.3.
- APs with active tunnels remain restricted to the controller platform limit.
- Enabling centralized licensing is required for this feature.
- Supported on 72xx, M3, and 3600 controllers, only.
  - 72xx controllers allow four times over-subscription.
  - M3 and 3600 controllers allow two times over-subscription.
- Mutually exclusive with the HA client state synchronization feature. Only one of the two features can be enabled at any time.
- Upon a failover, once the HA standby controller platform limit is reached by failed over active APs/tunnels, all standby tunnels are terminated and their corresponding APs will need to move to a different HA standby controller, if available.

The capacity extension feature finds its application in N+1 deployments, where one (1) HA standby controller acts as a backup to (N) HA active controllers of equal capacity.

## Standby AP Over-Subscription

[Table 4](#) describes the standby AP over-subscription by platform. It contains:

- A list of the platforms
- Maximum number of APs per platform
- Maximum GRE tunnels
- Over-subscription ratio

The over-subscription ratio is based on the platform:

- 70xx platforms - over-subscription is not supported
- 3600 and M3 platforms - over-subscription supported for up to two controllers
- 72xx platforms - over-subscription supported for up to four controllers

For example, a single 7220 controller acting as standby can support standby AP tunnels for up to four other 7220 controllers that are HA active. However, the combined total of the GRE tunnels cannot exceed the platform limit of 32,768 tunnels (a combination of active and standby GRE tunnels).

**Table 4:** *Standby AP Over-Subscription by Platform*

Platform	Max # APs	Max GRE Tunnels	Over-Subscription Ratio
7005	16	512	n/a
7010	32	1024	n/a
7024	32	1024	n/a
7030	64	2048	n/a
3600	128	16384	2:1
M3	512	16384	2:1
7205	256	8192	4:1
7210	512	16384	4:1
7220	1024	32768	4:1
7240	2048	65535	4:1



## Standby AP Over-Subscription Example

Following is an example of AP over-subscription using the 7210 platform:

- One HA Standby 7210 controller is backing up four 7210 HA Active controllers.
- Each Active controller is loaded with its maximum AP capacity of 512 APs.

Therefore, the Standby 7210 over-subscription ratio is:

$$512 * 4 = 2048 \text{ APs}$$

Let us consider two cases:

Case 1: Average basic service set identifiers (BSSIDs) per AP is 8 (or 4 dual radio SSIDs).

Tunnel capacity with 8 BSSIDs:

$$2048 * 8 = 16,384 \text{ tunnels}$$

Since the total tunnels for 8 BSSIDs does not exceed the 7210 total tunnel capacity, we conclude that the Standby 7210 is able to back up the 2048 APs from the four Active controllers.

Case 2: Average BSSIDs per AP is 10.

Tunnel capacity with 10 BSSIDs:

$$2048 * 10 = 20,480 \text{ tunnels}$$

The 20,480 tunnels needed exceed the 7210's total tunnels capacity. Therefore, we conclude that the 7210 Standby controller would not be able to back up all 2048 APs.

With an average of 10 BSSIDs per AP, the 7210 Standby controller will be able to back up a combined:


$$16,384 \text{ tunnels} / 10 = 1,638 \text{ APs only.}$$

## N+1 Over-Subscription

N+1 over-subscription is illustrated in [Figure 31](#), [Figure 32](#), and [Figure 33](#)


In [Figure 31](#) there are four active controllers with 512 APs terminating on each, and one Standby 7210 controller.

**Figure 31** Over-Subscription 1




In [Figure 32](#) the green standby tunnels terminate from all 2048 APs to the standby 7210 controller.

**Figure 32** Over-Subscription 2



[Figure 33](#) shows one active controller that failed, and its corresponding 512 APs switch the state of their tunnels to the HA standby controller from standby to active.

**Figure 33** Over-Subscription 3



## Master Redundancy

This section includes the following topics:

- [HA Support on page 51](#)
- [HA Constraints on page 51](#)
- [HA Failover on page 52](#)

### HA Support

Prior to ArubaOS 6.4, the Standby Master controller did not terminate AP tunnels nor forward any user traffic. All AP tunnels carrying user traffic terminated on the VRRP Active Master only.

The HA support for the Master Redundancy in ArubaOS 6.4 introduced the ability for the HA Standby Master to terminate AP standby tunnels only. No user traffic or active tunnels would reach the Standby Master until a VRRP failover takes place and its state changes to Active.

### HA Constraints

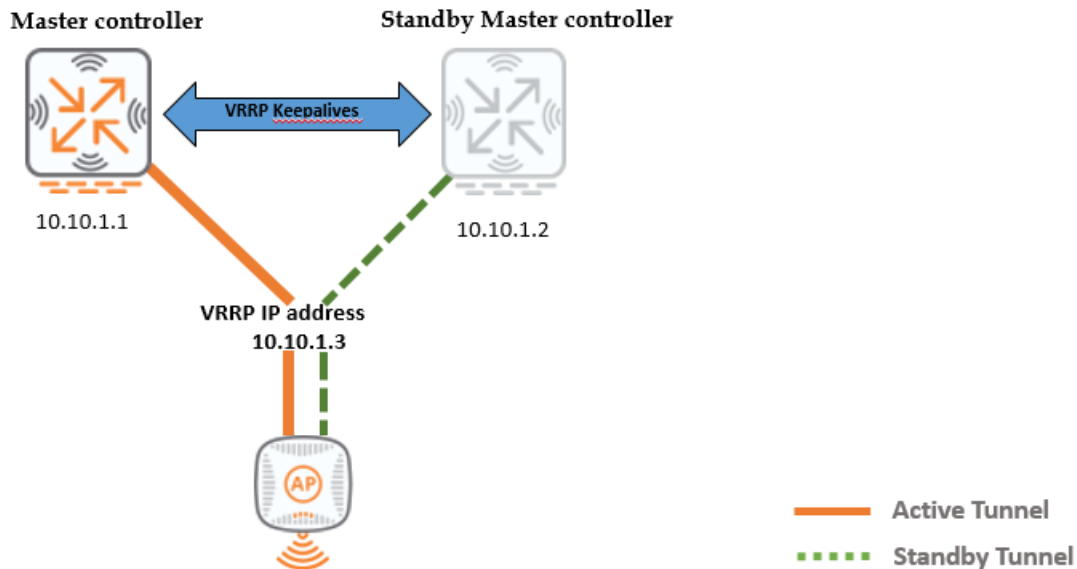
Following are the HA constraints with the master redundancy architecture:

- HA role must be 'dual' for both Active and Standby Master controllers. When configuring the HA group profile, both Master and Redundant Master controller-ip addresses must be listed with the 'dual' role, because at any time there could be a change of state and both controllers could assume the active or standby role. That is why the 'dual' role is the only role that qualifies in this architecture
- Inter-controller heartbeat is not supported. VRRP keepalives control the failover in the master redundancy architecture. Therefore, the inter-controller heartbeat is not supported and should not be enabled. As a result, sub-second failover cannot be achieved. However, the HA AP fast failover benefit resides with failover scalability with large number of APs due to pre-built CPSec and/or GRE tunnels.
- The capacity extension feature is not supported. There is no requirement for over-subscription in this architecture since both controllers need to have the same capacity to back each other up.

## HA Failover

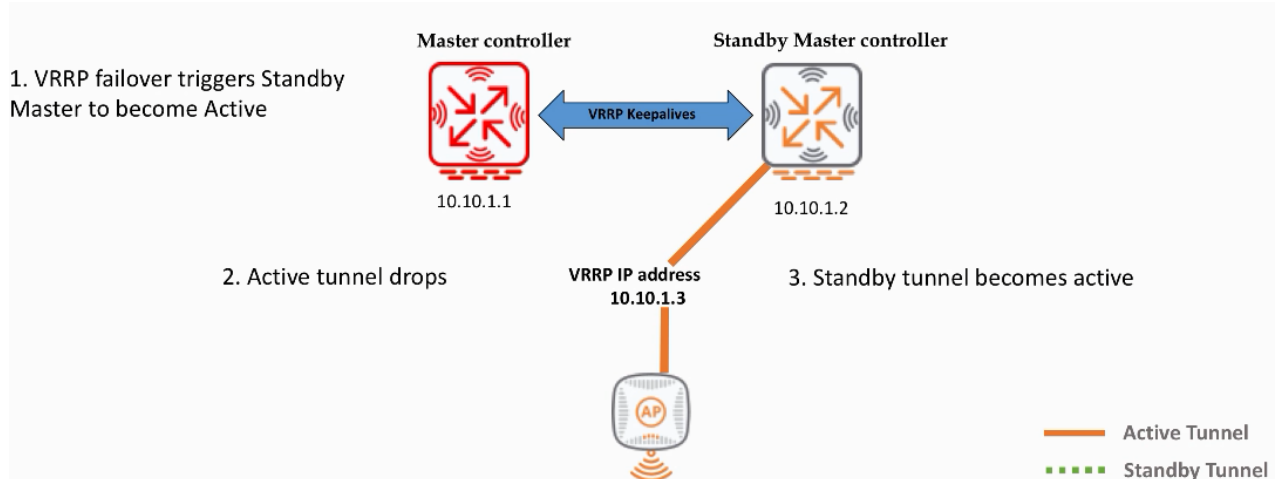
[Figure 34](#) and [Figure 35](#) illustrate the HA failover in the Master Redundancy architecture.

**Figure 34** *Master Redundancy*



In [Figure 34](#) there is an active master controller and a standby master controller. The active tunnel goes to the active master 10.10.1.1. The standby tunnel goes to the standby master 10.10.1.2.

**Figure 35** *Master Failover*



1. When the Active master fails, the Standby master detects the failure after three (3) consecutive keepalives missed, and changes its state from Standby to Active master.
2. The AP active tunnels to the failed master are terminated.
3. The standby tunnels from the AP to the new Master change from standby to active.

Although we lose the failover speed that we had with the inter-controller heartbeat, we gain in scalability by the fact that CPSec and GRE tunnels are already pre-built. A large number of APs failing over won't add any further failover delay compared to legacy redundancy.

## HA Support for Bridge Mode

HA support for Campus APs (CAPs) in bridge mode was introduced in ArubaOS 6.4.1 and beyond.

The following HA features are supported:

- Inter-controller heartbeat
- Client state synchronization - 802.1x authentication in bridge mode still takes place at the controller, and the PMK SA is still synchronized between the HA Active and Standby.
- Capacity extension

HA with Campus APs in bridge mode is only supported on the 72xx platform.



---

Bridge mode requires CPSec to be enabled.

---

---

Bridge mode supported on a maximum of 32 CAPs per controller.

---

## HA Supported AP Forwarding Modes

[Table 5](#) lists the HA supported AP forwarding modes for each redundancy feature. Each redundancy feature supports the tunnel, d-tunnel, and bridge forwarding modes.



---

HA features are not supported for the RAP split-tunnel mode.

---

**Table 5:** HA Supported AP Forwarding Modes

Redundancy Feature	Tunnel	D-Tunnel	Bridge
Active/Standby Tunnels	✓	✓	✓
Inter-Controller Heartbeat	✓	✓	✓
Client State Synchronization	✓	✓	✓
Capacity Extension (N+1 Over-Subscription)	✓	✓	✓
Master Redundancy	✓	✓	✓

The Centralized Licensing feature was introduced in ArubaOS 6.3 as follows:

- A master controller is designated as the license server.
- All attached local controllers are license clients.
- Licenses already loaded on existing local controllers are automatically shared as part of a global centralized license pool on the master controller.

The centralized licensing feature has these advantages:

- Savings in WLAN licensing costs. Often there are licenses on local controllers that go unutilized because they were bought in bulk, and not all of the licenses were used. These unutilized licenses are shared in the license pool and used by other controllers.
- Ease of license management. It reduces the complexity of quoting and maintaining licenses in a multi-controller environment.
- Elimination of tedious capacity planning. There is no need to plan for specific licenses on specific local controllers. You just plan the total licenses needed for the whole WLAN network. You purchase licenses and load them on the license server.



You can only designate a master controller as a license server; a local controller cannot be selected as a license server.

We strongly recommend selecting a backup license server from the other available masters. This is for redundancy purposes and as a best practice.

[Table 6](#) lists the licenses that can be distributed from the license pool.

**Table 6:** *Eligible Licenses to be Distributed from the License Pool*

Eligible Licenses	Ineligible Licenses
Access Point (AP)	Built-in
Policy Enforcement Firewall (PEF) / Policy Enforcement Firewall Next Generation (PEFNG)	Policy Enforcement Firewall Virtual Private Network (PEF-V)
RF Protect	
Extreme Security (xSec)	
Advanced Cryptography (ACR)	
Unexpired Evaluation	
Web Content Classification (WebCC) (ArubaOS 6.5)	



Built-in licenses are not sent to the license server. However, since they exist on the local controllers, they are used first before the local controller consumes licenses from the license server.

This chapter includes the following topics:

- [Supported Topologies on page 55](#)
- [Additional Supported Topology in ArubaOS 6.5 on page 56](#)
- [License Server Redundancy Failover on page 57](#)
- [Single License Server Failure - No Redundancy on page 58](#)
- [Frequently Asked Questions on page 59](#)

## Supported Topologies

The centralized licensing supported topologies are illustrated in [Figure 36](#), and include the Master-Local and All Independent Masters environments.

In the Master-Local environment, two master controllers are deployed in a redundancy mode with several attached local controllers. The active master acts as the license server with the standby master acting as the standby license server, while the locals are license clients.

In the All Independent Masters environment, one master is selected as the license server and all other standalone masters are license clients. However, for the sake of providing redundancy to the license server, it is highly recommended to select one of the standalone masters to act as a standby license server.

**Figure 36** *Supported Topologies*



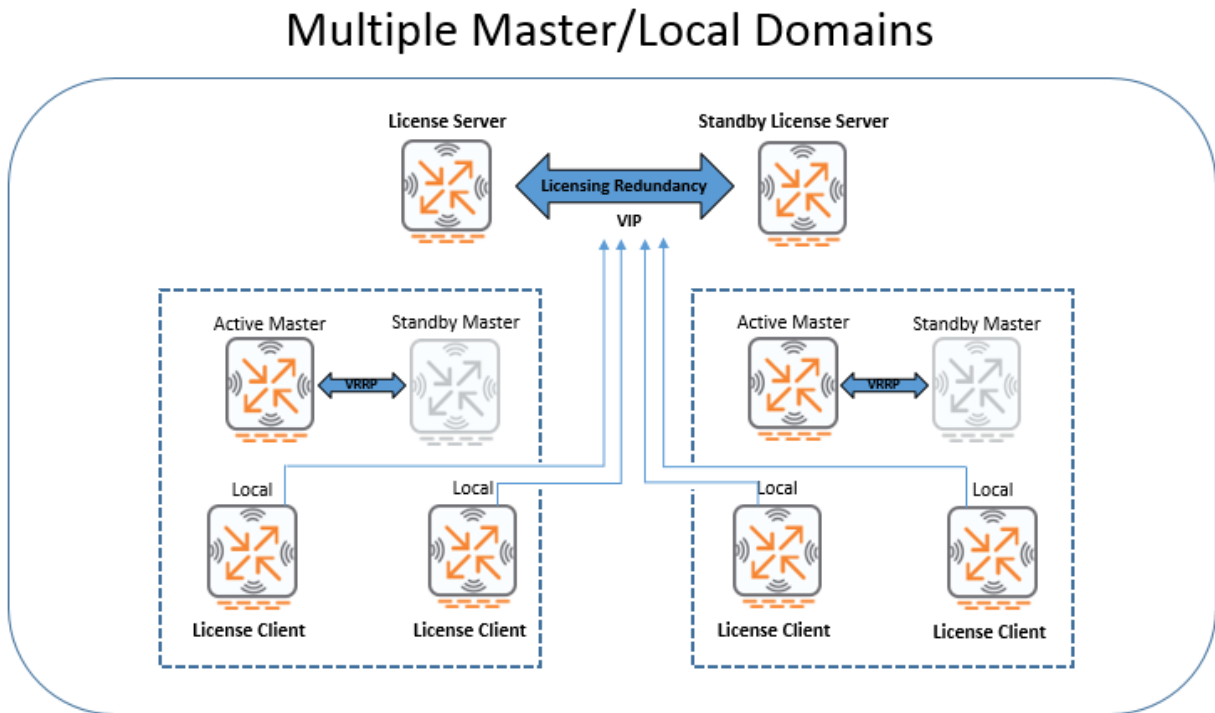
One other popular topology in large deployments is the multiple Master-Local domains. The master controller in each domain acts as the license server for the local controllers attached to it, but a global license server for all domains is not supported in ArubaOS 6.4.x. The next section will introduce such support starting with ArubaOS 6.5.

## Additional Supported Topology in ArubaOS 6.5

Starting with ArubaOS 6.5, a master controller could act as a license server for multiple Master-Local domains where each of the locals in the different domains are license clients to that Master license server.

For redundancy purposes, it is highly recommended to use a redundant Standby Master to act as the standby license server.

**Figure 37** *Additional Supported Topology*

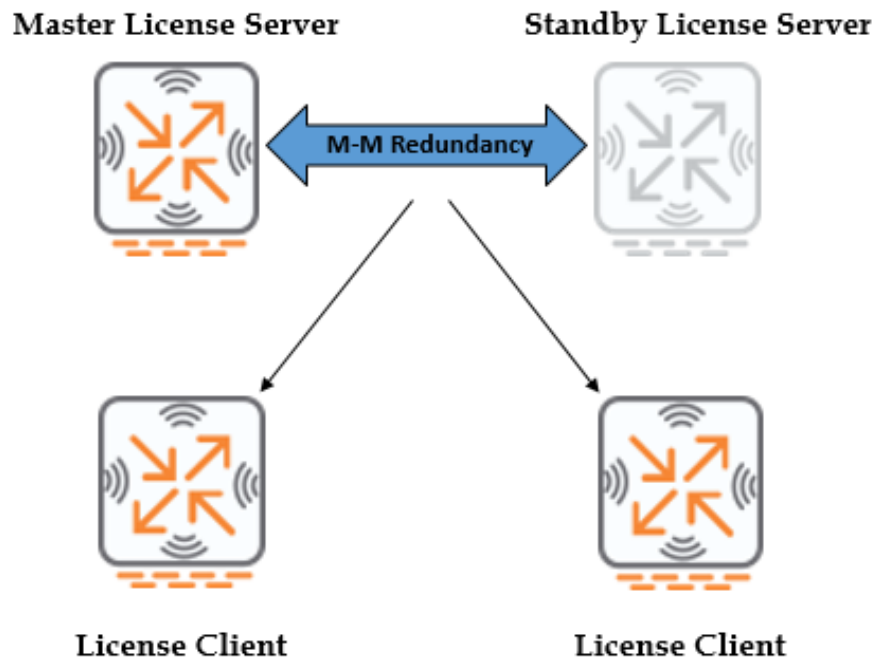




## License Server Redundancy Failover

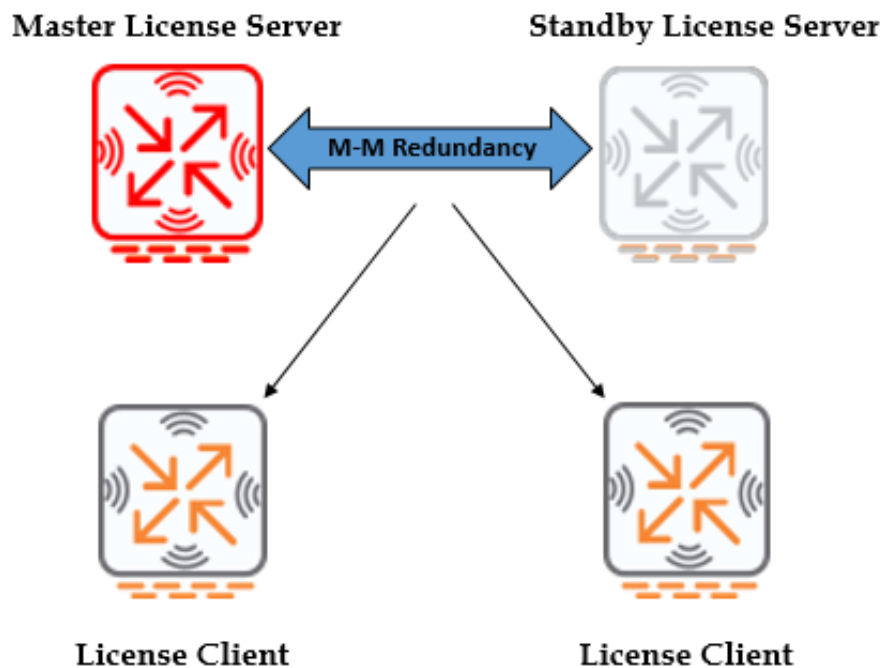
License clients receive licenses from their master license server as illustrated in [Figure 38](#).

**Figure 38** Master License Server



If the active master license server fails as illustrated in [Figure 39](#), then the standby license server acts as the license server.

**Figure 39** Standby License Server




## Single License Server Failure - No Redundancy

If a single master license server fails with no standby to take over its duties as the master license server:

- Existing controllers that received licenses from that master license server continue to function for up to 30 days while the master is down.
- Newly started controllers do not receive any licenses because there is not an active master license server.
- A warning message displays on the license client login screen if the license master is not available.
- After 30 days of no license client reach ability to the master, imported licenses from that client to the master license database are removed at the master level.

**Figure 40** *Single License Server*



## Frequently Asked Questions

Frequently asked questions regarding centralized licensing are listed in [Table 7](#).

**Table 7:** *Centralized Licensing Frequently Asked Questions*

Question	Answer
With which controllers does centralized licensing work?	The feature works for all controllers that support ArubaOS 6.3. Controllers running a previous ArubaOS version cannot use centralized licensing.
What if the standby license server is rebooted when the master is down?	License limit information is persistent in the standby license server.
Can I use evaluation licenses in the pool?	Yes, unexpired evaluation licenses are treated the same way as permanent licenses in terms of calculating total license counts. When an evaluation license expires, the limits are updated.
Where do I install licenses in a new installation?	For a new installation, install all licenses on the license server.
How do I export licenses that have been installed on a controller previously?	The master license server imports the installed licenses from the controller and updates the license limits. There is no need to manually move licenses between controllers.
What are the requirements for the standby license server?	The license server leverages VRRP to determine which controller becomes the master license server. This means that the master and standby license servers must be on the same L2 domain.
Does centralized licensing support multi-version?	Starting with ArubaOS 6.4.3.0, controller multi-version is supported in an all-master deployment provided that the license master is running 6.4.3, and the license client must be running a minimum version of 6.3.1.9, 6.4.2.3, or 6.4.3.x.
Can I use centralized licensing if I only have two controllers?	Yes, but most enterprises have more than two controllers. An existing master / standby master controllers would be good candidates to become license servers.
Do the local controllers need to be rebooted if new licenses are installed on the centralized license server?	There is no longer a requirement to reboot the controller after adding or deleting a new license type.
How do we "clear and release capacity" back to the pool?	The license server table is updated every 30 seconds. Changes to license usage are reported.
What happens if the standby license server or license client is rebooted when the active license server is not available?	The information is persistent in the standby license server across reboots.
In an all-master deployment if the customer has setup VRRP between peer controllers for redundancy, can they re-use the same VRRP IPs for centralized licensing?	Yes. For an all masters deployment use the command 'license server-redundancy' to configure the standby license controller.
Can I have a dedicated controller to provide licenses?	There are no requirements to have a dedicated controller to perform licensing functions. However, if a dedicated license master is desired, we would strongly suggest a 7200 series controller to accommodate additional use cases that may occur.
Do we need to reboot the controller after enabling centralized licensing?	No, enabling centralized licensing does not require controller reboot.

In previous chapters we discussed the legacy redundancy, using VRRP or LMS and Backup-LMS. We saw the pros and cons of those redundancies. We discussed the HA AP fast failover, which is the newer redundancy introduced in ArubaOS 6.3 and 6.4. We highlighted the advantages of the HA redundancy compared to the VRRP or LMS and Backup-LMS. The failover is faster, scalable with large numbers of APs when CPSec is enabled, and when the AP fails over, the radios and SSIDs do not disappear from the air.

This chapter includes the following topics:

- [HA AP Fast Failover Redundancy Guidelines on page 60](#)
- [Migration from VRRP to HA AP Fast Failover on page 61](#)
- [Migration from LMS and Backup LMS to HA AP Fast Failover on page 63](#)
- [Migration for Master-Local Deployment on page 65](#)

## HA AP Fast Failover Redundancy Guidelines

This section discusses the guidelines for migrating from the legacy redundancy to the HA AP fast failover redundancy.

- If a deployment model is using master redundancy through the redundancy feature or VRRP between two standalone masters, the VRRP VIP should be maintained as the AP master, discoverable via DNS, DHCP option 43 or ADP.
- When configuring the HA group-profile, always use the controller-ip address. This is a requirement to successfully deploy HA. This address is the same as the switch-ip and may be different than an interface IP.
- When migrating from the legacy redundancy, it is highly recommended to configure the lms-ip option in the ap system profile.
- The LMS IP address should match one of the 'dual' or 'active' controller IP addresses in the HA group profile. A word of caution: if the VRRP VIP is used, HA will not work correctly.
- Use the Backup-LMS in the AP system profile to recover from a double outage, meaning that the LMS controller is down while an AP is rebooting.

## Migration from VRRP to HA AP Fast Failover

In the legacy VRRP redundancy (VRRP Solution in [Figure 41](#)) a master redundancy is deployed and the APs terminate on the local controllers. VRRP is configured between the local controllers. VRRP redundancy is established by configuring the APs LMS-IP as the local VRRP VIP (10.70.217.5 in this diagram), and the APs terminate their tunnels on the local controller that is VRRP master, in this case Local1.

The migration from VRRP redundancy to HA AP fast failover redundancy in [Figure 41](#) is accomplished using the following steps:

1. Maintain the master redundancy and keep using the Master VIP (10.70.210.5) in AP master discovery.
2. Configure an HA group profile and add both local controllers' controller-ip (10.70.217.3 and 10.70.217) in DUAL role.
3. Configure the HA group membership on each Local controller to belong to the above HA group profile.
4. Change the LMS-IP value from VRRP VIP (10.70.217.5) to Local1 controller-ip (10.70.217.3).



---

It is recommended to execute the migration to HA AP Fast Failover during a maintenance window, as the above Step 4 will trigger AP bootstrap and cause a short outage until all tunnels are rebuilt to the controller-ip of Local1.

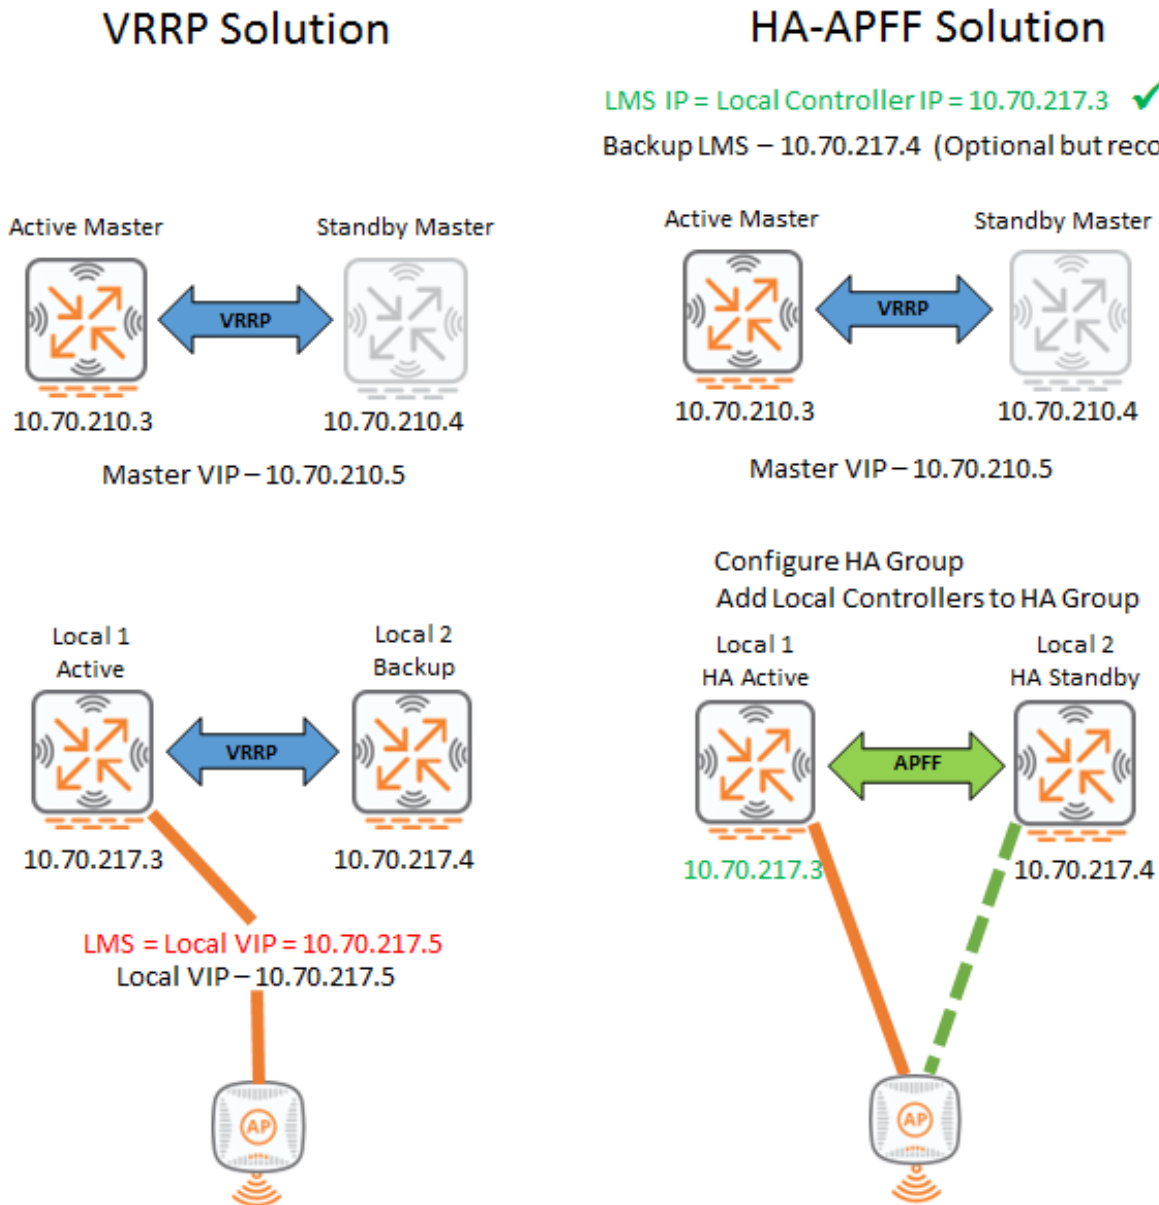
---

---

It is recommended to add Local2 controller-ip as the bkup-lms in the AP system profile.

---

**Figure 41** Migration from VRRP to HA-APFF



## Migration from LMS and Backup LMS to HA AP Fast Failover

In the LMS/Bkup-LMS Solution in [Figure 42](#), there is master redundancy and two local controllers. Rather than using VRRP between the local controllers, we are using LMS and Backup-LMS. The LMS IP can be an interface IP. The AP terminates on its LMS IP (Local 1 in this diagram).

The migration from LMS/Bkup-LMS redundancy to HA AP fast failover redundancy in [Figure 42](#) is accomplished using the following steps:

1. Maintain the master redundancy and keep using the Master VIP (10.70.210.5) in AP master discovery.
2. Configure an HA group profile and add both local controllers' controller-ip (10.70.217.3 and 10.70.217) in DUAL role.
3. Configure the HA group membership on each Local controller to belong to the above HA group profile.
4. Ensure that the LMS-IP value is set to the controller-ip of the local controller desired to be HA Active, Local1 in this case with controller-ip 10.70.217.3.



---

It is recommended to execute the migration to HA AP Fast Failover during a maintenance window, as Step 4 may trigger AP bootstrap and cause a short outage until all tunnels are rebuilt if the LMS-IP is changed from an interface IP to the controller-ip.

---


---

It is recommended to add Local2 controller-ip as the bkup-lms in the AP system profile.


---

Figure 42 Migration from LMS to HA

## LMS/Bkup-LMS Solution



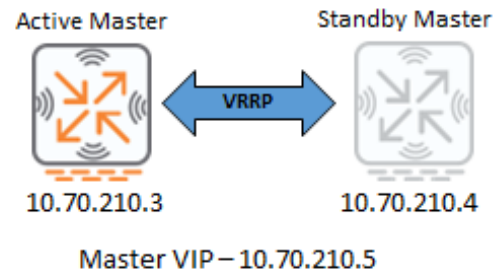
Ok if LMS IP is different than controller-ip



## HA-APFF Solution


LMS IP = Local Controller IP = 10.70.217.3 ✓

Backup LMS – 10.70.217.4 (Optional but recommended)



LMS IP is required to be a controller-ip

Configure HA Group  
Add Local Controllers to HA Group





## Migration for Master-Local Deployment

The Master-Local deployment is a possible deployment at sites that require no more than two controllers. In such deployment illustrated in [Figure 43](#), two controllers are utilized: one Master controller and one Local controller.

In Master-Local with VRRP in [Figure 43](#), the Local controller is the VRRP master, while the Master controller is the VRRP backup. In the VRRP legacy redundancy, the LMS-IP is set to the VRRP VIP (10.70.210.5).

To migrate to the HA AP fast failover as illustrated in Master-Local with APFF Deployment in [Figure 43](#), maintain the VRRP since there is only one master controller. To ensure that there is controller redundancy and that a controller will always respond to an AP when an AP boots up as its AP master, we use the VIP between the master and local controllers as the AP master. This is important for the AP to come up correctly. We deploy AP fast failover and add both the master and local controller IP addresses to the HA group profile. We recommend using the local controller as the HA active; use the LMS IP as the IP address of the local controller (10.70.210.4). Use the Backup-LMS IP as the IP address of the master controller (10.70.210.3).

**Figure 43** Migration Master-Local Deployment



This chapter includes the following topics:

- [Legacy VRRP Failover on page 66](#)
- [HA AP Fast Failover on page 66](#)
- [HA AP Fast Failover vs. Legacy on page 67](#)

## Legacy VRRP Failover

[Table 8](#) describes tests that were done with legacy VRRP failover. There are three different platforms: 7210, 7220, and 7240 controllers. The platform limit lists the maximum number of APs. The tests are based on Active-Standby and Active-Active failover.

**Table 8:** *Legacy VRRP Failover*

Controller Type	7210	7220	7240
Platform Limit	512	1024	2048
Failover: Active-Standby, CPSec ON, AP platform limit (m:s)	0:52	1:28	3:10
Failover: Active-Active (equal AP load), CPSec ON, AP platform limit (m:s)	Not Tested	0:52	1:25

## HA AP Fast Failover

[Table 9](#) describes the HA AP fast failover by controller type, number of AP loads, and the AP failover time.

The failover results in [Table 9](#) were identical with CPSec on or off.

Such results reflect the following facts:

- The short time taken for APs to failover (~ 1s).
- The AP number is not a factor in the failover time.
- No CPSec impact on AP failover time.

We can conclude that besides an aggressive fast AP failover, scalability of the HA solution is a big advantage when considering large number of APs in a campus environment where CPSec and GRE tunnels are pre-established compared to the legacy redundancy.

**Table 9:** *HA AP Fast Failover*

Controller Type	Number of APs	AP Failover Time
7210	512	0.8 sec
7220	1024	1.1 sec
7240	2048	1.3 sec

## HA AP Fast Failover vs. Legacy

[Table 10](#) compares side-by-side the technologies of HA AP fast failover vs. legacy. It provides details regarding the network topology, how fast an AP can detect controller failure, AP failover time, and 802.1X clients.

**Table 10:** *HA AP Fast Failover vs. Legacy*

Redundancy	VRRP	LMS / Backup-LMS	HA APFF (ArubaOS 6.4)
Network Topology	L2	L2 / L3	L2 / L3
Controller Failure Detection by AP	3 sec	8 sec	0.5 sec
AP Failover Time (1024 APs) min:sec	1:28	N/A	1.1 sec
802.1X Clients	Full EAP Exchange	Full EAP Exchange	4-way key exchange only (for same BSSID)

This chapter includes the following different levels of Switches Redundancy:

- [Component Level Redundancy on page 68](#)
- [Device Level Redundancy on page 71](#)
- [Link Level Redundancy on page 82](#)
- [Campus Network Redundancy - Implementation Guidance on page 89](#)

## Component Level Redundancy

This section includes the following topics:

- [Management Module Redundancy on page 68](#)
- [Power Supplies Redundancy on page 69](#)

### Management Module Redundancy

The Aruba 5400R series switch provides redundancy at the component level with redundant management modules (MM), which can provide non-stop switching in the event of a MM failure, synchronizing data between the “active” and “standby” MM.

**Figure 44** *Management Module Redundancy – 5400R*

Continuous switching in the event of a management module failure




**Aruba 5400R – Front View**

There are two modes for MM redundancy:

- Warm Standby Mode
  - Default mode
  - “Active” MM does not synchronize continuously with “Standby” module

- Non-stop Switching Mode
  - “Standby MM is synchronized continuously with “Active” module
  - Switching continues without interruption during failover

**Figure 45** Management Modules (MM) – Aruba 5400R




## Power Supplies Redundancy

Both the 5400R and 3810M switches also have power supply redundancy, providing backup power when one power supply fails.

**Figure 46** Power Supply Redundancy – 5400R/3810

Continuous power in the event of a power supply failure



Aruba 5400R – Rear View



Aruba 3810M – Rear View

With redundant power supplies, when each is connected to a different power source, one power supply can continuously supply power when the other fails.

**Figure 47** Redundant Power Supplies



---

It is important to plan the PoE budget accordingly when using redundant power supplies.

---

Refer to the *HPE ArubaOS-Switch PoE Planning and Implementation Guide*

[http://h20566.www2.hp.com/hpsc/doc/public/display?sp4ts.oid=1008605435&docLocale=en\\_US&docId=emr\\_na-c04344559](http://h20566.www2.hp.com/hpsc/doc/public/display?sp4ts.oid=1008605435&docLocale=en_US&docId=emr_na-c04344559) for more information.

## Device Level Redundancy

This section includes the following topics:


- [Stacking Overview on page 71](#)
- [Virtual Switching Framework on page 73](#)
- [Backplane Stacking on page 77](#)
- [Advantages on page 81](#)
- [Key Considerations on page 81](#)

### Stacking Overview

Aruba switches support two stacking technologies: Virtual Switching Framework (VSF) and Backplane Stacking. Stacking simplifies L2/3 design, configuration, and management; it replaces STP and VRRP. Stacking increases resilience and network bandwidth.


Virtual Switching Framework (VSF), also known as frontplane stacking, utilizes existing switch ports as stacking interfaces.

**Figure 48** *Virtual Switching Framework (VSF)*



Backplane stacking uses dedicated stacking modules and cables to create the switch fabric.

**Figure 49** *Backplane Stacking*






## Virtual Switching Framework

Virtual Switching Framework (VSF) is an ArubaOS-Switch stacking solution using standard switching ports. The following switches are supported:

- 5400R - supports stacking up to 2 members and allows for the switches to be connected in the chain topology.
- 2930F - supports stacking up to 4 members and allows for the switches to be connected in chain and ring topologies.

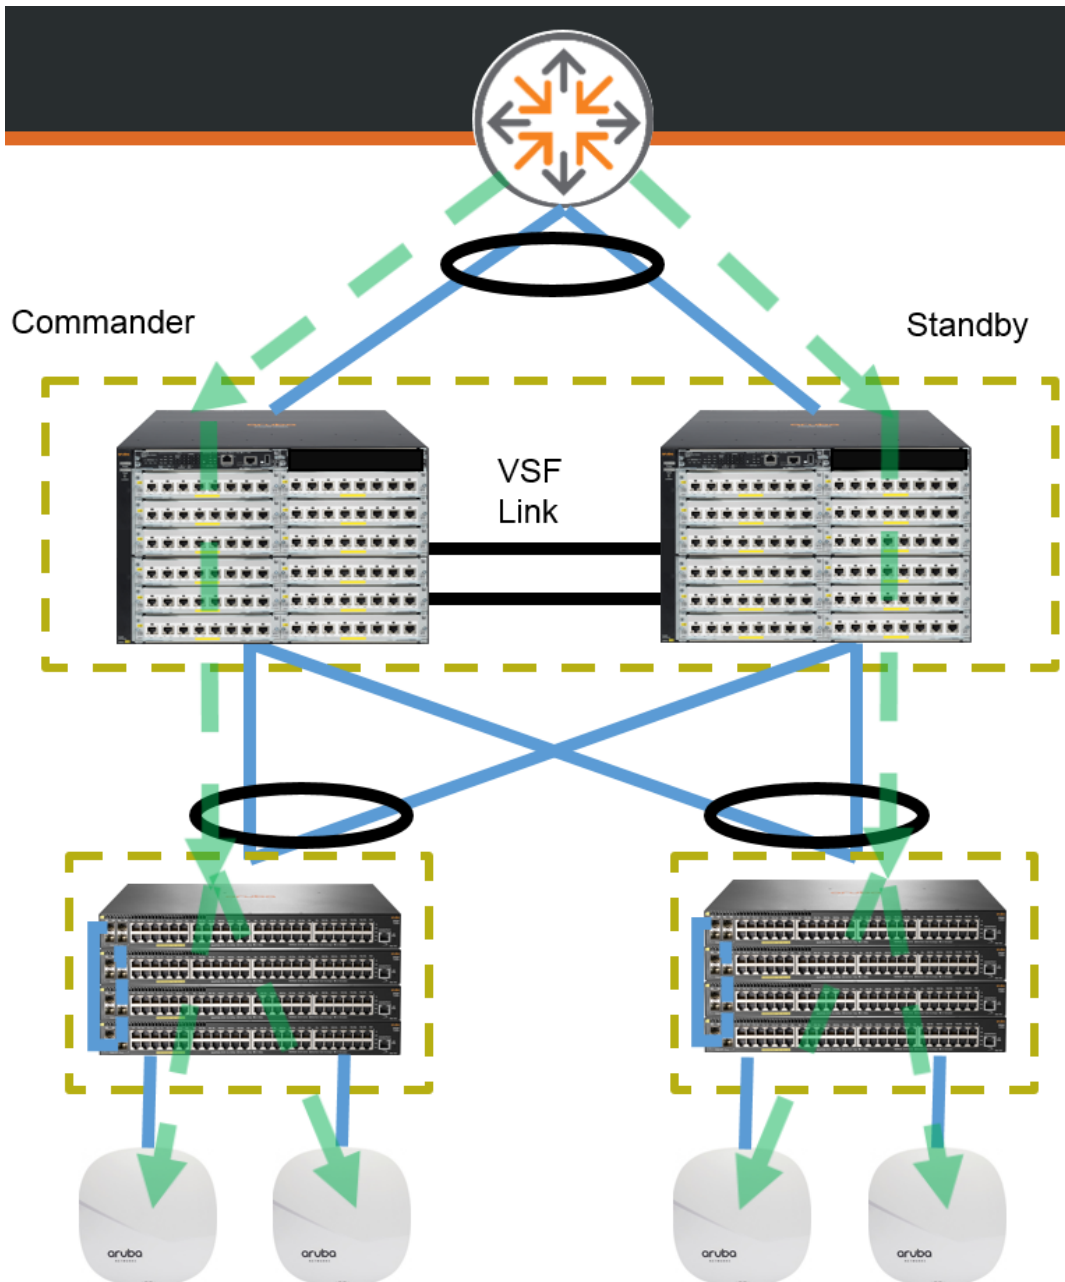
**Figure 50** VSF on 5400R and 2930F



In the following VSF example:


1. Each VSF member is connected to upstream and downstream devices using Link Aggregation Groups (LAGs).

**Figure 51** VSF 1




2. In the case of a stack member failure, traffic fails over to the other stack member. In the case of a stack commander failure, the standby becomes the new commander.

**Figure 52** VSF 2



3. Traffic then resumes on the other link of the LAG.

**Figure 53** VSF 3




## Backplane Stacking

Backplane stacking uses dedicated modules and stacking cables to create a dedicated physical stack link between each switch. The following Aruba stack links are supported:

- 2920 - can stack up to 4 switches in a chain or ring topology.
- 3810M - can stack up to 10 switches in a chain, ring, or 5 switches in a mesh topology.

**Figure 54** *Backplane Stacking on 2920 and 3810M*




This section includes the following topics:

- [Ring Topology on page 78](#)
- [Mesh Topology on page 79](#)


## Ring Topology

The following example shows the path of traffic flowing through a 5 member ring topology, entering stack member 2, and exiting stack member 3.


**Figure 55** *Backplane Stacking Ring Topology 1*



**Figure 56** *Backplane Stacking Ring Topology 2*




**Figure 57** *Backplane Stacking Ring Topology 3*



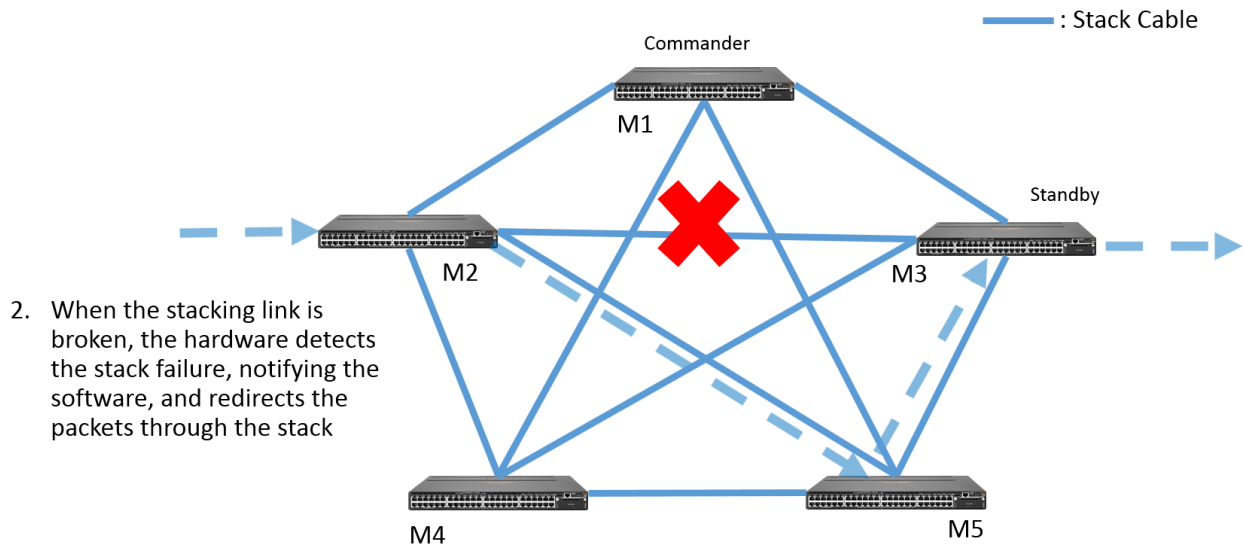
### Mesh Topology

The following example shows the path of traffic flowing through a 5 member mesh topology (3810M), entering a stack member 2, and exiting stacking member 3.


**Figure 58** *Backplane Stacking Mesh Topology 1*



**Figure 59** Backplane Stacking Mesh Topology 2



**Figure 60** Backplane Stacking Mesh Topology 3





## Advantages

Stacking helps eliminate the need for protocols such as VRRP and STP, since a stack has a single management interface.

Link redundancy is achieved using link-aggregation groups from downstream and upstream devices to each stack member. This way, in the case of a stack member failure, upstream and downstream switches will have continued connectivity with traffic failing over the other link of a LAG.

A stack has simplified management. A stack has a single forwarding and routing plane. Both switches act as single, virtual switch.

There is also simple configuration in a stack. Backplane stacking involves simply connecting stacking modules and cables (plug and play). Virtual Switching Framework (VSF) requires connecting stack members and minimal CLI command entry to form a stack.

Backplane stacking also uses a stack mesh topology, where each switch is connected to every other switch in the stack by stacking cables. In the case of a stack member failure, traffic can quickly reach its destination to other stack members with the stacking cable redundancy.

## Key Considerations

Stack software upgrades require network downtime. There is no true In-Service-Software-Upgrade (ISSU) feature. The recently introduced VSF Fast Software Upgrade feature helps minimize network downtime to less than a few seconds rather than normal downtime of approximately 2 minutes.



---

The Aruba 5400R introduced the Fast Software Upgrade feature in the 16.03 software (Jan 2017).

---

Additional stacking hardware can have an increased infrastructure cost. Stacking hardware for switches that support backplane stacking do not come with the switch when purchased by a customer. Additional stacking hardware can increase costs in a stack deployment.

Replacing, adding, or removing switch stack members can be tedious. Generally it is as simple as removing a stack member from the stack with a CLI command and provisioning a new member. The switch member will still need to be manually replaced in the case of a failed switch.

## Link Level Redundancy

This section includes the following topics:

- [Overview on page 82](#)
- [Virtual Router Redundancy Protocol on page 82](#)
- [Link Aggregation Control Protocol on page 83](#)
- [Distributed Trunking on page 85](#)

### Overview

The ArubaOS-Switch provides the following link redundancy features:


- **Virtual Router Redundancy Protocol (VRRP)** - Dynamic failover in the case of a loss of connection to a default router.
- **Link Aggregation Control Protocol (LACP)** - LACP link aggregations provide more bandwidth and redundancy between devices.
- **Distributed Trunking (DT)** - Enables a loop-free and redundant network topology without using STP. DT can be combined with VRRP to provide a L2/L3 HA solution.
- **Spanning Tree Protocols: Spanning Tree Protocol, Multiple Spanning Tree Protocol, and Rapid Per-VLAN Spanning Tree (STP/MSTP/RPVST)** - Spanning tree prevents network loops and provides redundancy in case of a link failure.

### Virtual Router Redundancy Protocol

The Virtual Router Redundancy Protocol (VRRP) includes the following features:

- In traditional networks, edge devices are configured to send packets to a static, default router or gateway.
- If this router becomes unavailable, the edge devices become isolated from the network.
- VRRP uses dynamic failover to ensure the availability of an end device's default router or gateway.
- The IP address used as the default router is assigned to a virtual router (VR).
- The VR includes: an owner router assigned to forward traffic designated to the VR - Master router, and one or more prioritized backup routers. Backup router forwarding traffic for VR replaces owner as Master router.

**Figure 61** Link Level Redundancy VRRP Example




As long as Router A remains available, it operates as the master. If it fails, then Router B takes over as master.

## Link Aggregation Control Protocol

The Link Aggregation Control Protocol (LACP) includes the following features:


- Link Aggregation allows you to assign physical links to one logical link that functions as a single, higher-speed link - a link aggregation group (LAG).
- Link Aggregation Groups are also known as port trunks.
- Redundancy occurs when a LAG link fails and the switch redistributes traffic originally destined to the failed link to the remaining links of the LAG.
- The LAG remains operational as long as there is at least one link in operation.
- The switch uses a hashing algorithm to balance traffic over the LAG.
- The LACP is the standard for Ethernet networks (IEEE 802.3ad).

**Figure 62** *Link Level Redundancy LACP Example 1*



1. Here are two links to each VSF member configured in a LAG to the downstream switch


**Figure 63** *Link Level Redundancy LACP Example 2*



2. When one of the links fail, traffic will failover to the other link

1. Here are two links to each VSF member configured in a LAG to the downstream switch

**Figure 64** *Link Level Redundancy LACP Example 3*



## Distributed Trunking

Distributed Trunking (DT) includes the following features:

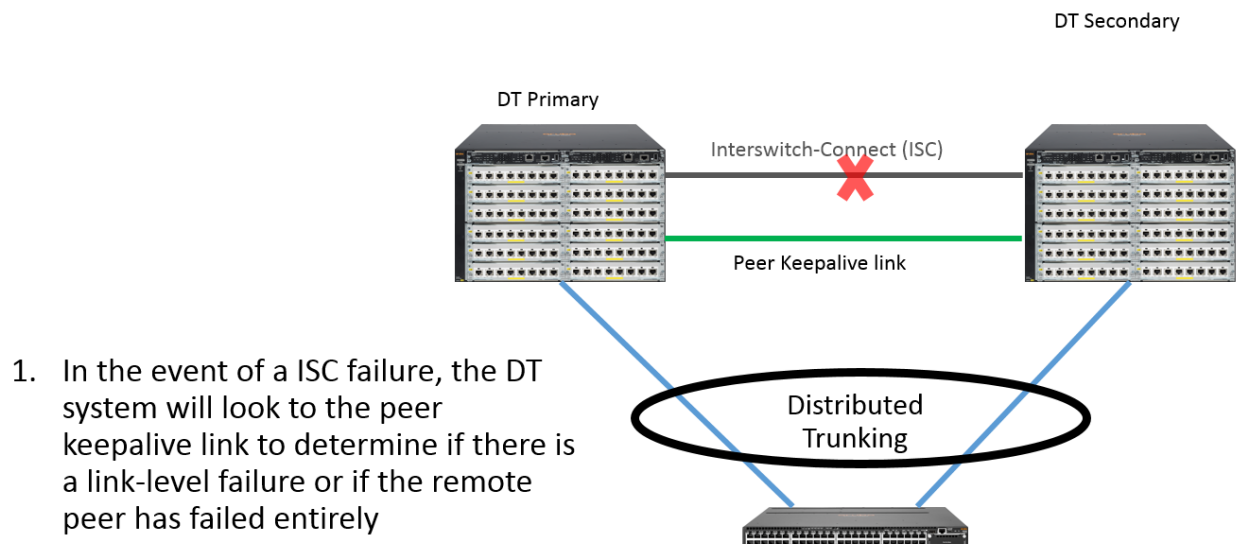
- Proprietary protocol that allows two or more port trunk (LAG) links distributed across two switches to create a trunk group.
- Grouped links appear to the downstream device as if they are from a single device.
- Distributed trunking provides device-level redundancy in addition to link failure protection.
- Each DT switch in a DT pair must be configured with separate Interswitch Connect (ISC) link and peer-keepalive link.
- Peer-keepalive link transmits keepalive messages when the ISC link is down to determine if the failure is a link-level failure or the complete failure of a remote peer.



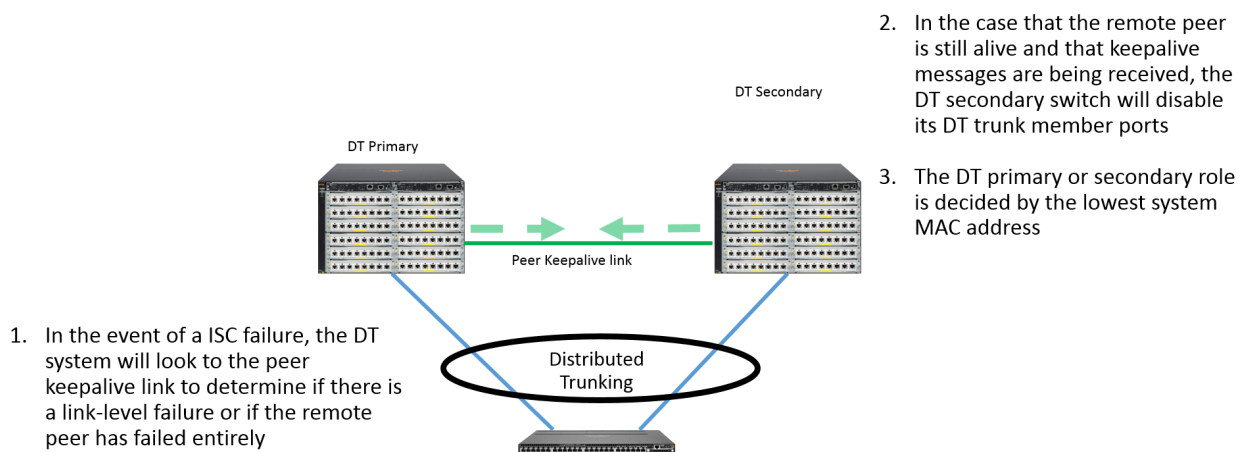
A DT can span a maximum of two switches.

DT is only recommended on older switches that do not support VSF.


**Figure 65** Link Level Redundancy DT Example 1



**Figure 66** Link Level Redundancy DT Example 2 and 3



**Figure 67** Link Level Redundancy DT Example 4




## Spanning Tree Protocol

Spanning Tree Protocol (STP) includes the following features:

- Adding redundant L2 links without a protocol to manage the links results in network loops.
- A STP provides the traditional solution for adding redundant links to a L2 network without causing loops.
- The default STP in ArubaOS-Switch is MSTP. Spanning tree is disabled by default on Aruba switches. Rapid Spanning Tree Protocol (RSTP) is also supported. MSTP supports multiple spanning tree instances.
- STP works by electing a root bridge and calculating the best path to that bridge, blocking alternative paths.
- STP Port Roles include the following:
  - **Designated** - Port closer to root than any other port in link.
  - **Root** - Port is not the closest port to the root on this link. Offers switch the best path to root.
  - **Alternate** - Link offers second or third best path, not connected to same switch.
  - **Backup** - Port is connected to the same switch. The switch has a looped connected to itself, or loop through downstream devices that do not support spanning tree.

**Figure 68** Link Level Redundancy STP Example



Alternate ports are blocked since they do not act as root or designated ports – A loop would be created if an alternate port were enabled – acts as a backup if other links fail

## Advantages

The VRRP, LACP, and STP protocols have the following advantages.

### VRRP Advantages

- Minimizing failover time and bandwidth overhead if a primary router becomes unavailable.
- Minimizing service disruptions during a failover.
- Providing backup for a load-balanced routing solution.

### LACP Advantages

- Aggregates multiple layer 2 connections directly between connected devices.
- Replaces STP.
- Balances traffic across multiple links.
- Simplified management. Configure on the logical interface and not at each physical interface.

### STP Advantages

- Provides link redundancy in the case of a link failure.
- Manages redundant paths by “blocking” unused links to prevent network loops.

## Key Considerations

The VRRP, LACP, and STP protocols have the following key considerations.

### VRRP Key Considerations

- Still uses multiple IP addresses to manage redundancy.
- VSF can replace VRRP with single IP management for multiple switches.

### LACP Key Considerations

- Can be complex to configure and troubleshoot in large-scale networks.
- Lost bandwidth if link fails.

### STP Key Considerations


- Decreases network bandwidth availability by logically blocking redundant physical links to avoid network loops.
- Slow failover to redundant spanning tree links.



# Campus Network Redundancy - Implementation Guidance


Figure 69, Figure 70, and Figure 71 illustrate our recommendations for implementing switches at the component level, device level, and link level.

**Figure 69** Implementation Guidance 1




1. Use redundant power supplies in both the Aruba 5400R and 3810M switches.

**Figure 70** Implementation Guidance 2



1. Use redundant power supplies in both the Aruba 5400R and 3810M switches.
2. At the core/agg. layer, use Aruba 5400R switches in a VSF stack – Use Aruba 3810M and 2920 backplane stacks in the edge or 2930F VSF stacks

**Figure 71** *Implementation Guidance 3*



1. Use redundant power supplies in both the Aruba 5400R and 3810M switches.
2. At the core/agg. layer, use Aruba 5400R switches in a VSF stack – Use Aruba 3810M and 2920 backplane stacks in the edge or 2930F VSF stacks
3. Each mobility controller should be connected by LAGs to each VSF member for redundancy – Each access layer stack should also be connected to each VSF member for redundancy

This chapter includes a review of five common HA deployment models along with the advantages and key considerations of each model:

- [Master / Standby Master with HA Active-Active Locals on page 91](#)
- [One Master - One Local on page 95](#)
- [Independent Masters / All Standalone Masters on page 100](#)
- [Master Redundancy \(Master / Standby Master\) on page 104](#)
- [N+1 \(Over-Subscription\) on page 108](#)

## Master / Standby Master with HA Active-Active Locals



---

Aruba's recommended best practices deployment.

---

This section includes the following topics:

- [Introduction on page 91](#)
- [Configuration Methodology on page 92](#)
- [Failover Scenario on page 93](#)
- [Benefits on page 94](#)
- [Key Considerations on page 94](#)

### Introduction

In a campus environment, full redundancy of the master controllers and AP redundancy via HA AP fast failover should be assured. The fast failover of APs minimizes the disruptions of WLAN clients in case of AP to controller communications failures.

Master redundancy provides non-interruption to these master functions:

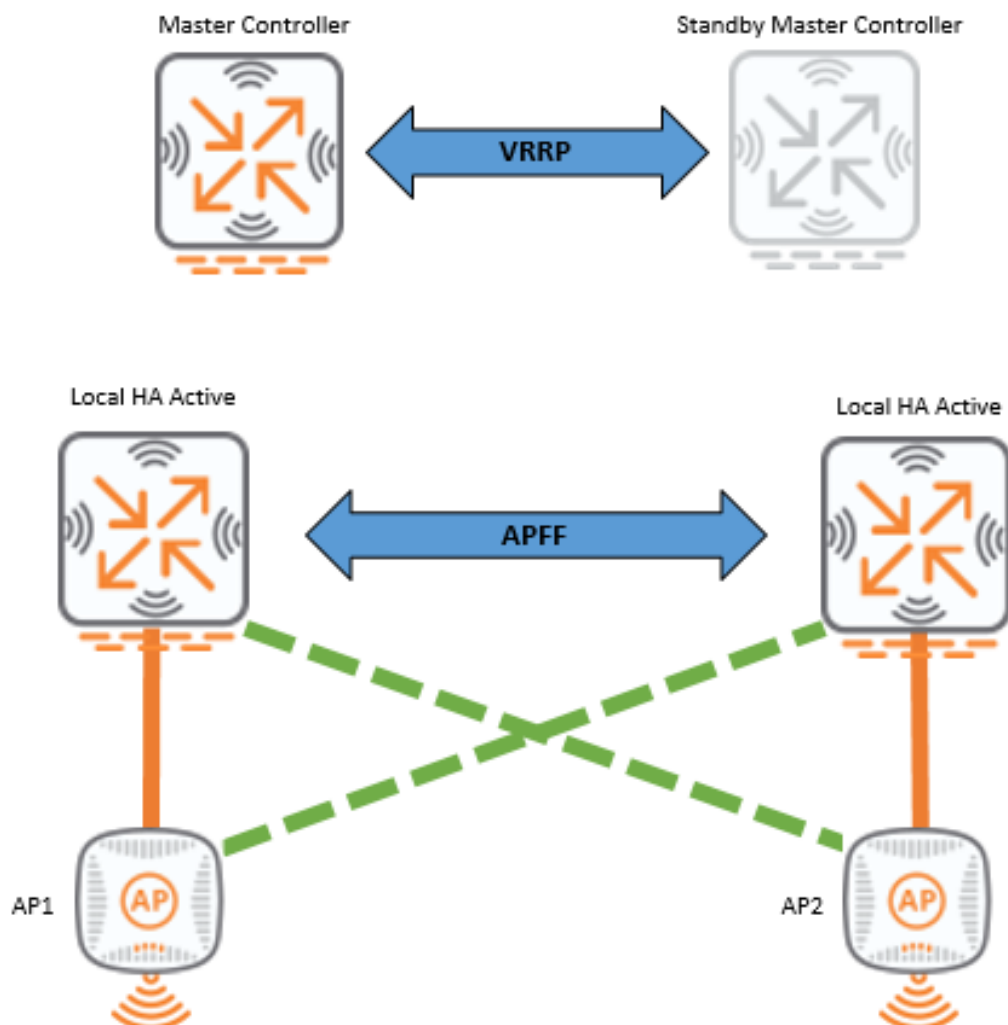
- WLAN configuration
- CPSec trust anchor
- Centralized management and monitoring
- Centralized databases: WMS, whitelists, local-user
- Centralized licensing
- AP master availability

Running HA AP fast failover between each pair of local controllers provides AP redundancy and fast failover in an active-active scenario with client state synchronization for 802.1X users.

## Configuration Methodology

Figure 72 illustrates the deployment model.

**Figure 72** Master / Standby Master with HA Active-Active Locals



This deployment model requires the following configuration guidelines:

- Configure the pair of masters in the master-master redundancy.
- Configure the VIP of that VRRP instance in DHCP Option 43 or DNS for AP master discovery.
- Configure HA AP fast failover between each pair of local controllers.
- Enable the HA features including: inter-controller heartbeat, client state synchronization (for 802.1X users), and HA preemption. The latter option ensures that APs that may have failed over move back to their original LMS controller and restore the AP load balance in the Active-Active scenario.
- Configure the LMS IP per ap-group to set the HA active controller for that group of APs.
- Configure the Backup-LMS IP to account for scenarios where the APs reboot while their LMS controller is unreachable.

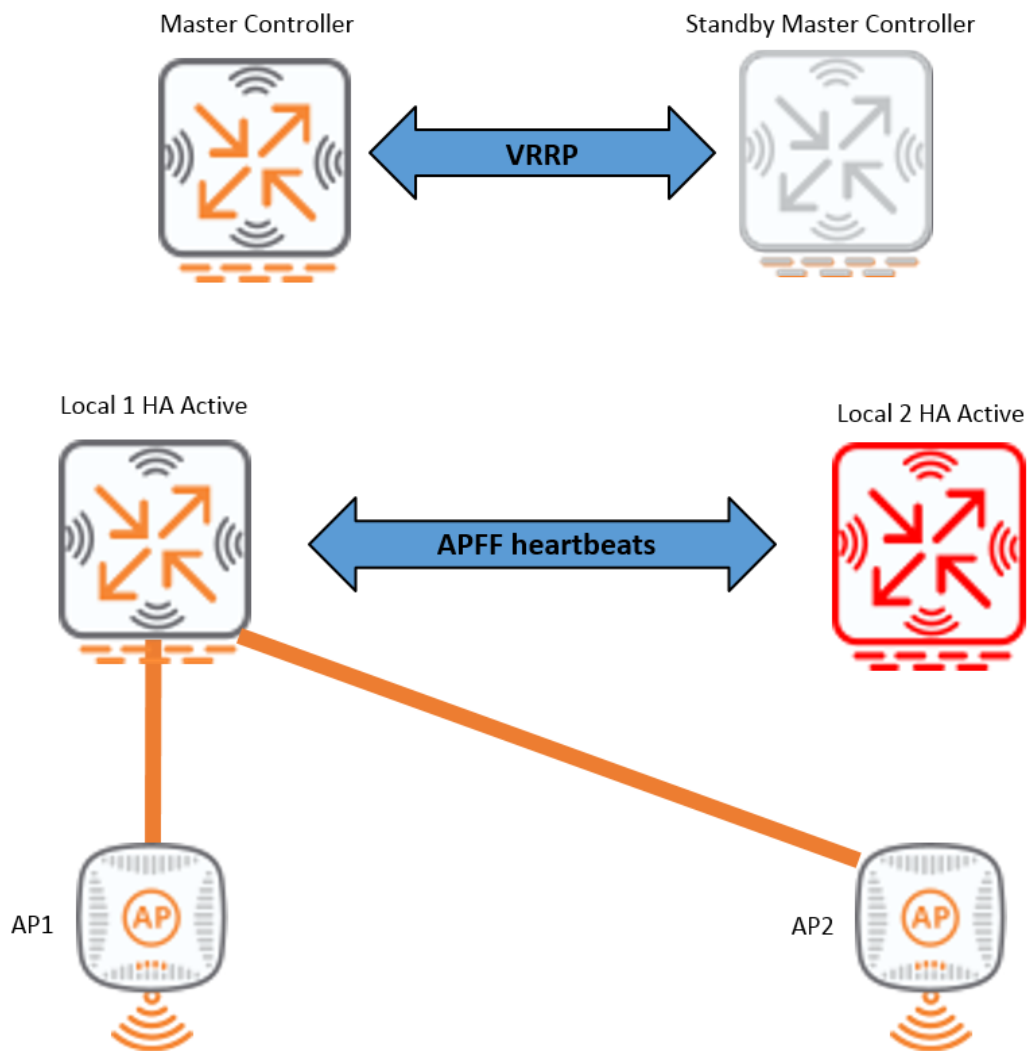


When you enable HA preemption, the parameter ensures that the APs keep monitoring the availability of their failed LMS controller after a failover, and move their active tunnels back to that controller after it is up and the LMS hold-down period timer has expired. This is an option in the AP system profile.

## Failover Scenario

[Figure 73](#) illustrates a failover scenario involving the local controllers.

**Figure 73** *Active-Active Locals with AP Fast Failover Operation*



Following are the steps in the failover scenario:

1. Local 2 fails.
2. Local 1 discovers the Local 2 failure within a sub-second.
3. Local 1 instructs AP2 to failover.
4. AP1 and AP2 tear down their tunnels with Local 2. Their active tunnel is on Local 1.

If HA preemption is enabled and when Local2 is back UP, AP2 active tunnels move back to Local2.

## Benefits

The benefits of this deployment model include the following:

- Redundancy of the master controllers.
- CPSec trust anchor redundancy.
- AP fast failover for each pair of local controllers.
- Active - Active deployment to maximize local controller AP capacity with both locals passing client traffic.
- Benefit of client state synchronization for 802.1X users in case of failover.

## Key Considerations

This deployment model requires a minimum of four Aruba WLAN controllers (two for the pair of masters and at least two for the pair of locals).



---

Master / Standby Master with HA Active-Active Locals is Aruba's recommended best practices deployment.

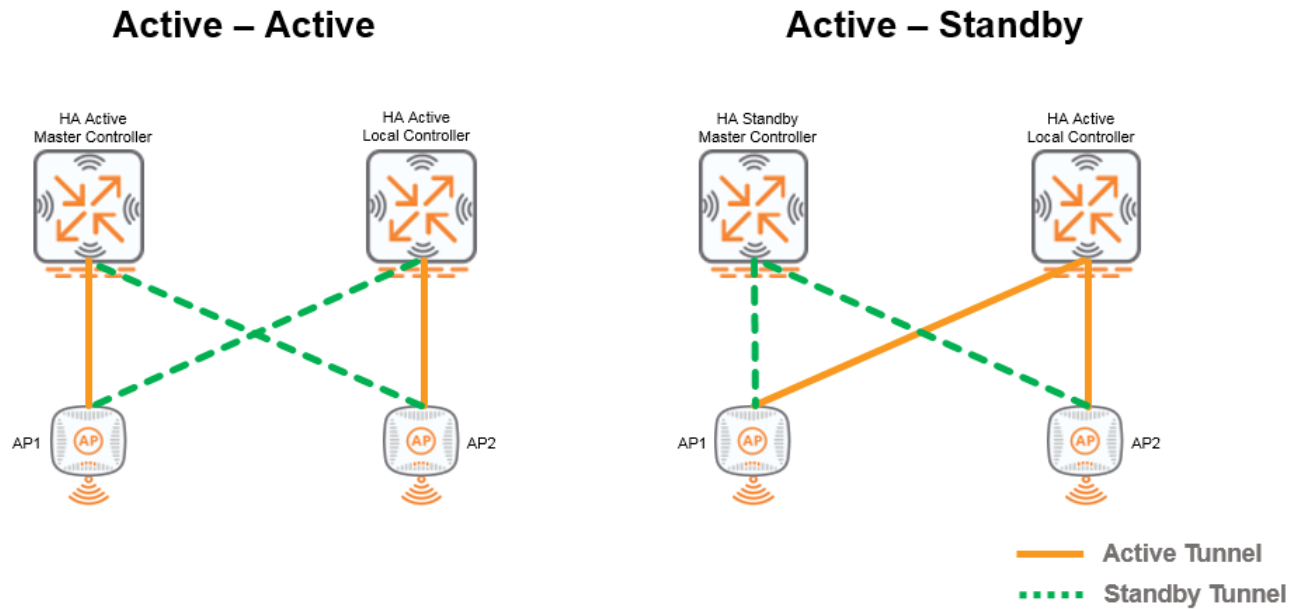
---

## One Master - One Local

There are two possible scenarios for a Master-Local deployment model as illustrated in [Figure 74](#):

- **Active-Active** – Both master and local controllers share the load of APs and clients.
- **Active-Standby** – Only one controller carries the load of APs and clients behind them.

**Figure 74** *Master-Local Deployment Models*




This section includes the following topics:

- [Master-Local \(HA Active-Active\) on page 96](#)
- [Master-Local \(HA Active-Standby\) on page 97](#)
- [Failover Scenario on page 98](#)
- [Benefits on page 99](#)
- [Key Considerations on page 99](#)

## Master-Local (HA Active-Active)

[Figure 75](#) illustrates the deployment model.

**Figure 75** Master-Local (HA Active-Active)



This deployment model requires the following configuration guidelines:


- Configure VRRP between the master and local controllers to ensure the redundancy for AP master discovery. If one of the controllers fail, there is one controller that the APs can go to in case they reboot or new APs come up.
- Configure the VIP in DHCP Option 43 or DNS for the AP master discovery.
- Configure HA AP fast failover with the HA dual role for both master and local controllers.
- Enable the HA features including: inter-controller heartbeat, client state synchronization (for 802.1X users), and HA preemption.
- Configure the LMS IP per ap-group to set the HA active controller (Master for AP1 and Local for AP2).
- Configure the Backup-LMS IP to account for AP reboots while the LMS controller is unreachable (Master for AP2 and Local for AP1).



## Master-Local (HA Active-Standby)

[Figure 76](#) illustrates the deployment model.

**Figure 76** Master-Local (HA Active-Standby)




This deployment model requires the following configuration guidelines:

- Configure VRRP between the master and local controllers to ensure the redundancy for AP master discovery.
- Configure the VIP in DHCP Option 43 or DNS for the AP master discovery.
- Configure HA AP fast failover with the HA dual role for both master and local controllers.
- Enable the HA features including: inter-controller heartbeat, client state synchronization (for 802.1X users), and HA preemption.
- In order to have an HA active and standby, choose one controller to be the LMS controller for all of the APs. Our preference is that the local controller is the LMS controller and HA active for all of the AP groups. All of the AP groups are set to use the local controller IP as the LMS IP, and therefore the local controller is the HA active controller. In this example, AP1 and AP2 active tunnels go to the local controller. AP1 and AP2 standby tunnels go the master controller.
- Configure the Backup-LMS IP to be the master controller in case of double failure.

## Failover Scenario

[Figure 77](#) illustrates a failover scenario.

**Figure 77** Master-Local (HA Active-Standby) AP Fast Failover Operation



Following are the steps in the failover scenario:

1. The local controller fails.
2. The master controller detects the local failure through the inter-controller heartbeat.
3. The master controller instructs AP1 and AP2 to failover. The AP1 and AP2 active tunnels go to the master controller.
4. AP1 and AP2 tear down their tunnels with the local controller.

## Benefits

The benefits of this deployment model include the following:

- Ease of deployment with just one Master and one Local.
- You can still leverage the inter-controller heartbeat for AP fast failover within 0.5 seconds, as opposed to the Master redundancy.
- Ability to use either one of the controllers or both to terminate the APs. There is flexibility to use either the active-active or active-standby scenario.

## Key Considerations

Key considerations of this deployment model include the following:

- With only a single Master controller, there is no master redundancy in case of Master failure and the ability to configure and manage the WLAN network is lost in addition to all the databases residing on the Master.
- Loss of CPSec trust anchor in case of master failure.

## Independent Masters / All Standalone Masters

This section includes the following topics:

- [Prerequisites on page 100](#)
- [Configuration Methodology on page 101](#)
- [Failover Scenario on page 102](#)
- [Benefits on page 103](#)
- [Key Considerations on page 103](#)

### Prerequisites


Independent masters HA redundancy requires the following:

- **Matching Global Configuration Profiles** - The master controllers are independent. Therefore, you configure them separately. Ensure that the global configuration profiles (WLAN configuration, AP groups, RF profiles, VAP profiles, and SSID profiles) match. In case of a failover, the APs need to find the same configuration on the other master controller.
- **Matching User VLANs** - When APs failover, users need to see the same VLANs on the other master controller.
- **Matching CPSec Whitelist** - If you enable CPSec, you need to ensure that APs failing over to another standalone master have their mac addresses in that standalone CPSec whitelist.
- **WMS Offload to AirWave is Recommended** - The WMS database is per independent master, so it may be wise to offload each WMS to AirWave.

## Configuration Methodology

Figure 78 illustrates the deployment model.

**Figure 78** *Independent Masters (HA Active–Active)*



This deployment model requires the following configuration guidelines:

- Configure AP master discovery. If Master 1 and Master 2 are in the same L2 domain, we recommend using VRRP. However, if Master 1 and Master 2 are not in the same L2 domain, then we recommend using the DNS round-robin configuration with two A records.
- Configure HA AP fast failover between the two independent masters.
- Enable the HA features including: inter-controller heartbeat, client state synchronization (for 802.1X users), and HA preemption.
- Configure the LMS IP per ap-group to split the AP load between the two master controllers and to set the HA active controller for that group. In the HA group profile set the controller IP of both masters with the dual role.
- Configure the Backup-LMS IP to account for AP reboots while the LMS controller is unreachable.




The AP system-profiles for both groups (one group terminating on Master 1 and the other group terminating on Master 2) should be identical except for the reversed LMS and Backup-LMS. For AP system-profile 1, the LMS IP is the Master 1 controller IP and the Backup-LMS is the Master 2 controller IP. For AP system-profile 2, the LMS IP is the Master 2 controller IP and the Backup-LMS is the Master 1 controller IP.

## Failover Scenario

[Figure 79](#) illustrates a failover scenario.

**Figure 79** *Independent Masters HA Fast Failover Operation*



Following are the steps in the failover scenario:

1. Master 1 fails.
2. Master 2 detects Master 1 failure within 0.5 seconds.
3. Master 2 instructs AP1 to failover.
4. AP1 tears down its GRE with Master 1 and changes the state of its standby GRE with Master 2 from standby to active. This is how it quickly fails over.
5. AP2 drops its standby tunnel with Master 1.

## Benefits

The benefit of this deployment model is its flexibility. Each master is independent. You can update its code. You can manage it separately.

## Key Considerations

Key considerations of this deployment model include the following:

- Cumbersome to configure, manage, and to meet HA requirements with all of the [Prerequisites on page 100](#) mentioned earlier.
- Potential to trigger rogue detection if APs on each master are within RF reach (when in the same campus). Be careful to mitigate the rogue detections between the two sets of APs.

## Master Redundancy (Master / Standby Master)

This section includes the following topics:

- [Introduction on page 104](#)
- [Configuration Methodology on page 105](#)
- [Failover Scenario on page 106](#)
- [Benefits on page 107](#)
- [Key Considerations on page 107](#)

### Introduction

A pair of controllers are used in this deployment. Rather than deploying them as master-local, some customers like the full redundancy of the master functionality and deploy the master-master redundancy. One master is the active master, and the other is the standby master. This leads to full redundancy of the databases on the master with periodic synchronizations of the databases to the standby master.

In legacy Master Redundancy, no standby tunnels could terminate on the Standby Master. HA support in ArubaOS 6.4 added the ability of the Standby Master to terminate Standby CPSec and/or GRE tunnels from APs with active tunnels to the Active Master. Such feature offers a huge benefit from an AP scalability perspective.


In this topology, the failover is not as fast as other HA deployments that rely on inter-controller heartbeats. Here Master Redundancy rely on VRRP keepalives timing out. Typically, it takes 3 seconds for VRRP to failover. However, the big advantage of using HA over Legacy rests with the pre-established tunnels that offers same failover time with various number of APs. The HA support does also provide for dot1x client state synchronization.



## Configuration Methodology

Figure 80 illustrates the deployment model.

**Figure 80** Master Redundancy (HA Master-Standby)



This deployment model requires the following configuration guidelines:

- Configure the master-redundancy feature and use the VIP in the DHCP option 43 or DNS for AP master discovery.
- Configure HA AP fast failover with the HA dual role for both master and standby master controllers.
- Enable the HA feature client state synchronization (for 802.1X users).




Inter-controller heartbeat and HA preemption are not supported in this topology.

- Configure the LMS IP to be the active master controller-ip. You cannot use the VRRP IP as the LMS IP. Use the active master controller-ip as the LMS IP in order for HA to come up.
- Configure the Backup-LMS IP to be the standby master controller-ip.
- VRRP preemption is recommended for the simple fact that the LMS-IP is set to the primary Active controller, and it will be advantageous for APs booting up to see the LMS controller as the Active Master controller.

## Failover Scenario

Figure 81 illustrates a failover scenario.

**Figure 81** Master Redundancy HA Failover



Following are the steps in the failover scenario:

1. The active master fails.
2. VRRP failover takes place after three seconds when the VRRP keepalives time out.
3. The standby master becomes active.
4. AP1 and AP2 standby tunnels become active. The tunnels are pre-set. The APs terminate on the standby master.

## Benefits

The benefits of this deployment model include the following:

- Master controller is fully redundant.
- Redundant CPSec trust anchor.

## Key Considerations

Key considerations of this deployment model include the following:

- Slower AP failover due to reliance on VRRP failover (three second failover versus a sub-second failover as seen in other scenarios).
- Only one controller can serve APs and Wi-Fi clients in a two-controller deployment. The load is reduced and the other controller is not fully used. However, the benefits of this deployment overcome considerations with this deployment.

## N+1 (Over-Subscription)

This section includes the following topics:

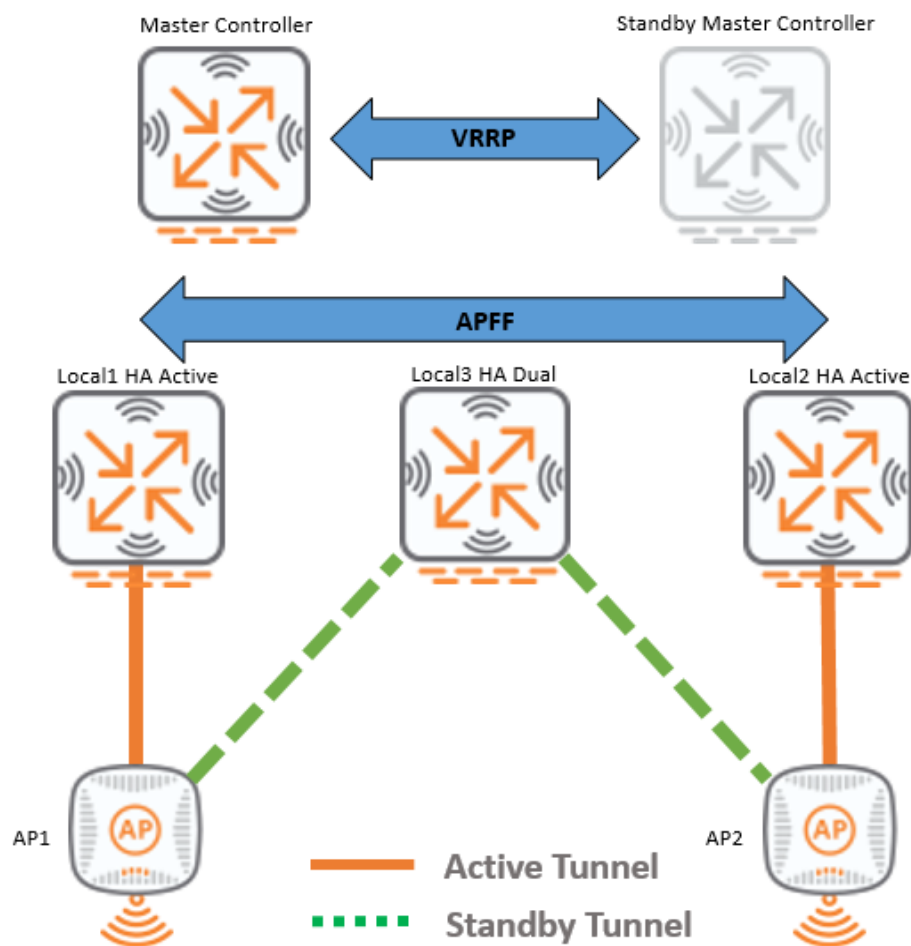
- [Configuration Methodology on page 108](#)
- [Failover Scenario on page 110](#)
- [Benefits on page 111](#)
- [Key Considerations on page 111](#)

### Configuration Methodology

[Figure 82](#) illustrates the deployment model. Three Local controllers attached to a pair of redundant Master controllers. Local1 and Local2 are set to be HA Active, while Local3 HA role could either be dual or standby. However, our recommendation is to set it to HA dual for the following reasons:

- In Standby role, Local3 cannot terminate tunnels carrying user traffic. Therefore, the controller in such role cannot act as a Bkup-LMS controller.
- In Dual role, Local3 will act as the HA Standby for Local1 and Local2, and be used as a Bkup-LMS controller.

**Figure 82** HA with 2+1 Deployment



This deployment model requires the following configuration guidelines:

- Configure the master-master redundancy feature with the VIP in DHCP option 43 or DNS for AP master discovery.
- Configure the HA AP fast failover group-profile with Local1 and Local2 in HA active role, and Local3 in HA dual role.
- Enable the HA features: inter-controller heartbeat and HA preemption.



---

Client state synchronization is not supported.

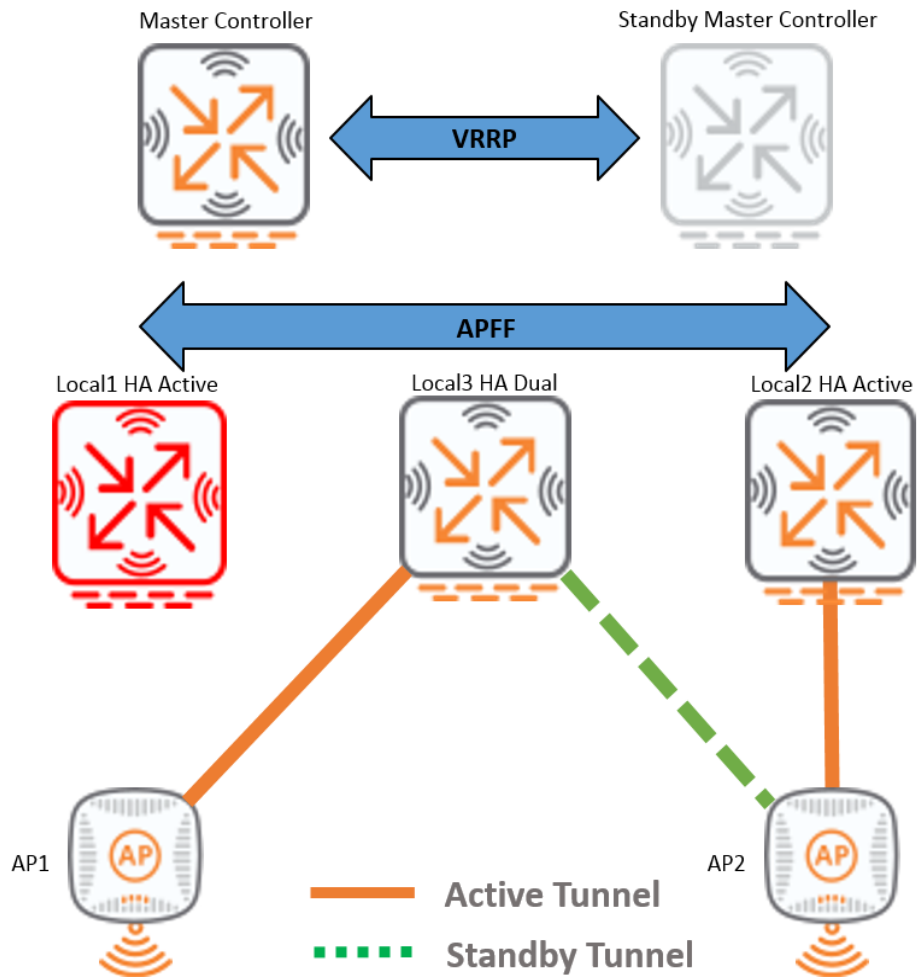
---

- Configure the LMS IP per ap-group to set the HA active controller for that group of APs. Some groups terminate on Local1 as their HA active. Other groups terminate on Local2 as their HA active. That is how we split the load among the multiple active controllers. Local3 is automatically selected as the standby for all deployed APs.
- Configure the Backup-LMS to be Local3 IP to account for AP reboots while the LMS controller is unreachable. Local3 can terminate active tunnels through the LMS and Backup-LMS legacy failover.

## Failover Scenario

Figure 83 illustrates a failover scenario.

**Figure 83** HA with 2+1 Deployment Failover



Following are the steps in the failover scenario:

1. HA active Local1 fails.
2. HA standby Local3 detects Local1 failure within 0.5 seconds.
3. Local3 instructs AP1 to failover.
4. AP1 fails over from Local1 to Local3.

## Benefits

The benefits of this deployment model include the easy migration from existing N+1 legacy redundancy deployments. It is easy to migrate such deployments to HA AP fast failover.

## Key Considerations

Key considerations of this deployment model include the following:

- Centralized licensing is required.
- Client state synchronization for 802.1x users is not supported.
- 70xx platforms are not supported. Only M3, 3600, and 72xx controllers are supported in this deployment model.