aruba
a Hewlett Packard Enterprise company

Aruba Instant

Chapter 2 – Branch Connectivity

Version 2.0.1

Authors:
Vishal Mann
Roopesh Pavithran
Andrew Tanguay

Contributors:
Sathya Narayana Gopal
Yan Liu

Validated Reference Design

Copyright Information

Open Source Code

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

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

www.arubanetworks.com

3333 Scott Blvd

Santa Clara, CA 95054

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

Fax 408.227.455

Contents

Branch Connectivity

Distributed Network Design

A branch office is a location other than a main office where business is conducted however the term has different connotations depending on the type of organization deploying the branch. In the context of a retail chain the term *branch* represents stores that serve customers, whereas for a traditional enterprise organization a branch would be defined as an offsite location where employees and contractors congregate for their daily work. Regardless of the type of organization a branch serves the main objective of any branch office network is to provide the following functions:

- Secure employee access
- Guest Access
- Support applications such as voice and video
- Devices like printers, mobile, kiosks, security cameras
- Comply with regulations such as PCI, HIPPA, and CALEA
- Secure sensitive data
- Provide a highly-available network

Branch offices generally have a need for secure communication with the centralized corporate network. This connectivity is typically provided through of WAN connectivity options such as leased lines, MPLS, or forming a VPN over the public Internet. Connecting branch offices through leased lines can be extremely cost prohibitive compared to options such as MPLS or VPN over Internet. The decision of whether to use VPN or MPLS for branch connectivity is dependent on numerous factors which need to be weighed including cost, security policies, and service availability.

Aruba Instant is a powerful platform which is fully capable of providing wireless connectivity for a branch office network. The number of IAPs required in a network depends on factors such as number of users, the size of the branches, and the type of services required at each branch. The physical design options that are available with Aruba Instant for branch office networks as follows:

- Single IAP branch
- Multi-IAP branch

Single IAP Branch

Single-IAP deployments consist of branches that are supported by a single AP. These branch locations typically have no more than 30 wireless users and a handful of wired devices.

Examples of single-IAP branch deployments include locations such as:

- Home offices,
- Home-based call centers,
- Small retail stores (i.e. a coffee shop or restaurant chain)
- Mobile clinics
- Offsite offices of law firms
- Realty groups

In addition to wireless access, some single-IAP deployments require support for wired devices. IAP models with extra wired ports are ideal for these deployments because they simplify the network design and eliminate the need for additional switching equipment.

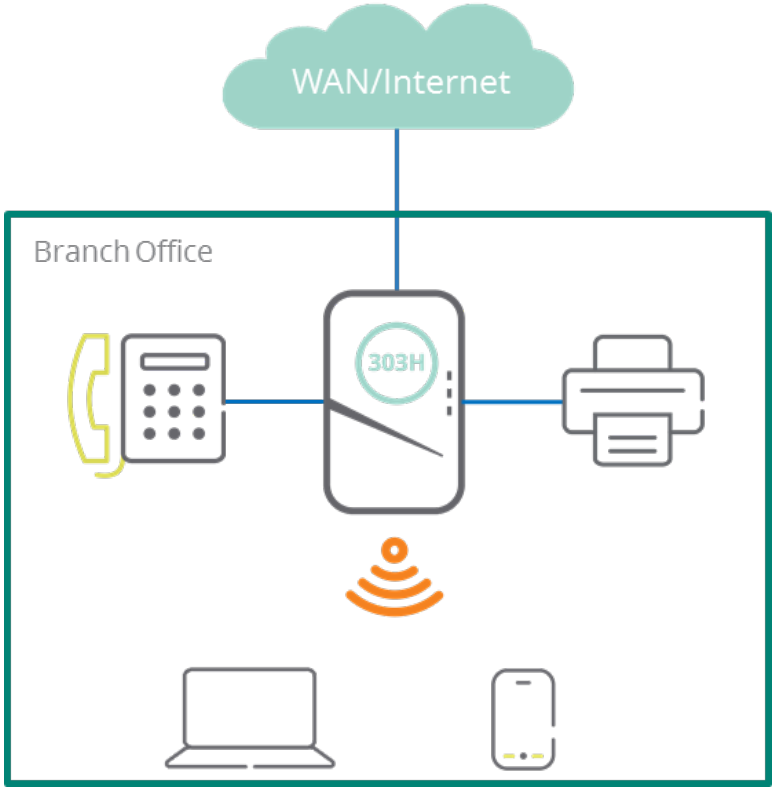

Figure 2-1 *Typical Single IAP Branch Deployment*

The uplink Ethernet port of the IAP is directly connected to the WAN uplink which eliminates the need for additional networking infrastructure at the branch. An IAP with a USB modem is also capable of acting as an uplink.

Multi-IAP Branch

The Multi-IAP Branch design consists of two options:

- Hierarchical Mode
- Flat Mode

Hierarchical Mode

In Hierarchical Mode one port of the multiport IAP acts as an uplink since it is connected to the WAN network. The remaining IAP ports are referred to as downlink ports and can be used to connect other multiport IAPs or wired devices. The IAP that is connected to the WAN through the uplink port is referred to as the *root IAP*. The root IAP provides DHCP services and well as a Layer 3 connection to the ISP WAN uplink through NAT. The root IAP will always win the Master election for the Aruba Instant cluster.

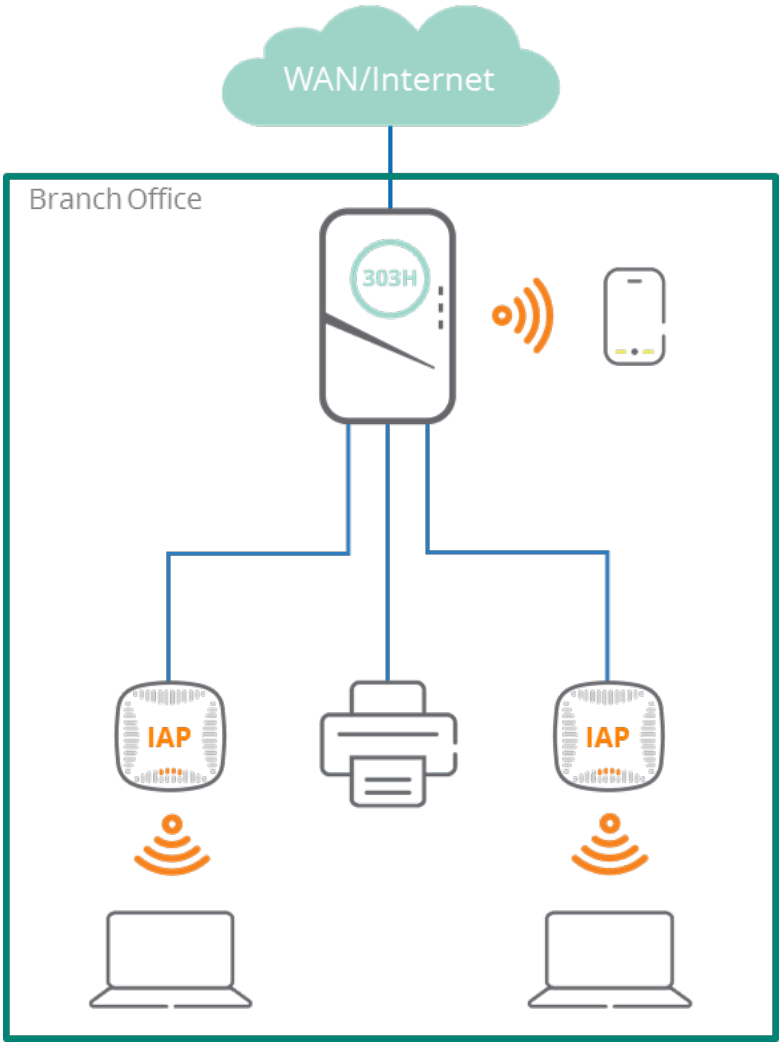

Figure 2-2 *Typical Multi-IAP Branch Hierarchical Mode Deployment*

Only the root IAP in the network uses its downlink port(s) to connect to the other IAPs. Other IAPs connected to the root IAP can only use their ports to connect unmanaged switches or end devices. Daisy chaining is not allowed.

NOTE

Aruba advises against using Hierarchical Mode if more than 5 IAPs are required in a network.

Flat Mode

The Flat Mode design is a default deployment model for a multi-IAP network and is recommended for all branch networks that require more than five IAPs. In Flat Mode, all of the IAPs deployed at the branch are connected to an uplink switch. If the Aruba Instant cluster is required to support multiple VLANs then the uplink switch must be a managed switch. In addition, the IAPs must be trunked to that uplink switch so that they may carry the appropriate VLAN tags.

E.g., if the AP VLAN, the employee VLAN, and the guest VLANs are VLANS 10, 20, and 30 respectively, then the IAPs should be trunked into the uplink switch with the native VLAN 10 and tagged VLANs of 20 and 30. The figure below represents what a typical Multi-IAP Flat Mode branch architecture would look like:

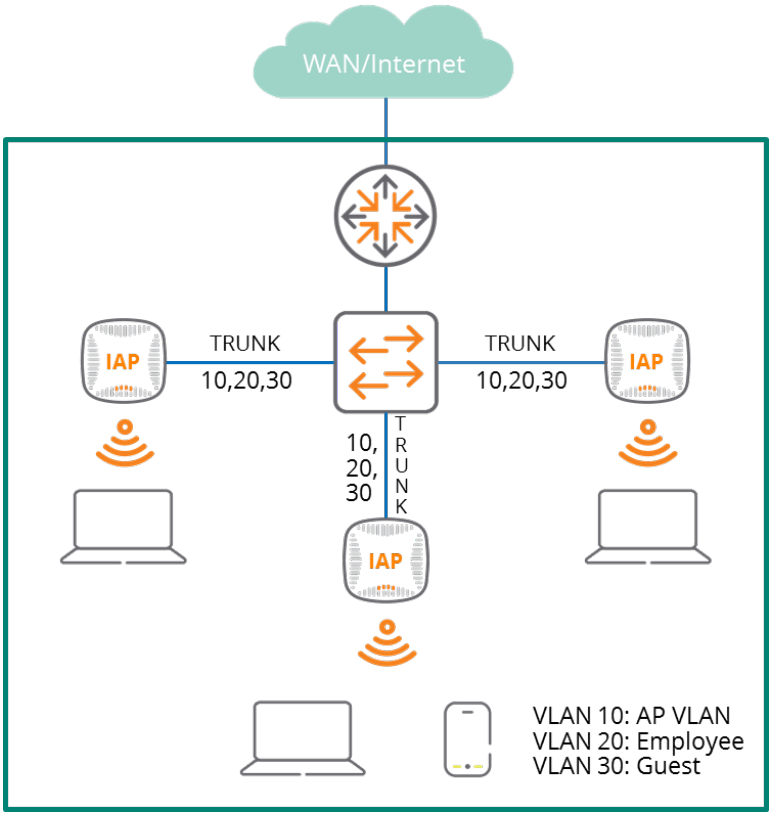

Figure 2-3 *Typical Multi-IAP Branch Flat Mode Deployment*

In general, the Hierarchical Mode design is more applicable to VPN-based branch deployments than it is to MPLS-based deployments. MPLS-based deployments typically employ either a single-IAP design or a Multi-IAP design in Flat Mode.

Aruba recommends deploying IAPs in a Flat Mode design if a managed switch is available.

VPN Branch Deployment

Connecting branches using MPLS certainly has advantages however is not always the best option for some branch office deployments due to cost and service availability concerns. Cost savings are a key driver in the adoption of distributed enterprise strategy and organizations seek a more cost-effective alternative to MPLS. Internet broadband service with high service availability and affordability provide an attractive alternative to an MPLS-based WANs.

In recent years both consumer grade and business grade broadband services have become faster, more reliable, and more affordable. This in turn has led many organizations to switch to a broadband service for branch and home office connectivity. In the case of organizations that support home-based employees broadband is the only choice as connecting home offices with an MPLS-based WAN is not a viable solution.

If an organization is using MPLS for branch connectivity the service provider ensures data security over the WAN on their behalf. However, when using the public internet for branch connectivity the responsibility for ensuring data security is the responsibility of the corporate IT team. The most common VPN technologies that provide secure remote access are SSL VPN and IPsec VPN. SSL VPN is well suited to provide remote access to a specific application, however it is not suitable for connecting enterprise networks. For that reason IPsec VPN is the most common choice for securely extending corporate networks and resources to remote sites. IPsec VPN protects sensitive data by interconnecting the remote sites with secure encryption tunnels over the Internet.

Historically, the implementation of IPsec VPNs were site-to-site. Implementing IPsec VPN requires IPsec-capable hardware at each remote site and involves complex configurations. Most branch sites have limited or no IT staff onsite, so interconnecting branches using IPsec VPN can be challenging.

Aruba Instant is designed to alleviate the complexity associated with deploying site-to-site IPSec VPNs. Aruba Instant's native VPN capabilities and zero-touch provisioning greatly reduce the challenges that normally come along with deploying IPSec VPN. Instant's zero-touch provisioning capabilities reduce deployment costs and eliminate the complexity that is normally associated with traditional IPsec VPN deployments.

Since IAPs have a virtual controller architecture the Instant network there is no need for a physical controller to provide the configured WLAN services at remote sites. However, a physical controller is required for terminating VPN tunnels from the Instant AP networks at branch locations to data centers where the Aruba controller acts as a VPN concentrator (VPNC).

When a VPN is configured, the IAP acting as the VC creates a VPN tunnel to an Aruba Mobility Controller in a corporate office. The controller exclusively functions as a VPN endpoint and does not supply the IAP with any configuration. Aruba recommends deploying IPSec VPNs with Instant for the following scenarios:

1. Enterprises with many branches that do not have a dedicated VPN connection to the corporate office.
2. Branch offices that require multiple Instant APs.
3. Individuals working from home and, connecting to the VPN.

The survivability feature of Instant APs with the VPN connectivity of Remote APs allows you to provide corporate connectivity on non-corporate networks

The WLAN controller is not responsible for anything other than acting as a VPNC for IAP networks in remote branches.

Architecture

The IAP TUNNEL architecture includes the following two components:

1. Instant APs at branch sites
2. Controller at the Data center

The Master IAP at the branch site serves as the VPN endpoint and the controller located in the datacenter serves as the VPN concentrator. When an IAP is set up for VPN it forms an IPsec tunnel to the controller in the datacenter to secure sensitive corporate data.

IPsec authentication and authorization between the controller and the IAP is based on the Remote AP whitelist configured on the controller. Only the Master IAP of the cluster forms the VPN tunnel to the VPNC. From the controller's perspective, the Master IAPs that form the VPN tunnels are considered VPN clients.

The controller's purpose in this scenario is to terminate VPN tunnels as well as route or switch VPN traffic. The IP cluster creates an IPsec or GRE VPN tunnel from the VC to a Mobility Controller in a branch office. The controller only acts as an IPsec or GRE VPN endpoint. It does not provide any configuration or management of any kind for the IAP. The figure below provides a visual depiction of the IAP TUNNEL architecture:

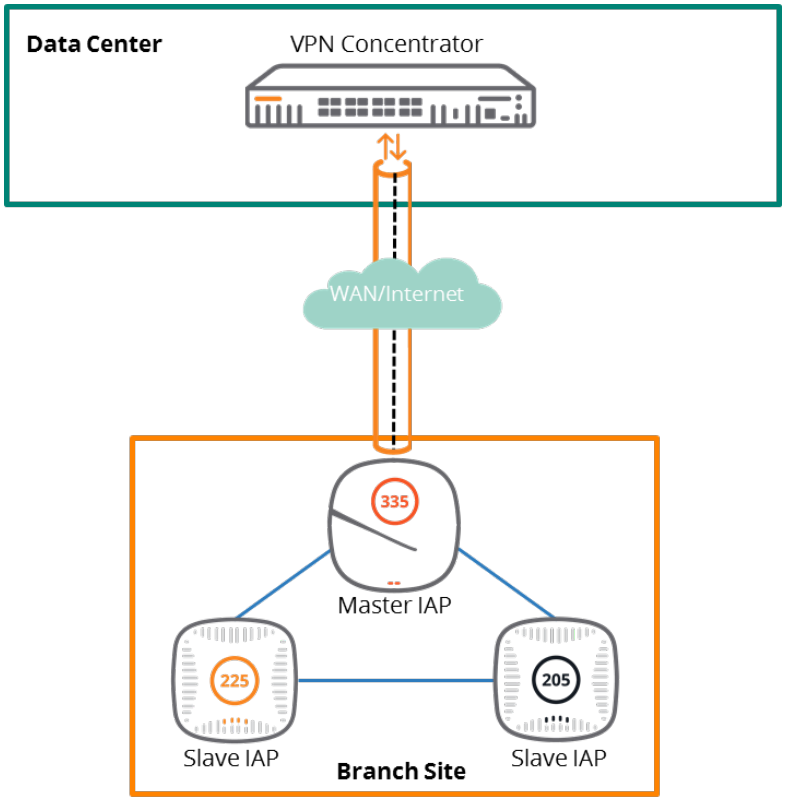

Figure 2-4 *IAP Tunnel Architecture*

When an IPsec connection is established between the WLAN controller and an IAP, each end of the IPsec tunnel has two IP addresses: an Inner IP address and an Outer IP address. By default, the WLAN controller assigns them the following roles:

- **Outer IP address:** Logon
- **Inner IP address:** Default VPN with an "allow all" access control list (ACL)

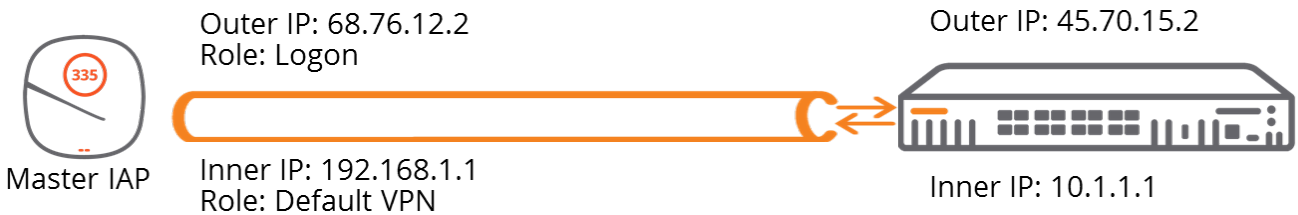

Figure 2-5 *IP Address Role Assignment*

Licensing

The WLAN controller considers the Master IAP that establishes a tunnel as a VPN client and not an AP which means that licenses such as the AP capacity license, PEFNG license, and RFProtect license are not required. However, a PEFV license is required in one of the following scenarios:

- Changing the ACLs in the default VPN role present in the controller
- Changing the role that is applied to the inner IP address and the ACLs within that role

Licenses	Features
Base ArubaOS	IAP can terminate a VPN tunnel and pass VPN traffic. Roles and polices cannot be edited.
ArubaOS with a PEFV license	IAP can terminate a VPN tunnel and pass VPN traffic. The default role in the default IAP VPN authentication profile of a WLAN controller can be edited. New user roles with custom firewall policies can be applied.

Table 2-1 *IAP License Functionality Descriptions*

Tunneling Options

Instant supports the following tunneling protocols for remote access:

- **Aruba IPSec** - IPsec is a protocol suite that secures IP communications by authenticating and encrypting each IP packet of a communication session. IPsec tunnels can be configured to ensure that the data flow between the networks is encrypted. However, a split-tunnel can also be configured which will only encrypt the corporate traffic. When IPsec is configured it is important to add the Instant AP MAC addresses to the whitelist database stored on the controller or external server. IPsec supports Local, L2, and L3 modes of IAP TUNNEL operations

Instant APs only support IPSec with Aruba Controllers.

- **Layer-2 GRE** - GRE is a tunnel protocol for encapsulating multicast, broadcast, and L2 packets between a GRE-capable device and an endpoint. Instant APs support the configuration of L2 GRE tunnel with an Aruba controller to encapsulate the packets sent and received by the Instant AP

  The GRE configuration for L2 deployments can be used when there is no encryption requirement between the Instant AP and controller for client traffic. Instant APs support two types of GRE configuration:

  - **Manual GRE** - The manual GRE configuration sends unencrypted client traffic with an additional GRE header and does not support failover. When manual GRE is configured on an Instant AP it is important ensure that the GRE tunnel settings are enabled on the controller

- **Aruba GRE** - Aruba GRE does not require any configuration on the controller except for adding the Instant AP MAC addresses to the whitelist database stored on the controller or an external server. Aruba GRE reduces manual configuration when the **Per-AP tunnel** configuration is required and supports failover between two GRE endpoints. Aruba GRE is only supported by Aruba Controllers running ArubaOS 6.4.x.x or later versions

Instant APs support manual and Aruba GRE configuration only for L2 mode of operations.

Instant APs can send IPsec and GRE heartbeat packets to Aruba Controllers. By default, Instant APs verify the status of heartbeat messages every 5 seconds and look for lost packets 6 times before marking the IPsec tunnel as down. The time intervals are fully configurable and can be modified according to the needs of network administrators.

- **L2TPv3** - The L2TPv3 feature allows the Instant AP to act as an L2TP Access Concentrator and tunnel all wireless client L2 traffic from the Instant AP to the L2TP Network Server (LNS). In a centralized L2 model the VLAN on the corporate side is extended to remote branch sites. Wireless clients associated with an Instant AP receive the IP address from the DHCP server running on the LNS. In order to receive the address the Instant AP has to transparently allow DHCP transactions through the L2TPv3 tunnel. Some important points to note about L2TPv3 in the context of Instant APs are as follows:
  - Instant supports tunnel and session configuration and uses Control Message Authentication (RFC 3931) for tunnel and session establishment. Each L2TPv3 tunnel supports one data connection and this connection is termed as an L2TPv3 session
  - IAPs only support tunneling over UDP
  - If the primary LNS goes down it will fail over to the backup LNS. L2TPv3 has one tunnel profile under which a primary peer and a backup peer are configured. If the primary tunnel creation fails or if the primary tunnel gets deleted then the backup is engaged. The following two failover modes are supported:
    - **Preemptive:** Preemptive mode means that if the primary peer comes back up while the backup is active then the backup tunnel is deleted and the primary tunnel resumes its role as an active tunnel. If preemption is configured when the primary tunnel goes down a persistence timer is triggered which will attempt to bring up the primary tunnel.
    - **Non-Preemptive:** In non-preemptive mode the backup tunnel will continue to operate after taking over from the primary tunnel even if the primary comes back up again.

L2TPV3 is not supported on IAP-205 devices.

Forwarding Modes

IAP forwarding modes determine whether the DHCP server and default gateway for clients reside in the branch or at the datacenter. These modes do not determine the firewall processing or traffic forwarding functionality. The virtual controller enables different DHCP pools (various assignment modes) in addition to allocating IP subnets for each branch. The virtual controller allows different modes of forwarding traffic from the clients on a VLAN based on the DHCP scope configured on the Instant AP.

The following forwarding modes are supported for IAP TUNNEL deployments:

1. Local mode
2. L2 Switching mode
3. L3 Routing mode

The DHCP scopes associated with these forwarding modes are described in the following sections. When configuring forwarding modes it is important to ensure that VLAN 1 is not configured for any of the DHCP scopes as it is reserved for a different purpose.

- **Local Mode** - In Local Mode the IAP cluster at that branch uses a local subnet and the Master IAP of the cluster acts as both the DHCP server and default gateway for clients. Local Mode provides access to the corporate network using the inner IP of the IPsec tunnel. The traffic destined for the corporate network is translated at the source with the inner IP of the IPsec tunnel and is then forwarded through the tunnel. All other non-corporate network traffic is translated using the IP address of the IAP and is forwarded through its uplink. When Local Mode is used for forwarding client traffic, hosts on the corporate network cannot establish connections to the clients on the Instant AP since the source addresses of the clients are translated.
- **Distributed L2 Mode** - In this mode, the Master IAP assigns IP addresses from the configured subnet and forwards traffic to both corporate and non-corporate destinations. The Master IAP acts as a DHCP server for the clients while the gateway for clients resides in the datacenter. Distributed L2 Mode can be thought of as an L2 extension of the corporate VLAN to remote sites. Either the controller or an upstream router can serve as the gateway for clients. Client traffic destined for datacenter resources is forwarded by the Master IAP through the IPSec tunnel to the default gateway in the datacenter. When an IAP registers with the VPNC it automatically adds the VPN tunnel associated to that IAP into the VLAN multicast table. This allows the clients connecting to the L2 Mode VLAN to be part of the same L2 broadcast domain on the controller.
- **Distributed L3 Mode** - Distributed L3 mode restricts all broadcast and multicast traffic to a branch which eliminates the cost and the complexity associated with a classic site-to-site VPN. Each branch location is assigned a dedicated subnet. The Master IAP in the branch manages the dedicated subnet in addition to serving as the DHCP server and default gateway for clients. Client traffic destined for datacenter resources is routed to the

controller through the IPsec tunnel, which then routes the traffic to the appropriate corporate destination. When an IAP registers with the controller a route is added to enable the routing of traffic from the corporate network to clients on the local subnet of the branch.

- **Centralized L2 Mode -** Centralized L2 Mode extends the corporate VLAN or broadcast domain to remote branches. The DHCP server and the gateway for the branch clients reside in the datacenter. Either the controller or an upstream router acts as the gateway for clients. Aruba recommends using an external DHCP server in this mode in lieu of the DHCP server on the controller. Client traffic destined for datacenter resources is forwarded by the Master IAP through the IPsec tunnel to the client's default gateway in the datacenter.
- **Centralized L3 Mode** - In Centralized L3 Mode the Master IAP acts as a DHCP relay agent by forwarding client DHCP traffic through the IPSec tunnel to a DHCP server located behind the controller in the corporate network. The Centralized L3 VLAN IP is used as the source IP. IP addresses are obtained from the DHCP server.

Forwarding Mode	DHCP Server	Client Default Gateway	Corporate Traffic	Internet Traffic	Branch Access from Datacenter
Local	Virtual Controller	Virtual Controller	Source NAT with inner IP of the IPSec tunnel	Source NAT performed with the local IP of the VC	No
Centralized L2	DHCP server in the datacenter	Datacenter controller or router	L2 reachable	Source NAT performed with the local IP of the VC	Yes
Centralized L3	DHCP server in the datacenter, VC acts as a DHCP relay	Virtual Controller	Routed	Source NAT performed with the local IP of the VC	Yes
Distributed L2	Virtual Controller	Datacenter controller or router	L2 reachable	Source NAT performed with the local IP of the VC	Yes
Distributed L3	Virtual Controller	Virtual Controller	Routed	Source NAT performed with the local IP of the VC	Yes

Table 2-2 *Forwarding Modes Feature Matrix*

NOTE

Local Mode, Centralized L2 Mode, and Distributed L3 mode are covered in depth as they are the most commonly employed forwarding modes.

Local Mode

Local Mode with Aruba Instant is similar to the local network of a home wireless router with the exception that it has VPN capabilities in addition to other enterprise grade features. The IAP cluster at the branch has a local subnet (e.g., 192.168.200.0/24) and the Master IAP of the cluster functions as the DHCP server as well as the default gateway for clients. Local Mode enables VPN capabilities by using the inner IP address of the Instant-VPN IPsec tunnel. Client traffic destined for corporate destinations is source NATed by the Master IAP using the inner IP address of the IPsec tunnel. Traffic that is destined for the Internet or local destinations is source NATed using the local IP address of the Master IAP. It is essential that the IP addresses that are defined in the VPN address pool of the WLAN controller (which is used for inner IP addresses of IPsec tunnels) are routable from the upstream router in the data center. If required, all client traffic can be forwarded through the IPsec tunnel or bridged locally.

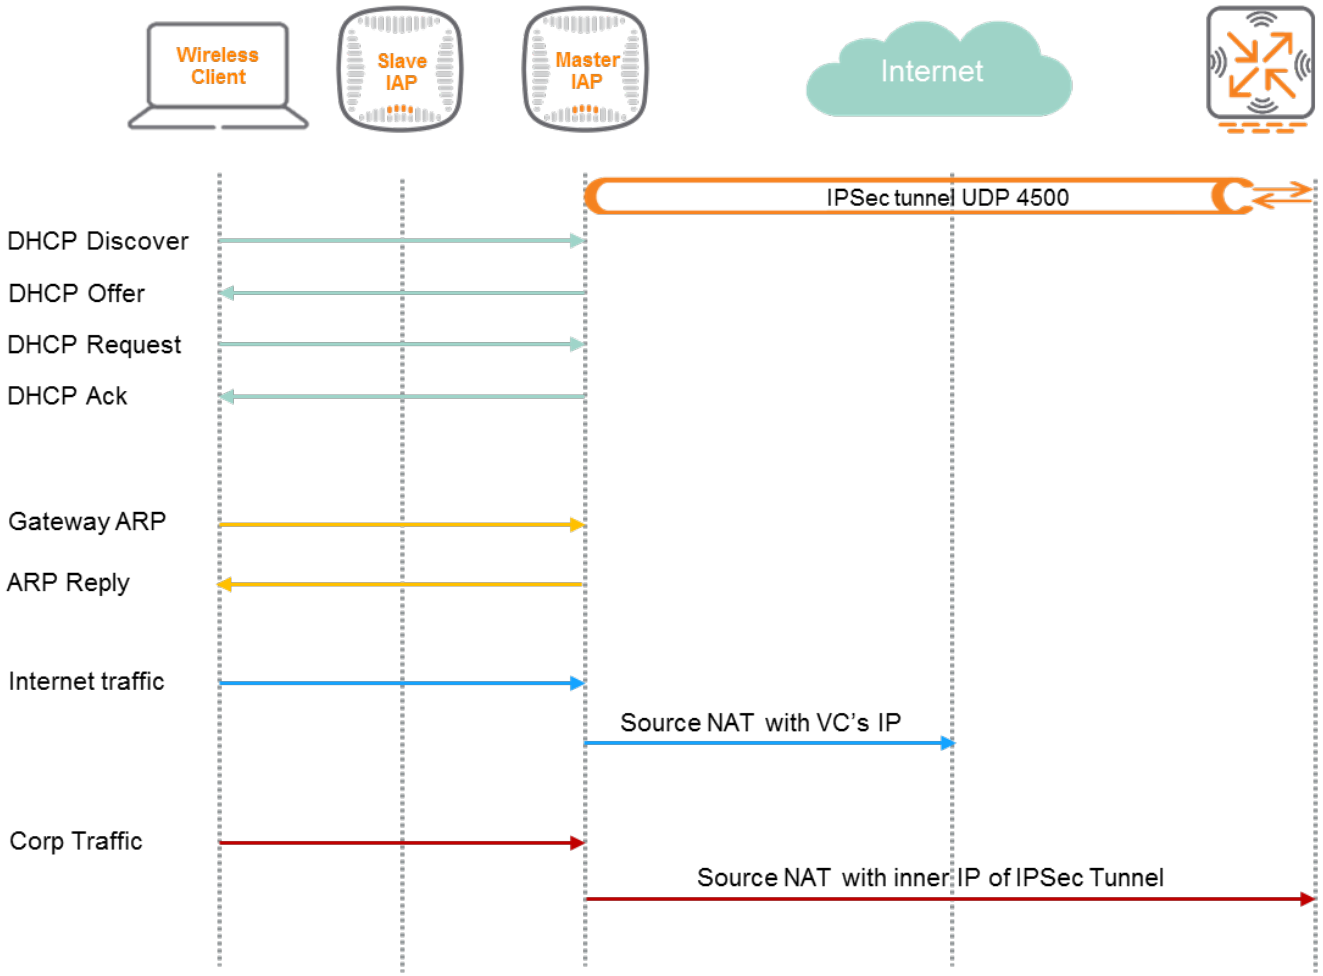

Figure 2-6 *Local Mode Forwarding*

In Local Mode, clients in the branch can initiate connections to a server in the data center, however connections cannot be initiated from the datacenter to remote clients. The behavior is similar to that of a NAT device. The WLAN controller and the upstream routers have no visibility or direct route to the branch subnet. Therefore, connections cannot be initiated from the data center to remote clients for troubleshooting purposes.

Local mode is ideal for branch guest networks using a captive portal sever in the datacenter for guest authentication.

Centralized L2 Mode

Centralized L2 Mode extends the corporate VLAN and broadcast domain to remote branches. The DHCP server and the gateway for branch clients both reside in the data center. Either the WLAN controller or an upstream router act as the gateway for clients. Aruba recommends using an external DHCP server for DHCP services in Centralized L2 mode rather than the DHCP server on the WLAN controller.

Centralized L2 mode has two options for forwarding traffic:

1. All traffic including guest traffic is forwarded by the Master IAP through the IPsec tunnel to the default gateway in the data center. This option is typically selected for organizations that prefer to exercise more control over guest traffic by having it forwarded to the data center
2. Alternatively, a split-tunnel can be used which forwards traffic destined for corporate resources through an IPsec tunnel to the VPNC while internet traffic is bridged locally. This option is effective for scenarios where control over guest traffic is less of a concern while also alleviating overhead on the corporate network

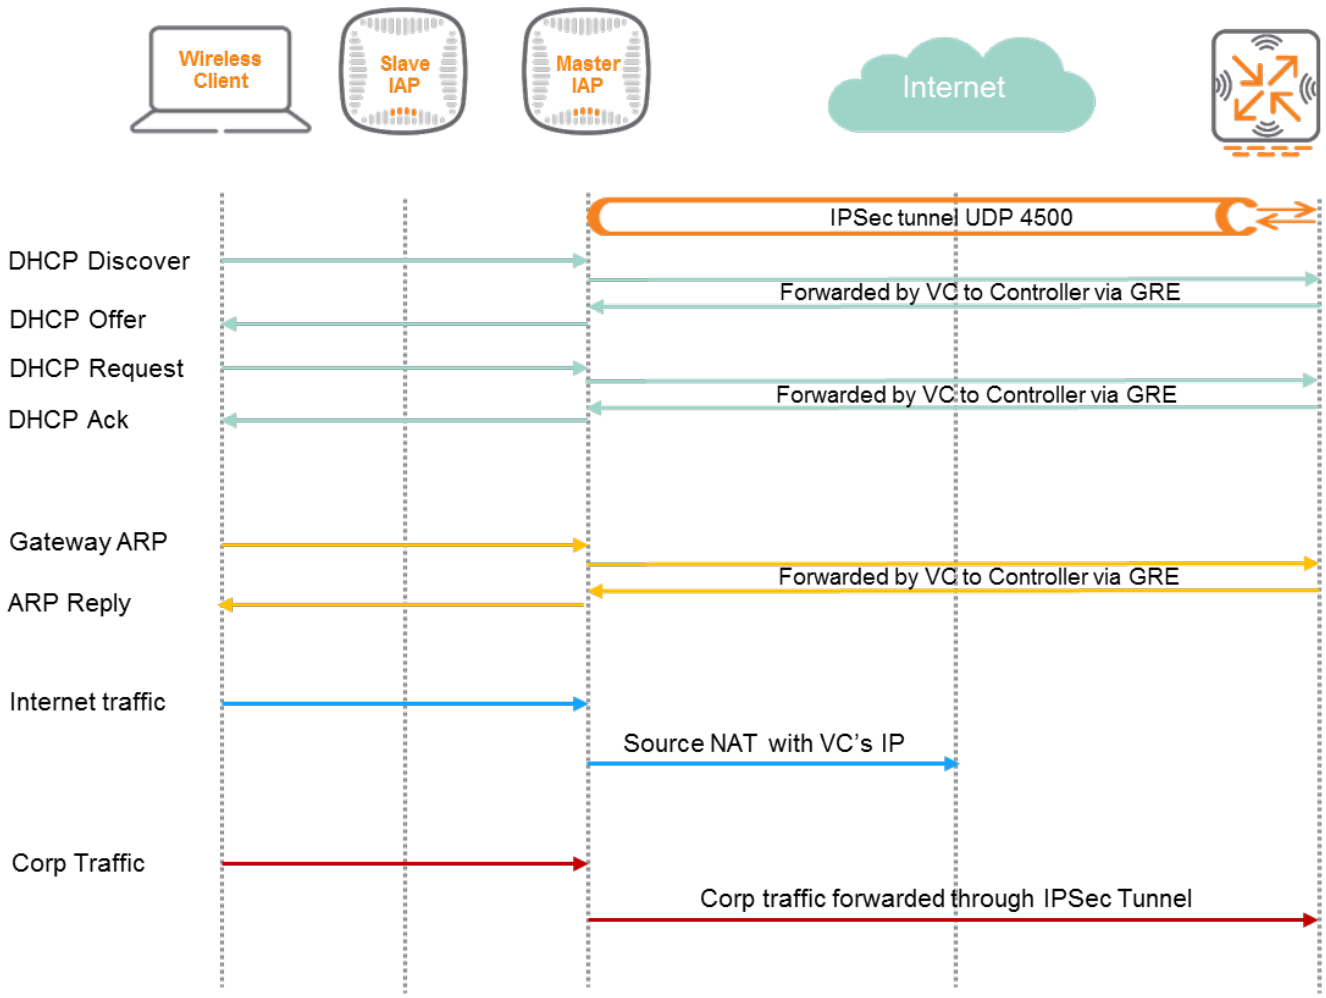

Figure 2-7 *Centralized L2 Mode Forwarding*

In Centralized L2 Mode connections can be initiated from the data center to remote clients for troubleshooting purposes. If RADIUS traffic is not source NATed at the WLAN controller, the VPN pool for inner IP addresses must be made routable for RFC 3576-compliance and 802.1X. A routable VPN address pool also allows access to the local WebUI of the Aruba Instant cluster from the datacenter.

NOTE

Aruba recommends using Centralized L2 Mode only if Layer 2 extension is mandatory for branches.

BID Allocation

When an Instant cluster in a branch comes up for the first time, one IAP is elected as the Master IAP through the master election algorithm. The designated Master IAP in a cluster generates a branch key by hashing its own MAC address and proceeds to distribute the key to all IAP cluster members. The branch key plays a significant role in ensuring that a branch is allocated the same subnet and IP addresses regardless of which IAP becomes the Master IAP of the cluster at a later point. This key is generated even for IAP clusters that are not configured for IAP TUNNEL.

After generating the branch key, the Master IAP forms an IPsec connection to the WLAN controller and obtains an inner IP address from its VPN address pool. The BID allocation process is initiated when the Master IAP sends a registration message to the WLAN controller. This registration message includes the following attributes:

- **Inner IP:** The inner IP address of the Master IAP that established the IPsec tunnel
- **Branch Key:** The key that was generated and distributed to all member IAPs by the Master IAP
- **MAC:** The MAC address of the Master IAP participating in the BID process
- **MAX_BID and subnet name:** The maximum number of subnets or IP address blocks that can be created based on the subnet size and the client count configured on the IAP

In addition to the *MAX_BID*, the IAP sends the corresponding subnet name to the VPNC. The subnet name is derived from IP address range in the configuration as well as the client count for each mode. E.g., if an organization uses 10.10.0.0/16 with 250 clients per branch, the IP configuration on the IAP is 10.0.0.0 - 10.10.255.255 instead of 10.10.0.0/16. The name of the L3 subnet will appear in the CLI as “10.0.0.0 – 10.10.255.255,250”.

The subnet name keeps track of which MAX_BIDs apply to which distributed mode configurations. If a branch is configured for multiple distributed modes, the IAP sends multiple combinations of MAX_BID and corresponding subnet names to the WLAN controller. This method allows a branch to have multiple SSIDs that use different distributed modes and different subnet sizes. For example an organization can have an SSID_1 with distributed L3 mode and a configuration of "10.10.0.0 /16” with 250 clients per branch and SSID_2 with distributed L3 mode and a configuration of "10.20.0.0 /16” with 100 clients per branch. The configuration on an Aruba IAP assumes the following definitions:

- **BID:** Value that specifies whether a branch is new or existing. A new branch uses a unique value in this field to specify that it requires a BID from the MAX_BID range. If the Master IAP of a branch that has already received a BID fails then a new IAP will be elected as the Master. When the newly elected Master IAP connects to the WLAN controller it will use the previously allocated BID in this field.
- **Backup:** Value that specifies whether the Master IAP is communicating with a backup host. A backup host is a backup WLAN controller where the Master IAP can initiate an IPsec connection. A backup host is similar to a backup local management switch BLMS controller in an ArubaOS campus network. It does not represent a VRRP backup to a WLAN controller.

The BID allocation process occurs between the primary host and the Master IAP. The WLAN controller serving as the primary host must be operational when a branch comes up for the first time. Any IAP TUNNEL branch brought up from a factory default configuration that is configured for Distributed L3 mode must exchange its first BID process with its primary host to receive its required address space and subnet.

Upon receiving a BID registration message, the WLAN controller determines whether a branch is new or existing by examining the BID field. If the branch is new, the WLAN controller verifies whether the branch key in the registration message is present in its database. If the branch key is not found then the WLAN controller selects an unused BID from the MAX_BID range and returns it to the Master IAP. If the branch key is already present in the WLAN controller's database, then the WLAN controller returns the BID that is already associated with the branch key. When the BID is allocated, the Master IAP uses the BID to determine the IP subnet or IP addresses that may be used. The following examples describe how the subnets are determined, based on BID value:

Consider an organization that uses a "10.10.0.0 /16" configuration with 250 clients per branch as the Distributed L3 mode configuration on IAPs in 200 branches. This configuration can support 256 branches. If a branch is assigned a BID of 0, it takes the first available /24 subnet. The subnet for the branch is 10.10.<bid>.0/24 = 10.10.0.0 /24. If a branch is assigned a BID of 50, the subnet for the branch is 10.10.<bid>.0 /24 = 10.10.50.0 /24.

After determining the IP address or the subnet that must be used, the Master IAP registers the IP addresses and IP subnet with the WLAN controller using ROUTE ADD and VLAN ADD messages. The ROUTE ADD and VLAN ADD messages notify the WLAN controller about the L3 subnet and L2 VLAN used at the branch. The Route ADD and VLAN ADD messages are not part of the BID process. These datapath programming messages are used to add the appropriate VLAN and route information to the WLAN controller datapath. If the Master IAP fails and a slave IAP becomes the new Master, then the branch key and BID will not change and the branch will continue to use the same subnet and IP addresses. The BID allocation process is depicted in the figure below:

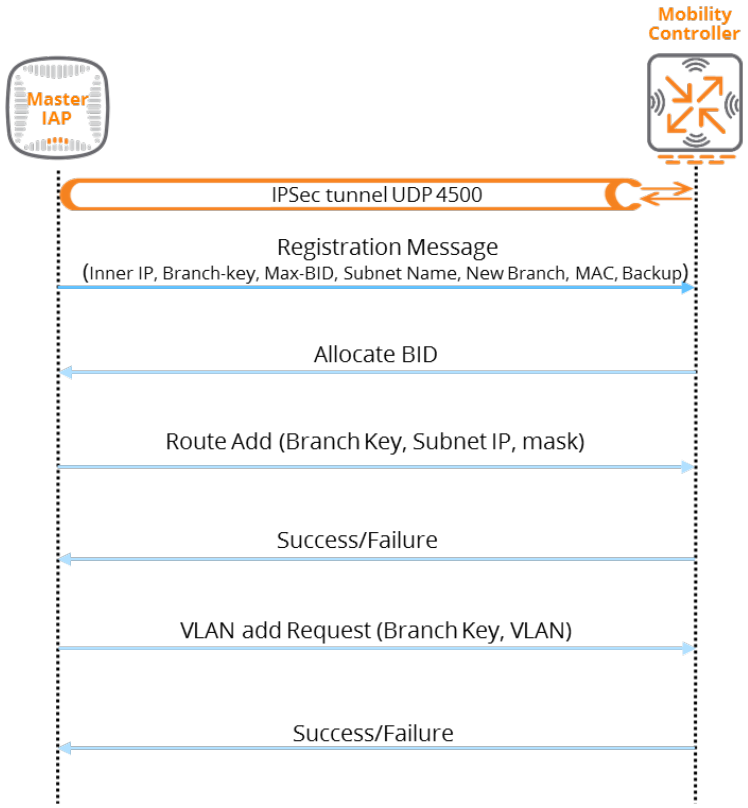

Figure 2-8 *BID Allocation process*

Distributed L3 Mode

In Distributed L3 mode, each branch location is assigned a dedicated subnet. The Master IAP in the branch manages the subnet, functions as the DHCP server, and acts as the default gateway for clients. Client traffic destined for datacenter resources is routed to the WLAN controller through an IPsec tunnel. The WLAN controller then routes the traffic to the appropriate corporate destinations as needed.

Any traffic destined for the Internet or a local destination is source NATed using the local IP address of the Master IAP and locally bridged. The WLAN controller in the datacenter is aware of the Layer 3 subnet at each branch and can redistribute these routes to upstream routers through the Open Shortest Path First (OSPF) routing protocol. All client traffic can be forwarded through the IPsec tunnel or bridged locally if required.

Distributed L3 mode allows connections to be initiated from the data center to remote clients for troubleshooting purposes. If RADIUS traffic is not source NATed at the WLAN controller then the VPN pool that is used for inner IP addresses of the IPsec tunnel must be routable for RFC 3576-compliance and 802.1X.

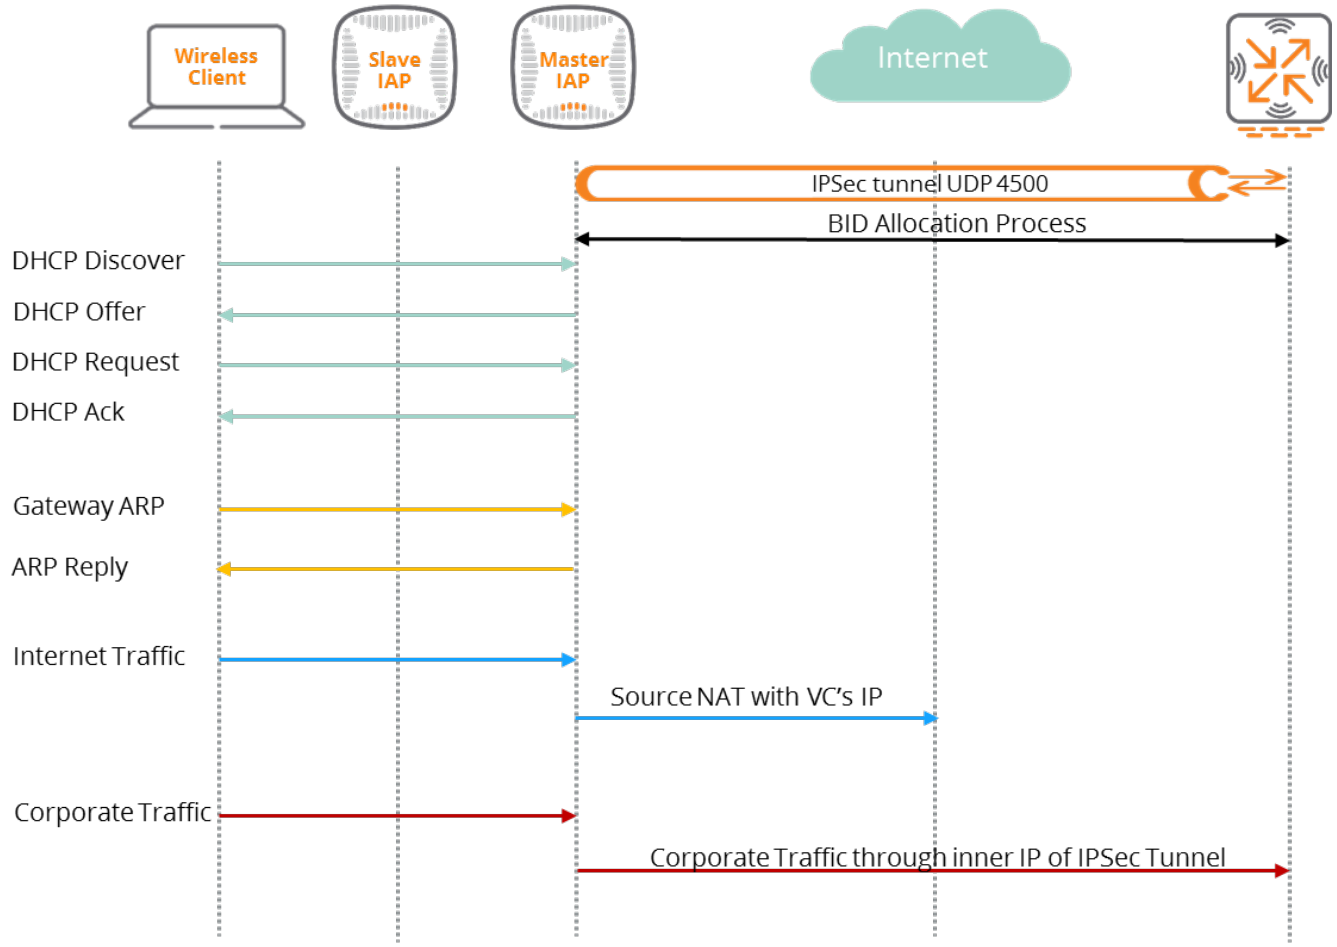

Figure 2-9 *Distributed L3 Mode Forwarding*

NOTE

Split-tunnel Mode cannot be configured in Distributed Layer 3 mode as it can in Centralized Layer 2 Mode. In Centralized Layer 2 deployments Split-tunnel Mode more frequent use case. Routing profiles can be used to achieve similar functionality with a Distributed Layer 3 deployment

NOTE

Distributed L3 mode is the recommended mode of operation for Instant-VPN networks. Centralized L2 mode should only be used by organizations which require the extension of corporate VLANs to branch networks.

Scalability

The table 9 below outlines scalability numbers for each controller model for an IAP TUNNEL deployment. The table assumes the following definitions:

- **IAP TUNNEL Branches** –The number of IAP TUNNEL branches that can be terminated on a particular controller platform
- **Route Limit** – The number of L3 routes supported on a controller
- **VLAN Limit** – The number of VLANs supported on a controller

Platforms	IAP TUNNEL Branches (Preferred)	Route Limit	VLAN Limit
7280	8192	32,769	4,094
7240	8192	32,769	4,094
7220	4096	32,769	4,094
7210	2048	32,765	4,094
7205	1024	16, 381	2,048
7030	256	8,189	256
7024	128	4,093	128
7010	128	4,093	128
7008	64	4,093	128
7005	64	4,093	128

Table 2-3 *IAP Tunnel Scalability*

IAP Tunnel Authentication

In an Instant tunnel branch network users can be authenticated either through a local RADIUS server or a RADIUS server in the data center. The Master IAP in an Instant tunnel branch determines which RADIUS server is used by checking its own routing profile. Traffic destined for a RADIUS server in the datacenter is sourced using the inner IP address of the IPsec tunnel. The VPN address pool that is used for inner IP addresses of IPsec tunnels must be routable from the upstream router in the data center. If dynamic authorization extension to RADIUS (RFC 3576) is not required then a rule can be placed on the WLAN controller to source NAT all RADIUS traffic with the IP address of that WLAN controller. If a branch network has a local RADIUS server and if dynamic RADIUS proxy (DRP) is enabled on the IAP, then 802.1X traffic is source NATed with the Master IAP's IP address.

RFC 3576-compliance dictates that CoA messages must be initiated by the RADIUS server. Rules should never be enabled on the WLAN controller to source NAT all RADIUS traffic with the IP address of that WLAN controller. If the RADIUS server is located in the datacenter then the inner IP addresses of the IPsec tunnels must be listed as RADIUS clients. If RFC 3576-compliance is used with a local RADIUS server, the Master IAP's IP address must be added as the RADIUS client. DRP must be enabled for Instant networks consisting of multiple IAPs to tunnel the RADIUS traffic from the member IAPs to the authentication server in the datacenter.

When DRP is enabled, the 802.1X transactions for clients connecting to the member IAPs are forwarded to the Master IAP functioning as a RADIUS proxy. With DRP enabled, the NASIP attribute in RADIUS packets destined for the RADIUS server in the datacenter contain the inner IP address of the IPsec tunnel. DRP is not required for single IAP deployments. However, if DRP is enabled in such a deployment then the NASIP attribute in RADIUS packets destined for the RADIUS server in the datacenter will contain the local IP address of the IAP rather than the inner IP address of the IPsec tunnel. As a best practice, Aruba recommends enabling DRP in single IAP deployments with RADIUS servers that use the NAS IP attribute as a filter for authentication. The following table outlines authentication options in various Instant deployment scenarios:

RADIUS Server Location	DRP	VPN Pool Routable From DC	ACL Source NATs Traffic to Controller IP	Source IP	NAS IP	RFC 3576 Compliant
DC	Enabled	Yes	No	Inner IP of IPsec Tunnel	VPN Tunnel IP	Yes
DC	Enabled	No	Yes	Controller	VPN Tunnel IP	No
Local	Enabled	N/A	N/A	Master IAP Local	Master IAP IP	Yes
Local	Disabled	N/A	N/A	Master or Slave IAP	Master/Slave IAP IP	Yes

Table 2-4 *802.1X and RFC 3576 Options*

IAP Tunnel DNS

In a typical IAP deployment without tunneling all DNS requests from a client are forwarded to the client's DNS server by default. However, this behavior changes if an IAP is configured for tunneling.

The DNS behavior for both wired and wireless clients on an IAP network configured for tunneling is determined by the enterprise domain settings. The enterprise domain setting on the IAP specifies the domains for which DNS resolution must be forwarded to the default DNS server of the client. E.g., if the enterprise domain is configured for arubanetworks.com, the DNS resolution for host names in the arubanetworks.com domain would be forwarded to the default DNS server of the client. The DNS resolution for host names in all other domains is source NATed to the local DNS server of the IAP.

If no enterprise domain configuration exists and the client is on an SSID for IAP VPN then all DNS traffic will be source NATed to the DNS server of the IAP. If a non-VPN SSID is present, its traffic will be forwarded to the default DNS server for the client that initiated the request. If Split-tunnel Mode has been disabled, all DNS traffic will be forwarded over IPSec tunnel to DNS server of the client regardless of the enterprise domain configuration. If an asterisk is configured in the enterprise domain list instead of a domain name then all DNS requests are forwarded to the default DNS server of the client.

Branch Connectivity Scenarios

Internet Connectivity

The Internet branch connectivity scenario consists of a deployment with multiple DCs with a requirement for redundancy between the DCs and branch offices. Internet traffic is locally bridged in branch offices while corporate traffic is secured and routed to the DC. Split-tunneling of client DNS traffic is preferred for the Internet branch connectivity scenario.

Aruba IPSec is used to securely transmit data between branch offices connected to the DC through the Internet. Distributed Layer 3 mode is the preferred forwarding mode as there is no need to extend the corporate VLAN or multicast traffic from DC to branch offices. The primary and backup IPSec VPN tunnels are configured with preemption and fast failover from the branches to the DC for redundancy purposes. The IPSec VPN tunnels terminate on the controllers in the DMZs. Clients connected to the branch offices obtain their IP address either from the local DHCP server on the switch or from DHCP server on the IAP (depending on whether the branch has a Flat or Hierarchical Mode topology).

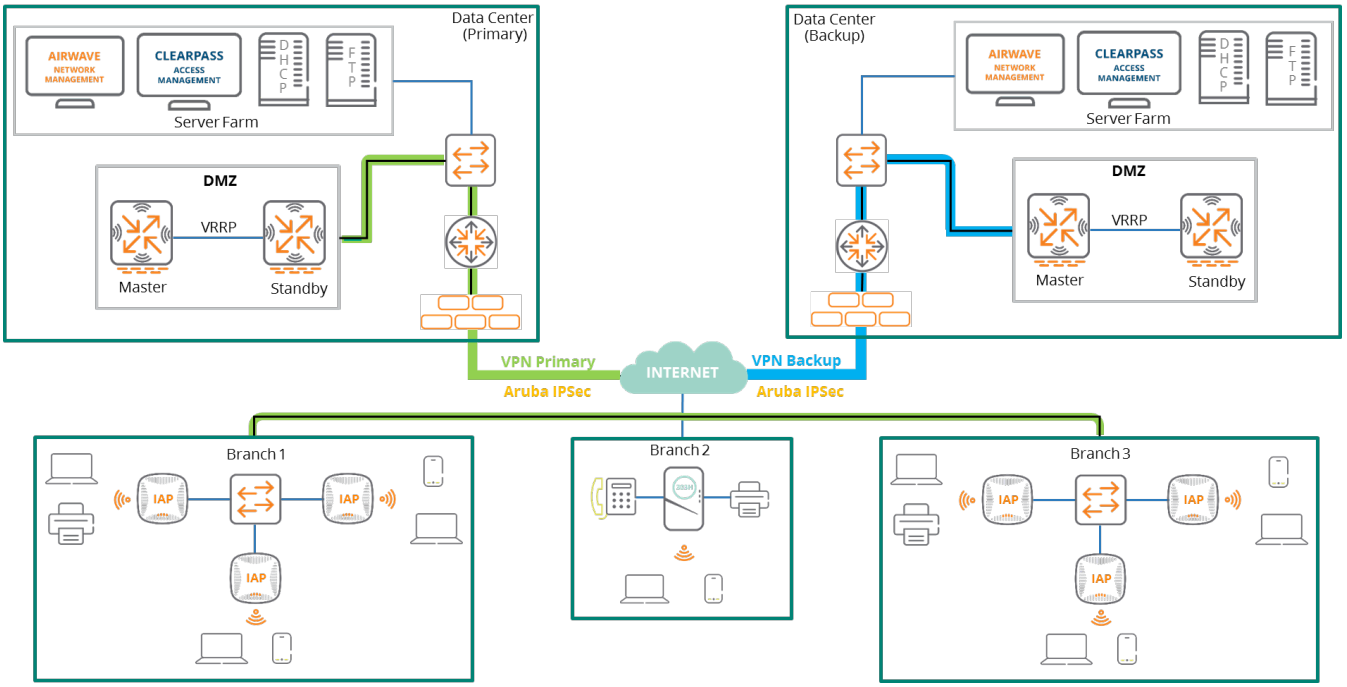

Figure 2-10 *Branch Connectivity through Internet*

Employees are authenticated with 802.1x through the ClearPass server in the DC. Guests can be presented with a captive portal from the ClearPass server as well with internet traffic bridged locally at the branch using a separate guest VLAN. Split-tunnel DNS is configured under the enterprise domain tab with a rule created only corporate domain name queries are tunneled to the DC.

MPLS Connectivity

The MPLS branch connectivity scenario consists of a deployment with multiple DCs with a requirement for redundancy between the DCs and branch offices. All employee and guest traffic is forwarded to the DC for processing.

Since the traffic from branch offices flows through MPLS network, the security of the data packets is handled by the service provider. Aruba GRE is the preferred tunneling mode and only control packets are encrypted using IPsec.

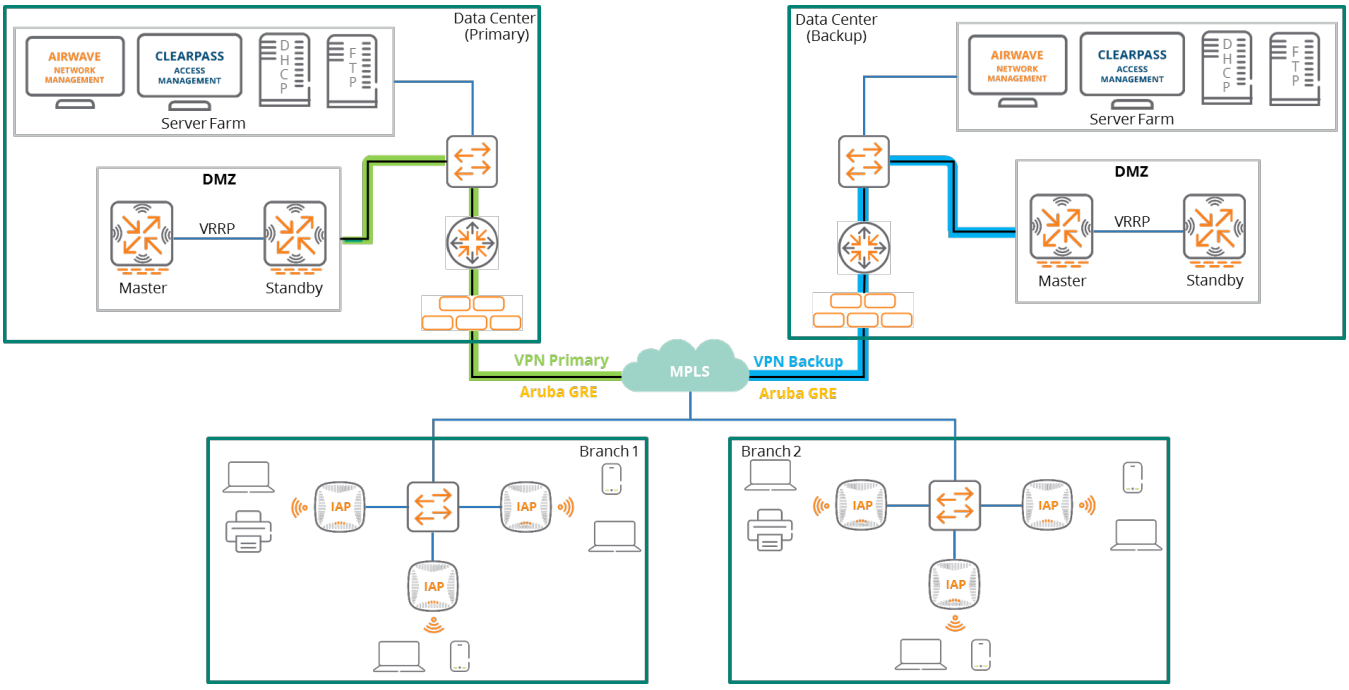

Figure 2-11 *Branch Connectivity through MPLS*

Centralized L2 Mode is the preferred forwarding mode for this scenario since the DHCP and DNS servers are centralized and employee as well as guest traffic needs to be forwarded to the DC. The ClearPass server provides 802.1X authentication so employees in the branch office can access the network.

Guest users are provided with a captive portal through the ClearPass server in the DC. All guest traffic is tunneled to the controller in the DMZ. Appropriate firewall policies should be applied to restrict guests from gaining access to internal resources. The IAPs where clients are connected tunnel the guest traffic to the controller in the DMZ. As the DNS server is also centralized an asterisk is placed in the enterprise domain list to ensure that all DNS queries are forwarded to the servers in the DC.