

# **Single VLAN Architecture for Wireless LAN**

**Author:**  
**Alap Modi**

**Contributors:**  
**Colin Joseph**  
**Michael Wong**  
**Partha Narasimhan**  
**Peter Thornycroft**  
**Shiv Mehra**



## Copyright Information

© Copyright 2016 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 Company. To obtain such source code, send a check or money order in the amount of US \$10.00 to:

Hewlett-Packard Company

Attn: General Counsel

3000 Hanover Street

Palo Alto, CA 94304

USA

Please specify the product and version for which you are requesting source code. You may also request a copy of this source code free of charge at [dl-gplquery@arubanetworks.com](mailto:dl-gplquery@arubanetworks.com).

---

<b>Contents</b> .....	<b>3</b>
<b>Figures</b> .....	<b>6</b>
<b>Symbols</b> .....	<b>7</b>
<b>About this Guide</b> .....	<b>9</b>
Acronyms .....	9
Scope .....	10
Reference Material .....	11
<b>Introduction to Single VLAN Design</b> .....	<b>12</b>
<b>VLAN Pooling for WLAN</b> .....	<b>13</b>
Use of VLANs .....	13
VLAN Pooling on Enterprise WLAN .....	13
Key Considerations with VLAN Pooling .....	14
802.11 Frames do not Identify VLAN Tags .....	14
IPv6 SLAAC Breaks .....	15
Roaming Complexity .....	16
Inconsistent Utilization of VLANs from VLAN Pool .....	17
<b>Single VLAN Design for WLAN</b> .....	<b>18</b>
Single VLAN Design .....	18
Advantages of Single VLAN Design .....	19
Simple Design and Easy to Support .....	19
IPv6 SLAAC Challenges Solved .....	19
Roaming Becomes Simple .....	19
Address Efficiency .....	19
Key Considerations with Single VLAN Design .....	19
Limiting Large Amounts of Broadcast Multicast Traffic .....	19
Infrastructure to Support Large VLAN .....	20
Routers and Switches .....	20
DHCP Servers .....	20
Firewall Rules .....	20
<b>Aruba's Single VLAN Design Solution</b> .....	<b>21</b>
Optimizing ARP and DHCP Traffic .....	21

Convert Broadcast ARP to Unicast .....	21
Suppress ARP .....	22
Optimize Duplicate Address Detection .....	23
Optimizing Broadcast and Multicast Traffic .....	25
Drop Broadcast and Multicast Traffic .....	25
AirGroup .....	26
IGMP Snooping and Dynamic Multicast Optimization .....	27
Optimizing ICMPv6 Traffic .....	27
Optimizing Router Advertisement and Router Solicitation Messages .....	27
Optimizing Duplicate Address Detection Messages .....	27
Summary of Recommended Settings .....	28
<b>Deployment Guidelines .....</b>	<b>29</b>
Recommended Design .....	29
Guidelines .....	30
Deployment Steps .....	30
Recommendations .....	31
Validation .....	31
<b>Configuration of Optimization Knobs .....</b>	<b>32</b>
Convert Broadcast ARP Request to Unicast .....	32
CLI .....	32
WebUI .....	32
Suppress ARP .....	33
CLI .....	33
WebUI .....	33
Optimize Duplicate Address Detection .....	33
CLI .....	33
WebUI .....	33
Drop Broadcast and Unknown Multicast .....	33
CLI .....	33
WebUI .....	33
AirGroup .....	34
CLI .....	34
WebUI .....	34
IGMP Snooping and DMO .....	35
CLI .....	35

---

WebUI .....	35
<b>Conclusion .....</b>	<b>36</b>

---

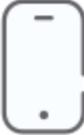
<i>Figure 1 VRD core technologies</i> .....	10
<i>Figure 2 WLAN design without VLAN pooling</i> .....	14
<i>Figure 3 WLAN design with VLAN pooling</i> .....	14
<i>Figure 4 Unicast and broadcast/multicast traffic on WLAN</i> .....	15
<i>Figure 5 IPv6 Stateless Address Auto-Configuration</i> .....	16
<i>Figure 6 Single VLAN design</i> .....	18
<i>Figure 7 Without converting broadcast ARP to unicast</i> .....	22
<i>Figure 8 Converting broadcast ARP to unicast</i> .....	22
<i>Figure 9 Without Suppressing ARP</i> .....	23
<i>Figure 10 With suppress ARP enabled</i> .....	23
<i>Figure 11 Without duplicate address detection</i> .....	24
<i>Figure 12 With duplicate address detection</i> .....	24
<i>Figure 13 Without Drop Broadcast and Multicast Traffic</i> .....	25
<i>Figure 14 With Drop Broadcast and Multicast Traffic</i> .....	26
<i>Figure 15 Large campus with the single VLAN design</i> .....	29
<i>Figure 16 Channel Utilization</i> .....	31
<i>Figure 17 Convert broadcast ARP request to unicast navigation</i> .....	32
<i>Figure 18 Convert broadcast ARP request to unicast checkbox</i> .....	32
<i>Figure 19 Suppress ARP navigation</i> .....	33
<i>Figure 20 Enable Suppress ARP checkbox</i> .....	33
<i>Figure 21 Drop broadcast and unknown multicast navigation</i> .....	33
<i>Figure 22 Drop broadcast and unknown multicast checkbox</i> .....	34
<i>Figure 23 AirGroup</i> .....	34
<i>Figure 24 IGMP navigation</i> .....	35
<i>Figure 25 Enable IGMP checkbox and Enable IGMP Snooping button</i> .....	35
<i>Figure 26 DMO checkbox and DMO Threshold field</i> .....	35

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](#)
- [Scope](#)
- [Reference Material](#)

## Acronyms

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

**Table 2:** *Acronyms*

Acronym	Definition
AP	Access Point
ARP	Address Resolution Protocol
BSSID	Basic Service Set Identifier
BYOD	Bring Your Own Device
DAD	Duplicate Address Detection
DHCP	Dynamic Host Configuration Protocol
DMO	Dynamic Multicast Optimization
GARP	Gratuitous Address Resolution Protocol
GTK	Group Transient Key
ICMP	Internet Control Message Protocol
IGMP	Internet Group Management Protocol
IP	Internet Protocol
LAN	Local Area Network
MAC	Media Access Control
QoS	Quality of Service
RA	Router Advertisement
RF	Radio Frequency
RS	Router Solicitation
SLAAC	Stateless Address Auto Configuration

**Table 2: Acronyms**

Acronym	Definition
SSID	Service Set Identifier
VLAN	Virtual Local Area Network
VRD	Validated Reference Design
WLAN	Wireless Local Area Network

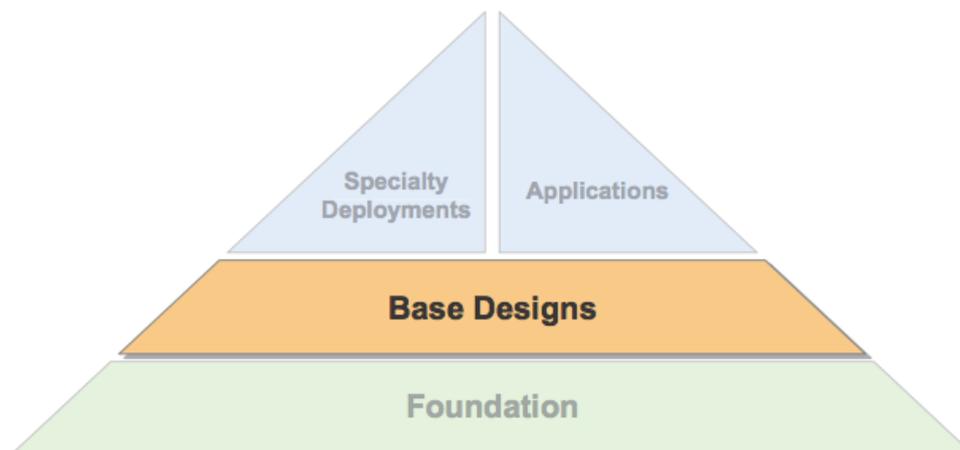
## Scope

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 Design:** 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



**Figure 1** VRD Core Technologies

This Single VLAN 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/>

The growth of Wi-Fi in enterprise Wireless Local Area Network (WLAN) started more than a decade ago. Initially, people used wired networks as a primary medium to connect to networks, and Wi-Fi was just an option. To support a handful of devices connecting to Wi-Fi, network administrators used to assign separate Virtual Local Area Network (VLAN) for wireless clients, which met the users' needs.

With the launch of mobile devices in 2007, the use of Wi-Fi increased exponentially. As the mobile devices did not have Ethernet ports, Wi-Fi became the essential medium for connectivity. To support increased numbers of clients on WLAN, network administrators started adding more VLANs for wireless users, and VLAN pooling became a popular concept. Network administrators created multiple smaller subnets to reduce broadcast multicast traffic rather than increasing the size of the subnet.

Today the requirements for campus and enterprise WLAN have changed quite a bit due to the following:

- Large number of consumer devices connecting to enterprise WLAN.
- Increased use of IPv6 devices on enterprise WLAN.
- Continuous roaming of mobile devices throughout the campus.

There are some challenges supporting these requirements with VLAN pooling. The challenges include: wireless Local Area Network (LAN) design and roaming complexity, controlling a large amount of broadcast multicast traffic, IPv6 Stateless Address Auto Configuration (SLAAC), and inconsistent use of VLANs from the VLAN Pool. Thus, network administrators need to find out a way to overcome these challenges.

At Aruba, we think using a Single VLAN design provides a great alternative to design enterprise Wireless LAN. This document discusses VLAN pooling in an enterprise WLAN, what are the challenges related to it, and how a Single VLAN architecture can help to address them. It also provides design and configuration guidelines to use with Single VLAN architecture, addresses some of the concerns that network administrators may have with the single VLAN design, and finally discusses migration strategy from VLAN Pooling to Single VLAN architecture.

This chapter includes the following topics:

- [Use of VLANs](#)
- [VLAN Pooling on Enterprise WLAN](#)
- [Key Considerations with VLAN Pooling](#)

### Use of VLANs

VLANs are a popular concept in the field of networking and there have been multiple reasons for network administrators to use VLANs. Here are a few common ones:

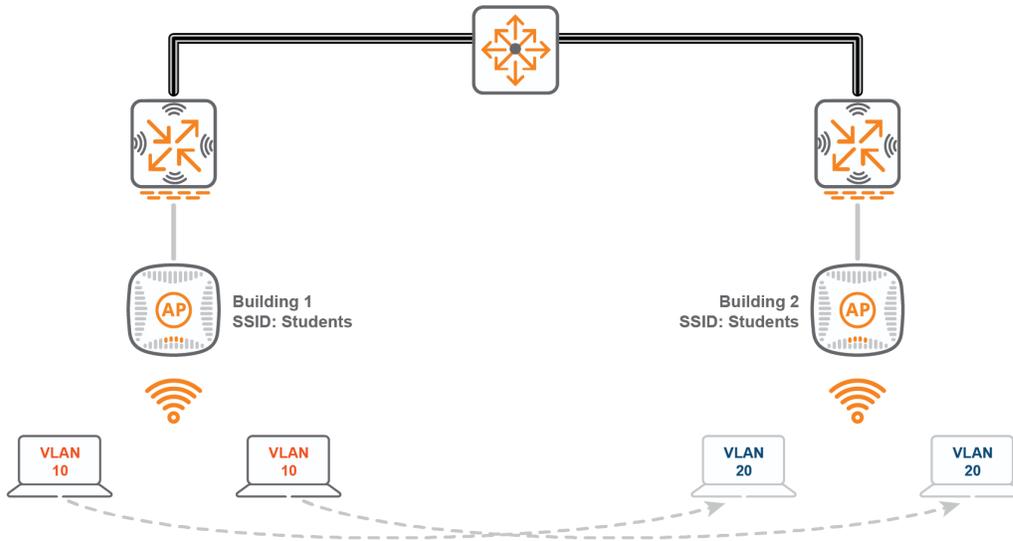
1. Limit broadcast domains by reducing the size of the subnet.
2. Segregate traffic for security - Not a very elegant way to provide security but sometimes that is the best option.
3. Quality of Service (QoS) - Really a way to segregate traffic, so that it can be queued properly even if not tagged.
4. Non-contiguous Internet Protocol (IP) space - Mostly the case where IPv4 public addresses are used for wireless clients.
5. Simpler administration - Looking at the IPv4 address can tell you something about the user, the user's group (employee, contractor, or guest), the location etc. However, IPv6 can be tricky.

### VLAN Pooling on Enterprise WLAN

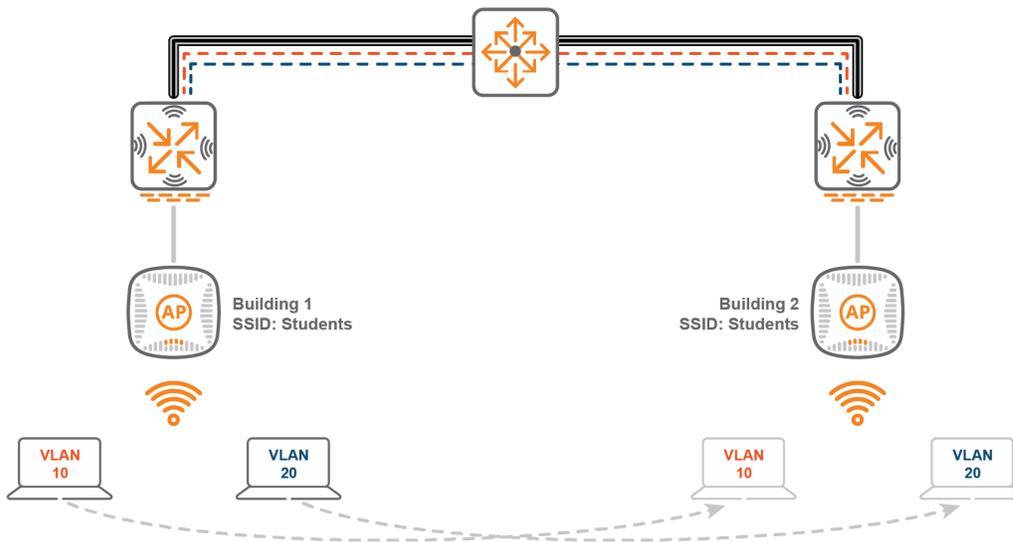
Due to the Bring Your Own Device (BYOD) trend, large numbers of consumer mobile devices are connecting to the enterprise WLAN. Network administrators need to ensure that there are enough IP addresses available to support these clients. On top of that clients are always roaming across the campus and as RF is contiguous throughout the campus, clients are not going to renew their IP address.

The VLAN pooling feature allows the network administrator to assign a “pool” of VLANs to a class of users. This class of users is mainly identified by the Service Set Identifier (SSID) and the location of the access points (APs). For example, as shown in [Figure 3](#), a pool of VLANs consisting of VLANs 10, 20, and 30 can be assigned to all the wireless clients connecting to “Students” SSID in the campus. When a client connects to a network of this configuration, the controller assigns a VLAN to the client from the configured pool. This VLAN derives from a hash algorithm based on the client's MAC address and the number of VLANs in the pool. The hash algorithm ensures that every time the client connects to the network via any AP to any controller, the controller assigns the same VLAN, thus maintaining the IP address of the client as it roams across the APs and controllers.

**Figure 2** WLAN design without VLAN pooling



**Figure 3** WLAN design with VLAN pooling



## Key Considerations with VLAN Pooling

While VLAN Pooling has been a popular feature for enterprise WLANs for many years, there are some challenges with it:

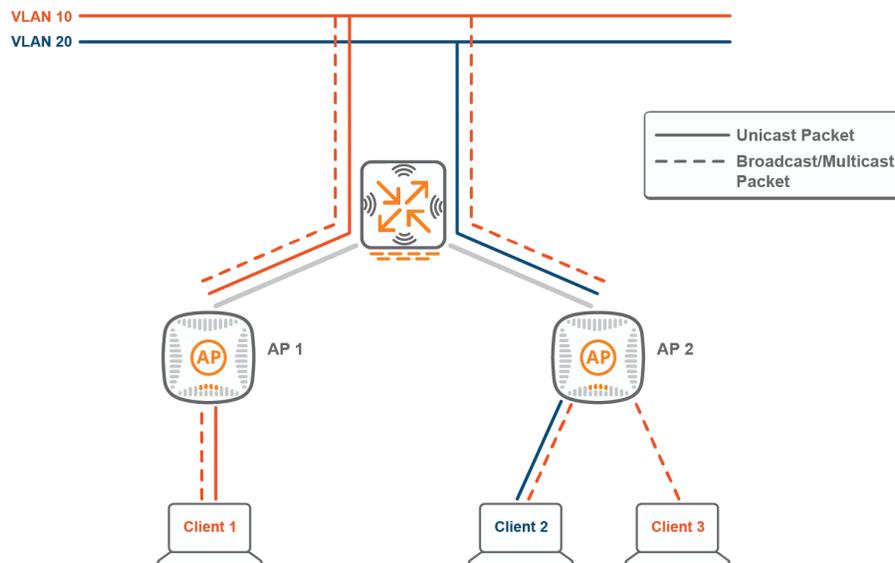
- [802.11 Frames do not Identify VLAN Tags](#)
- [IPv6 SLAAC Breaks](#)
- [Roaming Complexity](#)
- [Inconsistent Utilization of VLANs from VLAN Pool](#)

### 802.11 Frames do not Identify VLAN Tags

One of the reasons VLANs and smaller subnets are used in designing networks to reduce broadcast and multicast traffic. In the case of wired networks, switches and routers understand VLAN tags and thus restrict broadcast

domains to each VLAN only. However, when it comes to treating broadcast-multicast traffic for a WLAN, all the clients connected to a BSSID can hear broadcast-multicast traffic from all the VLANs being used on that SSID, irrespective of client's actual VLAN. Thus, logically BSSID creates a broadcast domain.

**Figure 4** *Unicast and broadcast/multicast traffic on WLAN*



Let us understand it in detail. In the example above, when a unicast packet for Client-1 in VLAN 10 comes to the controller, the controller forwards it to Client-1 via AP1. Over the air, when this frame is being sent to the client, it is sent to the destination MAC address, which is equal to Client-1's MAC address. Similarly, when the unicast packet for Client-2 in VLAN 20 comes to the controller, the controller forwards it to Client-2 via AP2 in a similar manner.

A problem occurs when the multicast or broadcast packet for Client-1 in VLAN 10 comes to the controller. For example, say the ARP packet for Client-1 comes to the controller. Assuming there are no optimizations done, ARP being a broadcast packet, by default the controller will forward it to all the APs where clients in VLAN 10 exists. The ARP packet will go to AP1 and AP2. When the packet comes to AP2, AP sends it out in the air with the destination MAC address as the broadcast MAC (FF:FF:FF:FF:FF:FF). As Client 2 and 3 are connected to the same BSSID, Client 2 and 3 both will receive it, irrespective of which VLANs they are part of.

So why is this concerning? Because, the use of smaller subnets / VLANs:

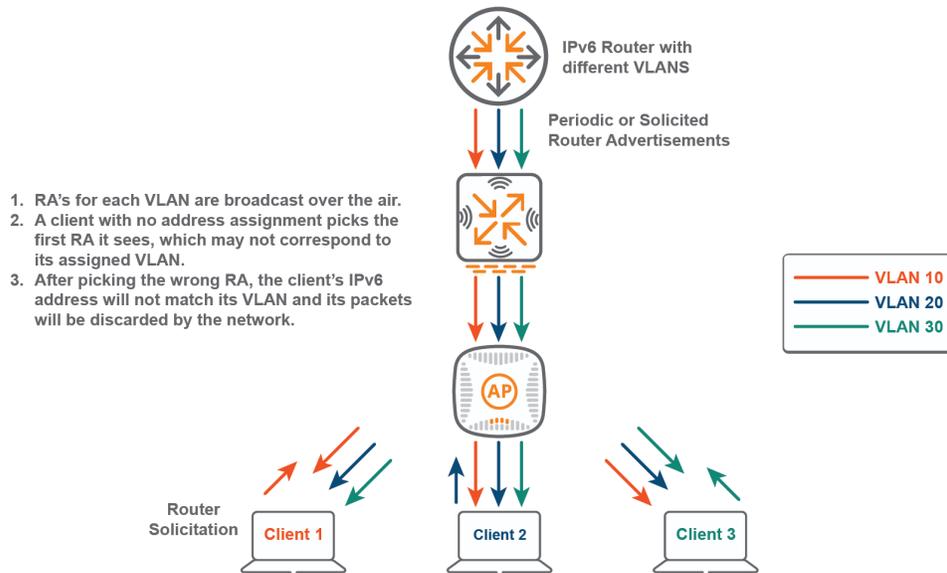
- To restrict broadcast domain does not completely hold true for WLAN as it depends on BSSID/Channel as well.
- To segregate traffic is also not the case as we explained in the example.

## IPv6 SLAAC Breaks

When network administrators started implementing IPv6 with Stateless Address Auto Configuration (SLAAC) on WLAN, they came across a problem where clients were getting IP addresses from a different VLAN than the one that the controller assigned.

A problem occurs when the clients on the same AP belong to different VLANs. This can occur when the WLAN uses VLAN pooling, IP Mobility, or the role-based VLAN derivation rule. (This problem will not occur if all the clients on an AP belong to same VLAN.)

**Figure 5** IPv6 Stateless Address Auto-Configuration



A newly authenticated client is assigned to a VLAN. The network infrastructure knows about the VLAN, but the client does not know it until it gets an IP address from the Dynamic Host Configuration Protocol (DHCP). In IPv6 SLAAC there is no DHCP, so the client does not know its assigned VLAN.

To obtain an IPv6 address, the client must find its parent router and obtain a 64-bit prefix, then add its 64-bit interface ID to form a 128-bit address. To find its router the client may send a broadcast Neighbor Discovery Protocol Router Solicitation (RS) or wait for a periodic Router Advertisement (RA).

Where an AP serves members of more than one IPv6 VLAN, each router configured with a VLAN will advertise RAs, using multicast. Un-assigned clients will respond to the first RA they see since 802.11 frames do not carry any VLAN tag information. If the first router belongs to a client's assigned VLAN, all is well; it will receive an address in that VLAN range. However, if the first RA is for a different VLAN, it will get the wrong address and the network will not route its source address to other destinations, so traffic will be 'black-holed'. On top of that, clients can have multiple IPv6 addresses. If the client hears RAs on multiple VLANs, it will assign multiple IPv6 addresses, where some of them will be invalid.

## Roaming Complexity

In an enterprise environment, based on the size of the campus, there might be multiple controllers to server APs and clients across the campus. If an SSID is available throughout the campus, then the client will not attempt to get a new IP address when it roams from one controller to another. To support seamless roaming of clients throughout the campus, network administrators need to configure VLAN (L2) or IP (L3) mobility.

In the case of VLAN Mobility, client VLANs are extended across all the controllers (as shown in [Figure 3](#)), which requires them to be configured across multiple switches in the path of the controllers and creates a big broadcast domain.

In the case of IP Mobility, different controllers will have different VLANs for wireless clients. However, to avoid clients renewing an IP address, you need to configure IP Mobility, which requires the setup of an IP address or GRE tunnels between all the controllers serving the campus.

Both VLAN and IP Mobility have proven to increase complexity of the design and scalability challenges for large campus networks.

## Inconsistent Utilization of VLANs from VLAN Pool

Most vendors use hash algorithms to assign a VLAN from the VLAN pool to the wireless client. As this hash algorithm is mainly based on the client's MAC address and number of VLANs in the pool, there have been many situations where some VLANs from the VLAN pool are completely exhausted due to a large number of clients in them, while other VLANs are quite empty.



Many vendors use round robin or even algorithms to achieve equal use of all the VLANs in the VLAN Pool. Over time, these algorithms have matured to work well in large campus networks.

Looking at all of the above challenges with VLAN Pooling, here at Aruba, we think that there is an another way to design a next generation of WLAN with just one flat large VLAN.

This chapter includes the following topics:

- [Single VLAN Design](#)
- [Advantages of Single VLAN Design](#)
- [Key Considerations with Single VLAN Design](#)

## Single VLAN Design

The Single VLAN design refers to one large subnet to serve all the clients connecting to an SSID in the campus environment with contiguous RF. Controller and tunneled APs make it possible to scale to larger subnets in the range of /22 to /16.

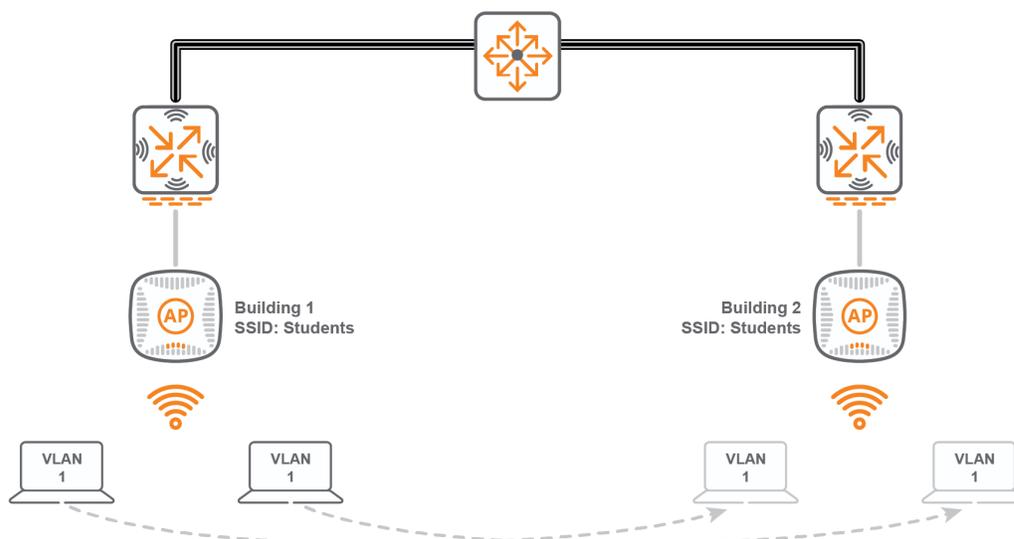
The Single VLAN design is simple and smart. It can greatly reduce the complexity of the WLAN design, and it addresses the IPv6, DHCP, and roaming challenges seen with VLAN pooling. At the same time, you can meet all of the requirements of the wireless LAN using the Single VLAN design. Large universities can use the maximum advantage of this design with their thousands of wireless clients across a large campus with contiguous RF followed by large enterprises with multiple buildings at a location.

The Single VLAN design recommends to use the same VLAN throughout the campus with contiguous RF. If you have multiple buildings in different locations (for example, school districts or corporate branches spread across different cities or towns), then you should use different VLANs and subnets.

The Single VLAN architecture recommends using one large subnet for all the clients connecting to an SSID. However, you can use separate VLANs for clients connecting to separate SSIDs. Ideally, in the campus WLAN, you should use separate VLANs for employee and guest SSIDs.

Lastly, wired and wireless clients should not be sharing the same VLAN. Use separate VLANs for wired clients and, if needed, use multiple smaller subnets to restrict broadcast domain for wired devices. The Single VLAN architecture is for wireless LAN only, as the controller has a lot of visibility and control over wireless users, but none for wired devices. The obvious problems related to large VLANs on wired networks still apply.

**Figure 6** Single VLAN design



## Advantages of Single VLAN Design

Reasons for using a single VLAN architecture include the following.

### Simple Design and Easy to Support

With just one large subnet for all the wireless clients connecting to an SSID, the network administrator does not need to configure multiple VLANs, DHCP scopes, extend VLANs across multiple devices, and configure redundancy for default gateways in each VLAN.

On top of that, the network administrator does not need to configure VLAN mobility or IP mobility for clients roaming across different controllers. As the client VLAN is going to be the same, the client keeps using the same IP address.

As wireless LAN design becomes simple, it is very easy for network administrators to support this design. If at the same time some issues come up, it is very easy to troubleshoot it as well. Network administrators do not need to figure out which VLAN the client falls into, try to isolate an issue related to just that VLAN, or network wide, and so on.

### IPv6 SLAAC Challenges Solved

With the single VLAN design, there will not be multiple RAs coming from IPv6 routers. Thus, there is no issue of the client getting the wrong v6 address.

### Roaming Becomes Simple

No matter where the client roams, either on a different AP on the same controller or on a different controller, the client does not need to change its IP address. The network administrator does not need to configure Layer 2 or Layer 3 roaming on the controller. This solves the roaming complexity for the network administrator.

### Address Efficiency

As there will be just one VLAN, there is no issue of some VLANs getting completely exhausted while other VLANs staying less utilized. In large campus environments with thousands of users, this helps to avoid the issue of clients not getting IP addresses due to some VLANs being completely exhausted.

## Key Considerations with Single VLAN Design

Many things come to a network administrator's mind before using a large subnet. Some of the most common questions that come up are:

### Limiting Large Amounts of Broadcast Multicast Traffic

This is the first thing that comes to the network administrator's mind when thinking about using the single VLAN design. As we discussed in section 3.3.1, in the case of WLAN, there is not a very big difference between VLAN Pooling and the Single VLAN design in terms of the amount of broadcast and multicast traffic.

In the case of the Aruba controller and tunneled APs, the controller is a central point who knows about all the associated clients, their IP addresses, and MAC addresses, to which AP they are connected, what role they are in, and what type of traffic they are allowed. Thus, the controller can make intelligent decisions on which packet needs to go to which clients to reduce unnecessary traffic in the air.

To do that ArubaOS has features and built-in intelligence that can reduce the amount of broadcast and multicast traffic. Chapter 5 in this document covers all the features available in ArubaOS to intelligently optimize broadcast and multicast traffic.

## Infrastructure to Support Large VLAN

To support a large VLAN, other devices on the network also need to be capable. Some of the most common concerns are related to:

### Routers and Switches

Most of the time the wireless client VLAN is configured as Layer 2 on the Aruba controller and uplink switch. The uplink router is configured as Layer 3 interface working as default gateway for that VLAN. While using large VLANs, the ARP table on the router can be very large depending on how many clients are connected to the WLAN. Thus, network administrators should find out any limitations on the router and address them accordingly.

The same goes for switches. Although the controller uplink switch does not need to handle ARP entries, the bridge/MAC address table entries can be very large. The network administrator should consider this.

### DHCP Servers

In the case of single VLAN design, the DHCP server needs to be able to support large DHCP scopes, as big as /16. Many Linux based DHCP servers are not capable of this. Thus, the network administrator needs to ensure that the DHCP server can handle a large DHCP scope.

At the same time, the DHCP server needs to be powerful enough to lease out IP addresses quickly. During a failover situation, it becomes essential.

### Firewall Rules

For most of the enterprise networks, existing firewall rules are based on IP subnets. There are many network administrators still entrenched in this way of thinking. Even inside the network, routing, access control, and QoS are done based on the VLANs and IP subnet through routers and legacy firewalls. Obviously next generation firewalls are user and application centric and are reasonably deployed at the Internet gateway and data centers for large enterprise customers.

With the Single VLAN Design, all the traffic of wireless clients is in one subnet only. Hence user and application centric firewalls are needed to apply proper firewall rules. One more option is to use the Aruba Controller's built in stateful firewall to manage firewall policies for wireless clients.

The biggest concern that network administrators may have with the single VLAN design is about how to restrict broadcast and multicast traffic. In the case of a large subnet with thousands of users, there will be a large amount of broadcast multicast traffic. This broadcast and multicast traffic can affect client performance if not restricted properly by the wireless controller. At the same time, not all the broadcast and multicast traffic should be dropped by the wireless controller or network infrastructure, otherwise some of the basic functionalities like VRRP, IPv6, DHCP, ARP, MDNS, etc. may break. In the case of Aruba controllers, tunneled APs, and software intelligence built into ArubaOS, these make it possible to deploy the single VLAN design. In this chapter, we will look into different knobs available in ArubaOS to limit broadcast and multicast traffic.

This chapter includes the following topics:

- [Optimizing ARP and DHCP Traffic](#)
- [Optimizing Broadcast and Multicast Traffic](#)
- [Optimizing ICMPv6 Traffic](#)
- [Summary of Recommended Settings](#)

### Optimizing ARP and DHCP Traffic

When using large subnets ARP and DHCP traffic can significantly impact the performance of a WLAN. A single ARP packet in a VLAN will be flooded to all the AP tunnels where that VLAN exists. Thinking in a real world scenario, if you have 1000 APs on the controller, and if an ARP packet comes to the controller (from wired side or from one of the wireless clients) on a wireless client VLAN, then that ARP packet will be flooded to all 1000 AP tunnels and sent out in the air. This can affect the performance of wireless clients and thus the broadcast filter ARP knob should be used.

A similar behavior will be seen for DHCP packets as well.

Gratuitous ARP can also impact the WLAN performance. Whenever a client sends a GARP frame, that frame goes to all other wired and wireless clients. Each client receiving this GARP frame will update its ARP cache entry for it. In the network where there are thousands of clients sending GARP, we noticed problems with mobile devices and laptops where they ran out of their ARP table limits as they do not support too big of a size of ARP table. Mostly clients support ARP table entries in the range varying from 512 to 4000 entries or less. Thus after a certain point, clients started losing default-gateway entries and thus showing connectivity issues.

The Aruba controller has features that can help to reduce the amount of ARP and DHCP packets flooded to wireless clients.

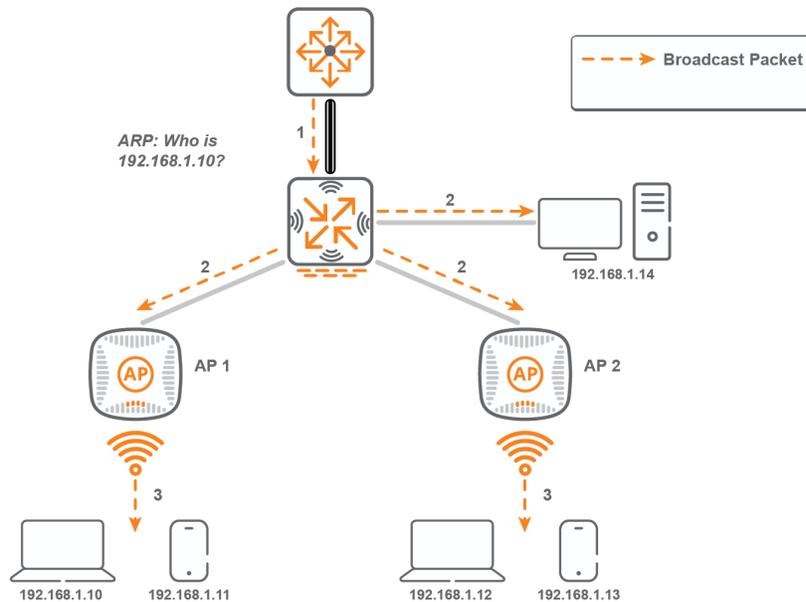
### Convert Broadcast ARP to Unicast

**Function:** Converts broadcast ARP and DHCP packets to unicast before sending to wireless clients.

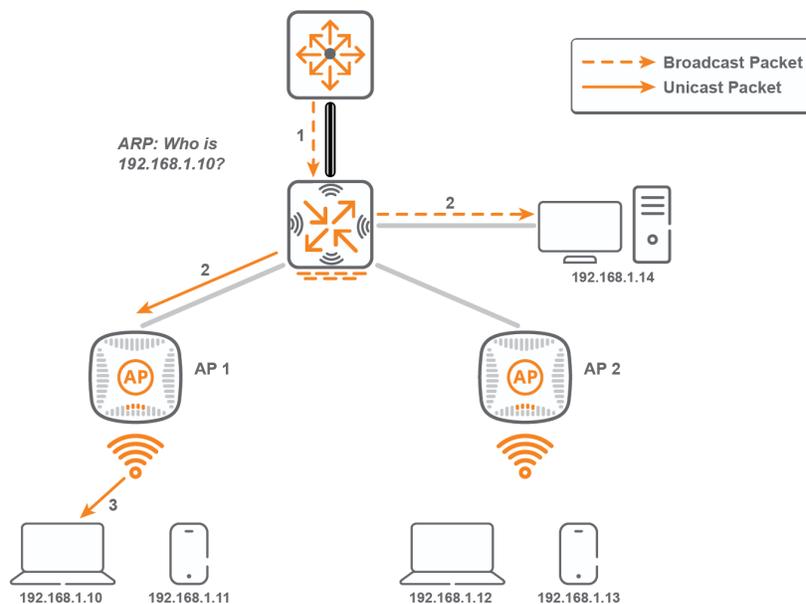
**Default:** Turned ON

**Recommended:** Turned ON

**Figure 7** Without converting broadcast ARP to unicast



**Figure 8** Converting broadcast ARP to unicast



### Details

This knob will enable ARP conversion on all of the VLANs. If this knob is enabled, all the broadcast ARP destined to wireless clients, that are part of the user-table, are converted to a unicast ARP request and sent to a particular AP where that client is associated.

Apart from ARP, it takes care of DHCP packets as well in the same manner.

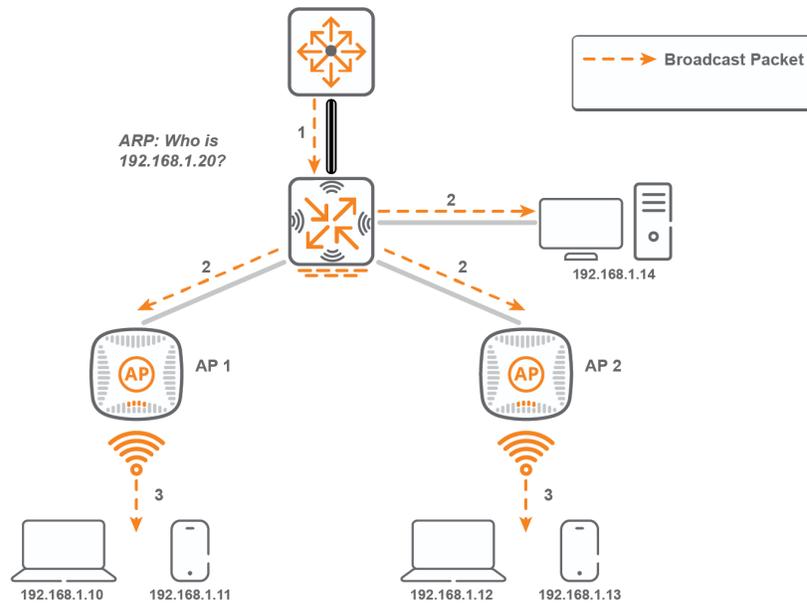
### Suppress ARP

**Function:** Stops flooding of unknown ARP request to wireless clients.

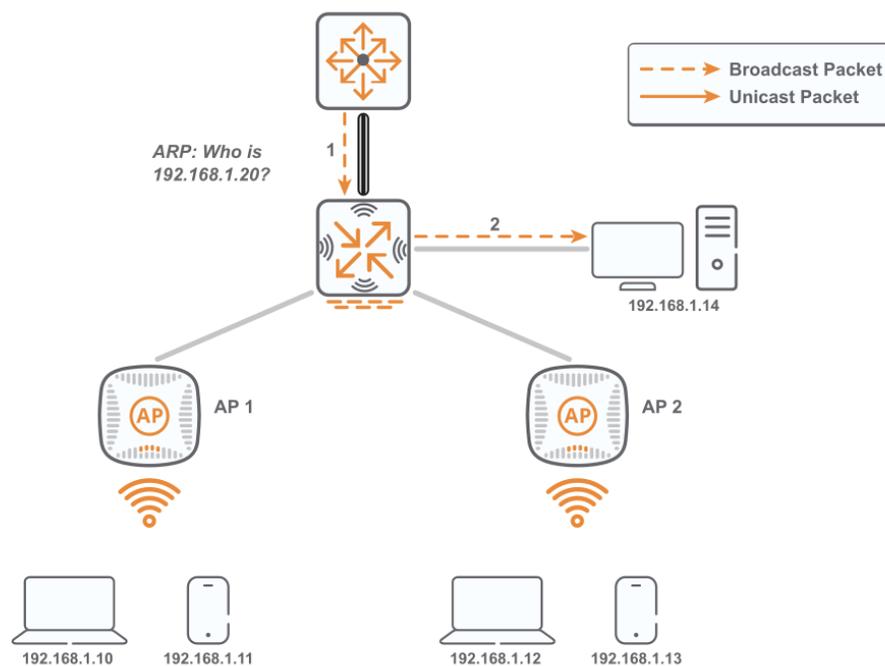
**Default:** Turned ON

**Recommended:** Turned ON

**Figure 9** Without Suppressing ARP



**Figure 10** With suppress ARP enabled



### Details

This knob will stop flooding of “unknown” ARP requests to AP tunnels (tunnel and d-tunnel mode only). The unknown ARP request will still be flooded out of LAN ports, wired AP ports, and split-tunnel VAPs.

### Optimize Duplicate Address Detection

**Function:** Controls the flooding of Gratuitous ARP and IPv6 Duplicate Address Detection frames to wireless clients.

**Default:** Turned ON

## Recommended: Turned ON

Figure 11 Without duplicate address detection

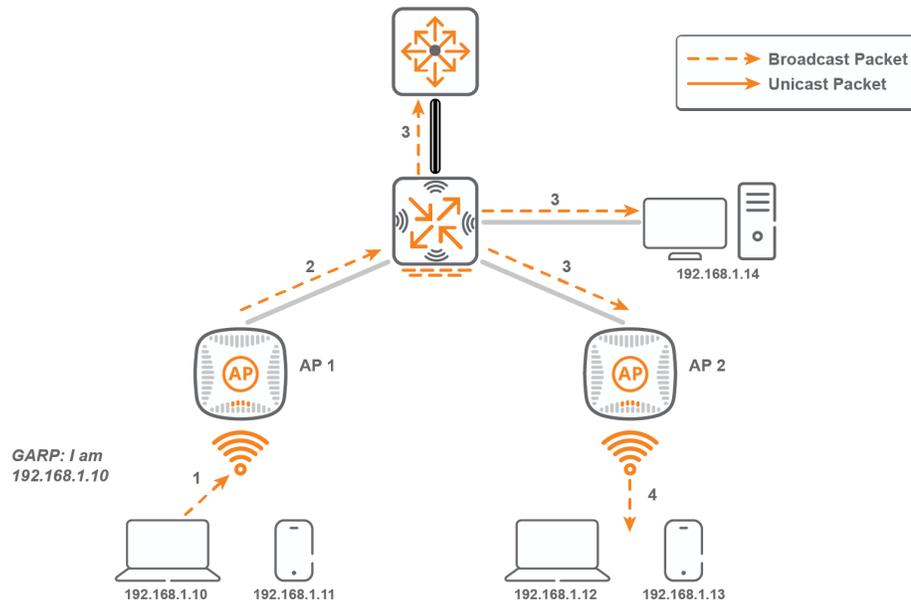
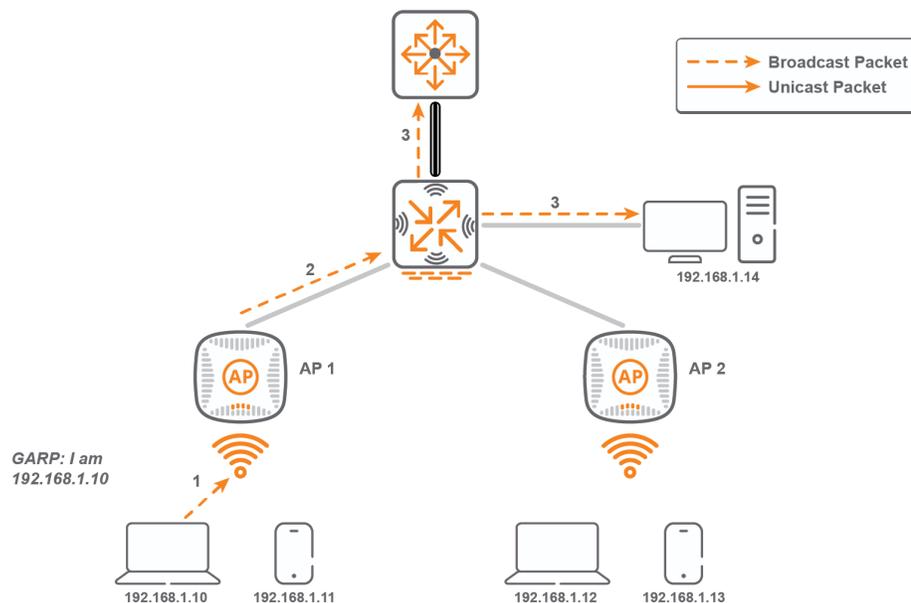


Figure 12 With duplicate address detection



## Details

Broadcast Filter ARP and Suppress ARP features take care of known and unknown ARP frames but not GARP. The “Optimize Duplicate Address Detection” feature stops forwarding GARP and IPv6 DAD frames to wireless clients. However, there is an exception to this. If the GARP or DAD frame is coming from the router (default gateway), they will still send it to all the wireless clients as that is very important information for clients. This especially helps when customers have default gateway redundancy. If the active gateway goes down, then the standby gateway will start owning the gateway IP and send out the GARP frame to update clients about the new IP to Media Access Control (MAC) address mapping.

## Optimizing Broadcast and Multicast Traffic

Once ARP and DHCP packets are converted to unicast, the next step is to restrict broadcast and multicast that might be generated as a part of some applications running on the client devices. Most common applications use either NetBIOS, MDNS or DLNA based services, which are multicast based. The Aruba controller has features that can completely drop all of these broadcast multicast traffic or allow specific types of traffic by converting it to unicast.

Many companies or universities have multicast video streams running in their network for business and educational purpose. There might be other custom applications, which run on multicast. To meet these requirements, the Aruba controller provides knobs to optimize multicast traffic over the air.

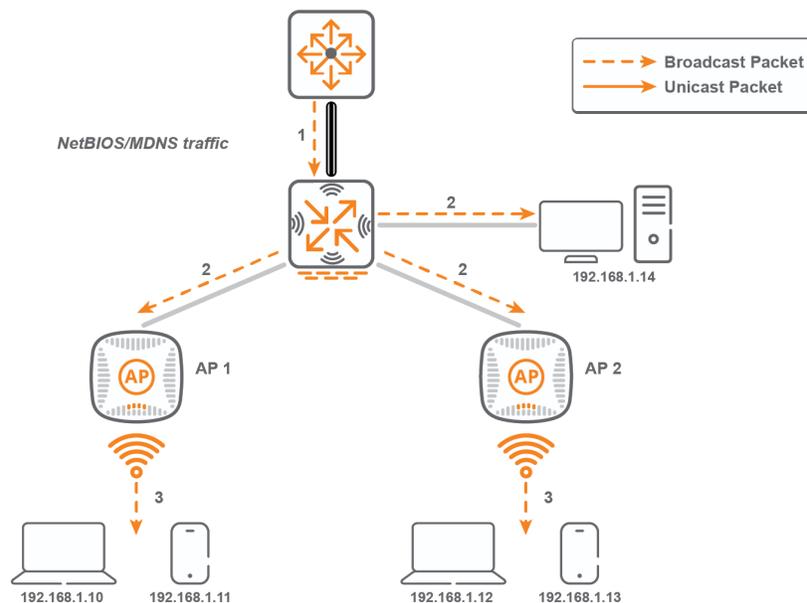
### Drop Broadcast and Multicast Traffic

**Function:** Stops flooding of broadcast and multicast traffic to wireless clients.

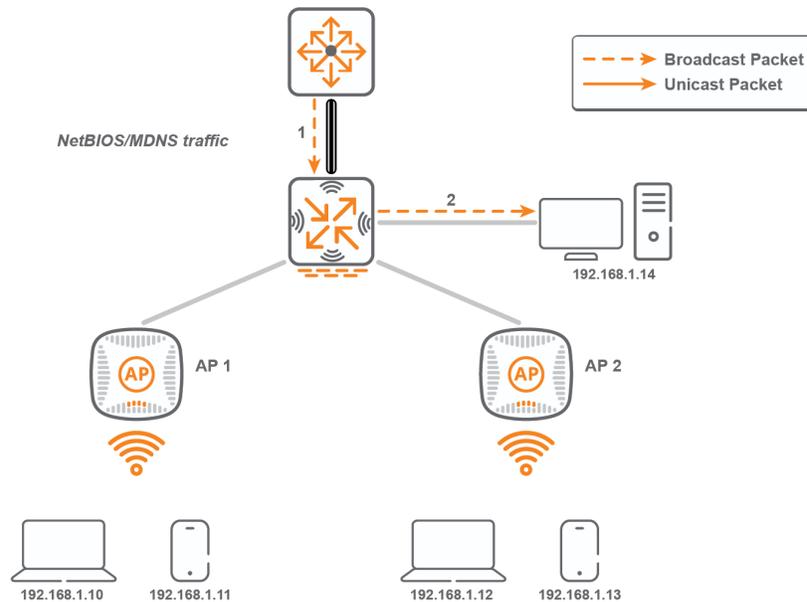
**Default:** Turned OFF

**Recommended:** Turned ON

**Figure 13** Without Drop Broadcast and Multicast Traffic



**Figure 14** With Drop Broadcast and Multicast Traffic



### Details

ArubaOS has “Drop Broadcast and Multicast traffic” knob in Virtual-AP profile. When enabled, all of the broadcast and multicast traffic on the WLAN will be dropped.



Drop Broadcast Multicast knob will drop broadcast ARP and DHCP packets if “Convert broadcast ARP to unicast” knob is not enabled.



Even if a multicast based application needs to work on the WLAN, it is safe to enable the “Drop broadcast Multicast” knob. The upcoming sections give instructions on how to permit the known multicast traffic even when this knob is enabled.

### AirGroup

**Function:** Allows mdns and upnp traffic by converting them to unicast.

**Default:** Turned Off

**Recommended:** Turn ON to allow mdns and universal plug and play traffic.

### Details

While we drop the broadcast-multicast traffic in the air, in enterprise or campus deployments DLNA, MDNS, and other zero-configuration services are essential. To allow such services for airplay and Chromecast kind of applications, Aruba WLAN with AirGroup technology enables context aware access to DLNA, Apple Bonjour, and other shared devices without constraining WLAN performance.

To understand in detail how the AirGroup feature works, please review:

- Aruba User Guide  
<https://support.arubanetworks.com/Documentation/tabid/77/DMXModule/512/Default.aspx?EntryId=8863>
- AirGroup Deployment Guide  
<http://community.arubanetworks.com/aruba/attachments/aruba/unified-wired-wireless-access/15478/1/ArubaAirGroup-6136-DG.pdf>

## IGMP Snooping and Dynamic Multicast Optimization

**Function:** Keeps track of clients subscribed to multicast stream and converts multicast packets to unicast before sending it to wireless clients.

**Default:** Disabled (both IGMP snooping and DMO)

**Recommended:** Enable if multicast streaming needs to be allowed on Wi-Fi

### Details

Dynamic Multicast Optimization (DMO) knob converts multicast packets to unicast and transmits it at higher unicast rate over the air.



IGMP snooping or IGMP proxy needs to be enabled for DMO to work.



Multicast stream should be prioritized by configuring uplink ACL and correct WMM parameter to match DSCP values.

## Optimizing ICMPv6 Traffic

In the case of IPv6 (Internet Control Message Protocol (ICMP)), RA, RS, and DAD are the main control messages that go to the IPv6 multicast address. If they are not optimized properly, then it can affect client performance the same way as IPv4 multicast traffic.

### Optimizing Router Advertisement and Router Solicitation Messages

In the case of IPv6, clients can auto configure the IPv6 address either by listening to periodic RA messages or by sending out RS messages when they come up and then in turn getting RA from the IPv6 router. All these RA and RS messages are sent to specific multicast addresses, so that all the routers and clients can hear them. However, at the same time it can affect the wireless client performance.

In the case of the Aruba controller, if all the wireless clients fall into the same VLAN (Single VLAN design), then periodic RAs are sent as it is - multicast packets as impact on the client performance is not too much. However, if VLAN pooling is used, then depending on the number of VLANs in the pool, periodic RAs can flood the network. To avoid that, the controller intercepts RAs, converts them to unicast packets, and sends them to all the clients in that VLAN.

For the RS message, even though it goes on a multicast address, the controller does not put that frame into AP tunnels; thus, they do not go to wireless clients. When the response (RA) comes back from the default gateway, the controller just forwards that packet to that particular client that sent out the RS message.

Thus, ArubaOS is by default optimized to handle RA and RS messages without enabling any special features.

### Optimizing Duplicate Address Detection Messages

The purpose of DAD packets are the same as GARP packets in the case of IPv4. Whenever the client gets an IP address, either by DHCPv6 or by SLAAC, it sends out DAD packets to inform others about its IP address. As DAD packets are sent on the IPv6 multicast address, all the clients on that VLAN can hear them. Especially in the case of the Single VLAN design, if each and every client sends out a DAD packet when it gets the IPv6 address, and if all other clients in that VLAN listen to it, then clients will need to handle a large number of ARP tables. There is a limitation to how many ARP table entries the client can handle. This depends on the type of client. It is usually in the range of 512 to 4000 entries. If you have more than 4000 clients on your campus and if all of them send DAD packets, then a situation can occur, where the client runs out of its ARP cache entries and thus faces intermittent connectivity issues.

To address that the Aruba controller has the “Optimize Duplicate Address Detection” feature, which works as described in [Optimize Duplicate Address Detection](#)..

## Summary of Recommended Settings

[Table 3](#) summarizes the recommended settings.

**Table 3:** *Summary of Recommended Settings*

Feature / Setting	Default Setting	Recommended Setting	Note
Convert broadcast ARP to Unicast	ON	ON	Converts broadcast ARP and DHCP packets to Unicast before forwarding to wireless clients
Suppress ARP	ON	ON	Stops flooding of unknown ARP request to wireless clients
Optimize Duplicate Address Detection	ON	ON	Controls the flooding of Gratuitous ARP and IPv6 Duplicate Address Detection frames to wireless clients
Drop Broadcast and Multicast traffic	OFF	ON	Stops flooding of broadcast and multicast traffic to wireless clients
AirGroup	OFF	ON*	* Turn it on to allow MDNS and UPnP traffic
IGMP Snooping and Dynamic Multicast Optimization	OFF	ON*	* Turn it ON to allow multicast streaming on WLAN

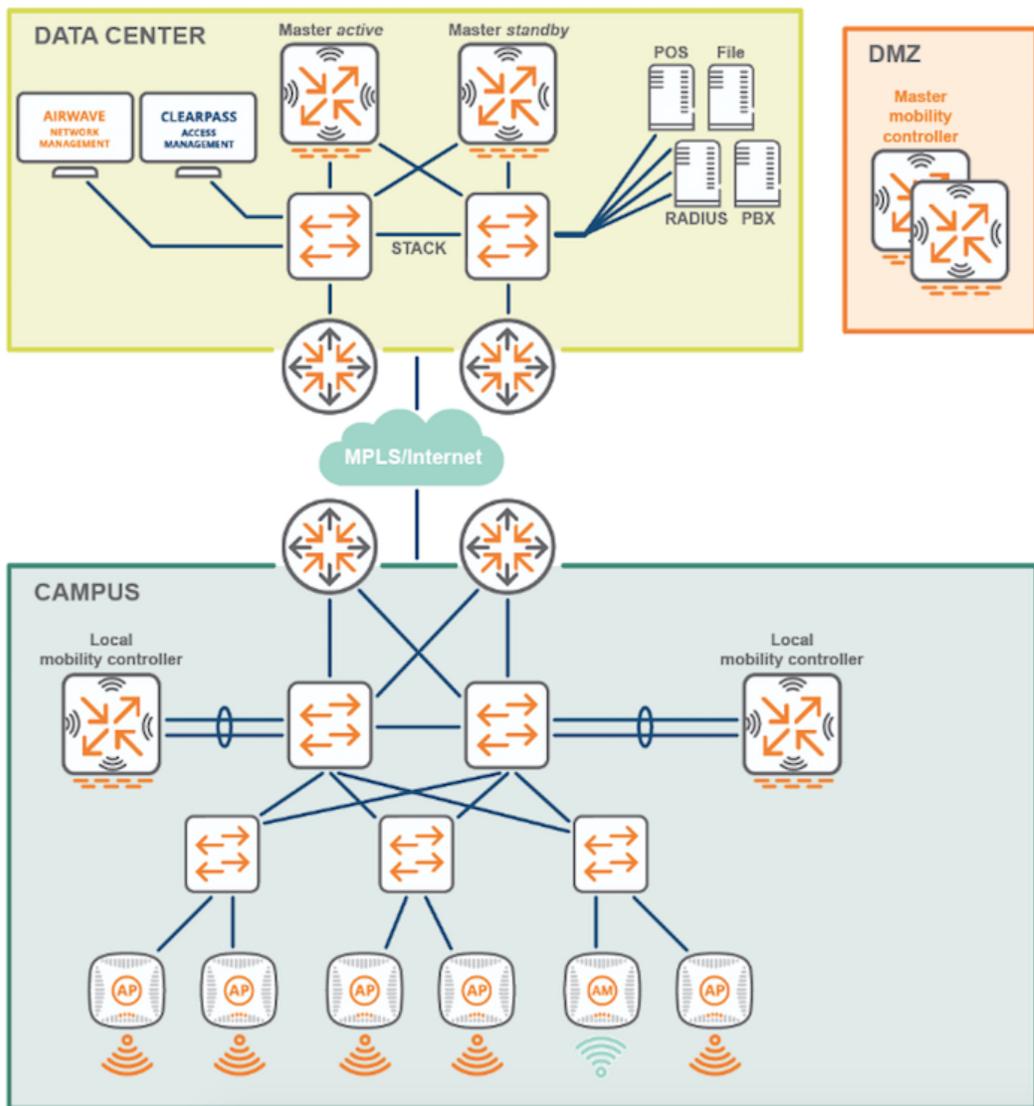
This chapter includes the following topics:

- [Recommended Design](#)
- [Guidelines](#)
- [Deployment Steps](#)
- [Recommendations](#)
- [Validation](#)

## Recommended Design

A typical large campus with the single VLAN design will look like [Figure 15](#) below.

**Figure 15** Large campus with the single VLAN design



As shown in the diagram:

- APs are connected at the access layer to the access switch. As all of the client traffic is tunneled between the controller and the AP, the client VLAN is not needed on the access switch.
- Local controllers are connected to the distribution switch in large campuses, as they will have multiple controllers serving different parts of the campus.
- The master controller is typically located in the data center, connected to the core switch.
- The client VLAN is Layer 2 to the Aruba controller and distribution switches. The core switch or firewall will have a Layer 3 interface for the client VLAN.

## Guidelines

1. The single VLAN design is intended for wireless clients connecting APs in the same campus with contiguous RF. If RF is not contiguous, then it is completely fine to use different VLANs/subnets for different RF domains.
  - a. If the client can roam from one building to another building in the campus without dropping off from SSID, then a single VLAN should be used in that campus.  
The typical use cases using the Single VLAN design include university campuses, large enterprise headquarters, and hospitals with multiple buildings.
  - b. If the customer has multiple buildings and offices spread across different geographic locations, then RF is not going to be contiguous. Thus, they should use separate VLANs.  
The typical use case is branch offices or home offices.
  - c. Customers must not use VLAN derivation by user role or server-rule. If clients connected to the same SSID are in different VLANs, then the behavior is the same as VLAN pooling. Thus, all the concerns mentioned related to VLAN pooling will be applied to such a design.
2. The large subnet is for wireless clients connected to the same SSID only. Use different VLANs for different SSIDs (for example, Employee vs. Guest SSID).
3. If customers have multiple controllers serving a campus network with contiguous RF, then they should extend the client VLAN all the way back to their core switch, router, or firewall where the default gateway exists.
4. Customers using the Single VLAN design must ensure that their uplink switches, router, or firewalls are able to support the required number of ARP and MAC address table entries depending on the subnet size as discussed in section 4.3.
5. Ensure that the DHCP server can support DHCP scope as large as the size of the client VLAN.

## Deployment Steps

Customers who are planning to use the single VLAN design for their wireless clients should follow the steps below:

1. Find out the number of wireless clients on the network (in the same campus with contiguous RF) and add 20% more to incorporate future growth. AirWave or any other network management system can help.
2. Decide the size of the subnet and find out IP space that can accommodate it.
3. Create this new VLAN on your controller, distribution switch, and core switch/router/firewall as needed.
4. Create a DHCP scope for this new VLAN on your DHCP server and configure the IP helper address on your switch or router as needed.
5. Enable optimization on the WLAN controller to drop unnecessary broadcast multicast traffic. Refer to Chapter 6 about all of the knobs and their recommended settings.
6. Assign a new VLAN to the Virtual AP profile on the controller during the maintenance window. Usually universities do it over the summer. For enterprise customers, they need to do it during the maintenance window.
7. Monitor the network for the next few weeks and then retire the older VLAN scopes if everything is working fine. Make sure there are no performance, roaming, connectivity, DHCP, IPv6, or firewall policies related issues. Once that is assured, older VLAN scopes can retire.

## Recommendations

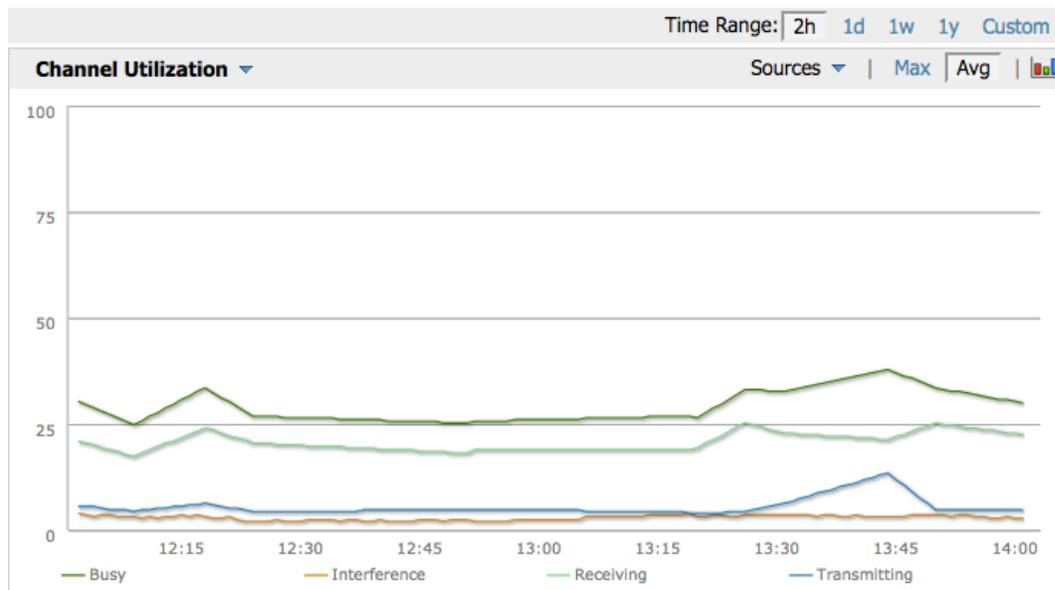
1. Use ArubaOS 6.4.3.4 or higher.
2. Size of the subnet - Typically in the range of /22 to /16.
3. Distribution Switch - Need to support a large number of MAC address table entries based on the size of the subnet selected.
4. Core Switch or Firewall (default gateway) - Needs to support a large number of ARP table entries depending on the size of the subnet selected.
5. DHCP server - Needs to support DHCP scope as large as the size of the subnet selected for wireless clients.

## Validation

Verify the RF/Channel utilization in a Network Management System like AirWave. If there is a lot of broadcast and multicast traffic flooding to wireless clients, then the RF utilization will be significantly high. Normally on 2.4 GHz channel utilization will be between 20 to 40%, and on 5 GHz it will be between 10 to 30%. Of course the number may vary if you are in a multivalent facility or due to another source of interference.

[Figure 16](#) below provides an example of channel utilization.

**Figure 16** Channel Utilization



This chapter provides information about enabling optimization knobs on the Aruba controller as discussed in [Chapter 5](#).

This chapter includes the following topics:

- [Convert Broadcast ARP Request to Unicast](#)
- [Suppress ARP](#)
- [Optimize Duplicate Address Detection](#)
- [Drop Broadcast and Unknown Multicast](#)
- [AirGroup](#)
- [IGMP Snooping and DMO](#)

## Convert Broadcast ARP Request to Unicast

### CLI

```
(config)# wlan virtual-ap <profile-name>  
(Virtual AP profile <profile-name>)# broadcast-filter ARP
```

### WebUI

**Figure 17** Convert broadcast ARP request to unicast navigation

**Configuration > AP Group > Edit "default"**

**Figure 18** Convert broadcast ARP request to unicast checkbox

Virtual AP profile > default Show Reference Save As Reset

Basic Advanced

**General**

Virtual AP enable

VLAN  ?

Forward mode

**RF**

Allowed band

Band Steering

Steering Mode

**Broadcast/Multicast**

Dynamic Multicast Optimization (DMO)

Drop Broadcast and Unknown Multicast

Convert Broadcast ARP requests to unicast

## Suppress ARP

### CLI

```
(config)# interface vlan <id>  
(config-subif) #supress-arp
```

### WebUI

Figure 19 *Suppress ARP navigation*

**Network > IP > IP Interface > Edit VLAN (101 )**

Figure 20 *Enable Suppress ARP checkbox*

The screenshot shows the configuration page for VLAN 101. At the top, there is a dropdown menu for 'IPv4' and a text input field containing '101'. Below this is a section titled 'DHCP Helper Addresses' with a button 'Add' and a text input field containing 'No Helper Addresses Conf'. Underneath, there are several configuration rows: 'Option-82' with a dropdown menu set to 'None', 'MTU [1280 - 1500]' with a text input field containing '1500', and 'Enable Suppress ARP' with a checked checkbox. The 'Enable Suppress ARP' row is highlighted with a black border.

## Optimize Duplicate Address Detection

### CLI

```
(config)# firewall optimize-dad-frames
```

### WebUI

Not available from the WebUI.

## Drop Broadcast and Unknown Multicast

### CLI

```
(config)# wlan virtual-ap <profile-name>  
(Virtual AP profile <profile-name>)# broadcast-filter ALL
```

### WebUI

Figure 21 *Drop broadcast and unknown multicast navigation*

**Configuration > AP Group > Edit "default"**

Figure 22 Drop broadcast and unknown multicast checkbox

Virtual AP > default Show Reference Save As Reset

Basic **Advanced**

**General**

Virtual AP enable	<input checked="" type="checkbox"/>
VLAN	<input type="text"/> ?
Forward mode	tunnel

**RF**

Allowed band	all
Band Steering	<input type="checkbox"/>
Steering Mode	prefer-5ghz

**Broadcast/Multicast**

Dynamic Multicast Optimization (DMO)	<input type="checkbox"/>
Drop Broadcast and Unknown Multicast	<input checked="" type="checkbox"/>
Convert Broadcast ARP requests to unicast	<input checked="" type="checkbox"/>

## AirGroup

Refer to the following for more details:

- AirGroup Deployment Guide  
<http://community.arubanetworks.com/aruba/attachments/aruba/unified-wired-wireless-access/15478/1/ArubaAirGroup-6136-DG.pdf>
- Aruba User Guide  
<https://support.arubanetworks.com/Documentation/tabid/77/DMXModule/512/Default.aspx?EntryId=8863>

## CLI

(Config): airgroup enable/disable

## WebUI

Figure 23 AirGroup

**Advanced Services > AirGroup**

AirGroup service details **AirGroup Settings**

**Global Setting**

AirGroup Status	Disabled
AirGroup MDNS Status	Disabled
AirGroup DLNA Status	Disabled
AirGroup CPPM enforce registration	Enabled
AirGroup IPV6 Support	Disabled
AirGroup CPPM query interval	1 Hours
AirGroup location discovery	Enabled
AirGroup Active Wireless Discovery	Disabled

# IGMP Snooping and DMO

## CLI

```
(config)# interface vlan <id>
(config-subif)#ip igmp snooping

(config)# wlan virtual-ap <profile-name>
(Virtual AP profile <profile-name>)#dynamic-mcast-optimization
(Virtual AP profile <profile-name>)#dynamic-mcast-optimization-thresh 80
```

## WebUI

Figure 24 IGMP navigation

**Network > IP > IP Interface > Edit VLAN ( 101 )**

Figure 25 Enable IGMP checkbox and Enable IGMP Snooping button

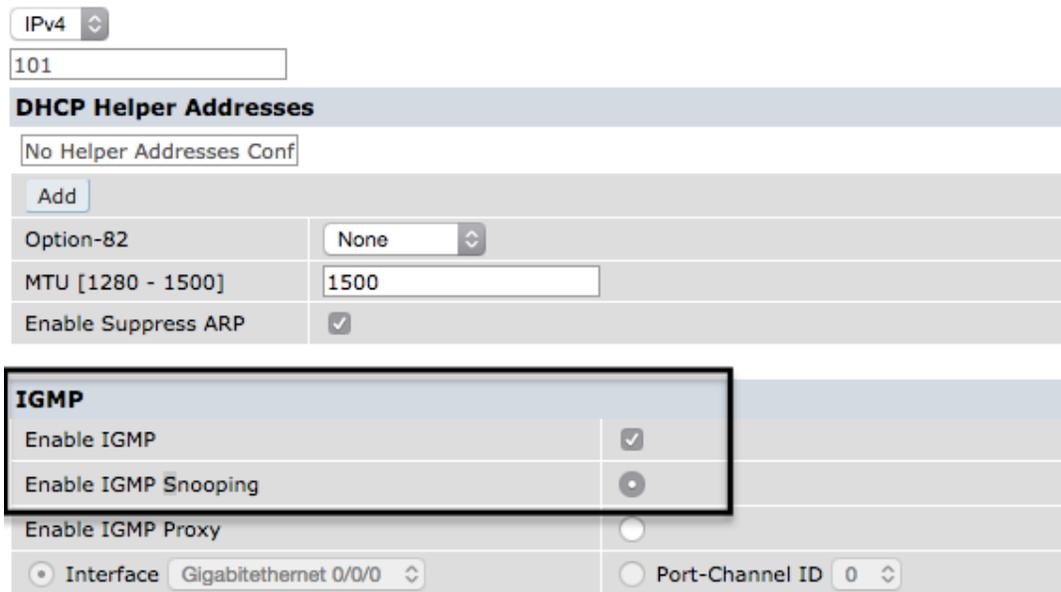
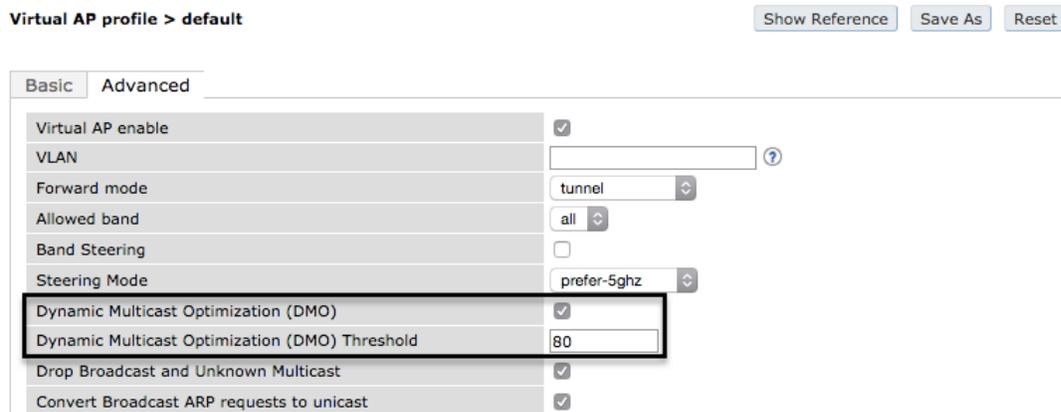


Figure 26 DMO checkbox and DMO Threshold field



As discussed in this document, there has been a traditional way of designing wired and wireless networks using multiple VLANs and VLAN pooling. Although VLAN pooling has been working out fine for many years, it has some design complexities and challenges especially with IPv6. At the same time, the Single VLAN design provides a great alternative to make design simple and avoid many complexities associated with VLAN pooling as discussed in this document.

We think there are going to be cases and requirements where VLAN pooling needs to be used and it should be used, but at the same time for many other requirements (especially universities, public Wi-Fi, and large enterprise campuses) the Single VLAN design should be considered.

The Aruba controller and AP has features to control and optimize broadcast and multicast traffic associated with a large VLAN. AirGroup and Dynamic Multicast Optimization features make it possible to use multicast based applications and streaming on wireless without affecting the performance of the wireless clients. The Aruba controller takes care of optimizing multicast traffic related to IPv6 as well.

Network architects should consider this simplified design and follow the guidelines and recommendations provided in this guide to implement it successfully.