

Validated Solution Guide

Aruba Solution TME

May 30, 2024

Table of Contents

Aruba ESP SD-WAN & SD-Branch Design Guide	5
Introduction to Aruba SD-WAN and Branch	6
Purpose of This Guide	6
Design Goals	7
Customer Use Cases	7
Aruba EdgeConnect SD-WAN Benefits	9
Aruba EdgeConnect SD-WAN & SD-Branch WAN Transport Design	12
Commodity Internet	12
MPLS	12
3G/4G/5G	12
Private Layer 2	13
Starlink / Low Orbit	13
Recommended Pairings	13
Zero Touch Provisioning	14
Hub Designs	15
Appliance Deployment Models	15
Branch Designs	18
Internet Egress	19
Overlay Components	20
Cloud Integration	21
Microsoft Azure	21
Amazon Web Services	22
EdgeConnect SD-WAN Solution Fundamentals	24
Aruba Orchestrator	24
EdgeConnect SD-WAN Appliances	25
Business Intent Overlays	26
Boost	28
Licensing	29
EdgeConnect SD-WAN Overlay Design	31
Traffic Classification	31
Topology	31
Regional Routing	31
Interface Selection	31
Service Level Objectives	32
Link Bonding Policy	32
QOS, Security, and Optimization	33
Cloud Integration	36

EdgeConnect SD-WAN Hub Design	38
WAN Transport Integration	38
Inline Deployment Methods	38
Out-of-Path Deployment Methods	39
High Availability	40
LAN Routing Integration	41
WAN Routing Integration - Overlay	41
WAN Routing Integration - Provider	43
EdgeConnect SD-WAN Branch Design	45
Design Overview	45
SD-WAN Gateways	47
Switching	49
Wireless	50
Zero Touch Provisioning	51
Management	52
Segmentation and Policy	52
Aruba EdgeConnect SD-Branch Solution Fundamentals	55
Aruba Central	55
SD-WAN Orchestrator	55
Device Overview	58
EdgeConnect SD-Branch WAN Overview	62
WAN Topologies	62
WAN Monitoring	67
WAN Policies	69
EdgeConnect SD-Branch Hub Design	81
EdgeConnect SD-Branch Branch Design	85
Aruba SD-LAN	92
Microbranch	98
Customer Profile	101
EdgeConnect SD-WAN Reference Design	103
Component Selection	103
EdgeConnect Gateways	103
Licensing	104
Example Hub Design	107
Example Branch Design	108
Example Overlay Design	113
EdgeConnect SD-Branch Reference Design	116
Component Selection	116
Licensing Options	116
Overlay Design	117
Hub Design	117
Branch Design	118

What's New in This Version

122

Aruba ESP SD-WAN & SD-Branch Design Guide

This guide provides IT professionals information and guidance with design considerations for the SD-WAN environment:

- Hardware selection
- Software selection
- Topology
- High availability
- Scalability
- WAN application performance
- Security

Introduction to Aruba SD-WAN and Branch

Aruba's SD-WAN solutions include the right mix of software and components to implement a service-provider-independent wide-area network (WAN) that delivers enterprise-level performance and security across a disparate range of technologies and devices.

With SD-WAN, organizations can convert mission-critical applications into multiple cloud-based services so employees can work and collaborate effectively from any location. Cloud applications are secure, fast and reliable, but the WAN connections used to access them must be optimized and automated for maximum availability, productivity, and efficiency.

SD-WAN significantly enhances network capabilities for organizations with widely distributed locations and a wide range business priorities and IT needs.

Organizations often operate multiple distributed, heterogeneous networks with small centralized teams. Distributed networks require many services reliant on the WAN connectivity. Branch networks require wired and wireless LANs, security and policy enforcement, multiple levels of access for different user groups, efficient management of multiple applications with various functions, and the flexibility to change or expand a branch's infrastructure as new technology develops.

SD-Branch extends the benefits of SD-WAN operation to branch deployment, operation and maintenance. SD-Branch delivers a full-stack solution that includes SD-LAN, artificial intelligence, and security to address network connectivity needs at all levels.

Migration to SD-WAN resolves many immediate network and IT challenges cost effectively and efficiently. However, it is important to consider overall technology goals and develop strategies to meet current business requirements and align with long-term financial and organizational goals:

- Purchase as much WAN bandwidth as possible to alleviate potential bottlenecks during the busiest times of the day and allow for continued traffic growth.
- Increase Internet bandwidth instead of buying additional private bandwidth to support.
- Use cloud-based tools to simplify the configuration, operation, and management of the WAN.

Purpose of This Guide

This design guide provides information and guidance for designing SD-WAN in the Edge Services Platform (ESP) architecture and associated hardware and software components. The guide provides an explanation of the requirements that shape the design and benefits for the organization. Example architecture designs are provided for additional reference.

The guide uses an example of a single unified infrastructure that integrates headend gateways, branch gateways, switches, access points, cloud-based orchestration, artificial intelligence, and network management with access-control and traffic-control policies.

Design Goals

The overall goal is to create a simple, scalable design that can be replicated easily across all sites in the network. Selection of solution components is limited to a specific set of products proven to deliver streamlined, cost efficient operation and maintenance. Key features provided by Aruba SD-Branch include:

- **Simplicity with Zero Touch Provisioning:** SD-WAN devices can be factory-shipped directly to a remote site by automatically matching orders to an Aruba customer account.
- **Unified policy management:** For Aruba and third-party network infrastructure, Aruba ClearPass delivers a common policy framework for multivendor wired and wireless networks. The software-defined approach makes it easy for the network administrator to distribute changes quickly to meet current and emerging corporate risk management and compliance requirements. ClearPass Device Insight (CPDI) adds AI-powered device profiling to automate the discovery and monitoring of mobile and IoT endpoints.
- **Predictive analytics and assurance:** Aruba Central's artificial intelligence (AI), machine learning (ML), and automation capabilities identify issues and notify IT personnel of problems with recommended changes. With the shift to a cloud-hosted model, data are collected and crowd-sourced from Aruba's extensive base of installations to take advantage of Aruba's collective data science expertise.
- **Secure WAN connectivity:** SD-WAN technology uses the Internet to replace or augment private WAN services. The solution includes path quality monitoring (PQM) to track the available paths, stateful firewall with application fingerprinting to identify traffic flows, dynamic path selection (DPS) to identify the optimal path, and the route orchestrator to facilitate routing decisions without requiring admin intervention.

Use this guide to design new networks or optimize and upgrade existing networks. The guide is not intended as an exhaustive discussion of all options. It presents commonly recommended designs, features, and hardware.

Customer Use Cases

Branch networks are changing rapidly. The most pressing challenges include an increasing number of mobile and IoT devices, growing bandwidth requirements for business operations, and modern users who expect connectivity for work and personal use anywhere, any time.

IT teams responsible for managing distributed networks are facing staff reductions, even as organizations significantly increase network capabilities and reduce implementation times. Busy IT departments must improve service levels, reduce costs, manage an increasing workload and shift spending from capital expense to operating expense.

This guide discusses the following use cases:

- Employing artificial Intelligence to augment operator oversight with Smart Telemetry.

- Accommodate highly secure WAN communications using IPsec tunnels over an independent transport.
- Use Zero Trust Provisioning for all networking components in the branch.
- Institute switch stacking for simplified management, high availability, and scalability.
- Add link aggregation for high bandwidth, redundancy, and resiliency between switches and gateways.
- Incorporate wireless as the primary access method for branch employees.
- Invoke Zero Trust Security to secure the network against inside and outside attacks.

Aruba EdgeConnect SD-WAN Benefits

This section presents the benefits of SD-WAN and provides an overview of the EdgeConnect SD-WAN and EdgeConnect Branch solutions, with:

- A general description of SD-WAN
- Benefits of EdgeConnect SD-WAN
- Benefits of EdgeConnect Branch.

SD-WAN Overview

A Software-Defined Wide Area Network (SD-WAN) is a virtual WAN architecture that enables enterprises to combine different transport services such as MPLS, LTE, and broadband Internet. The expanded capabilities to connect users securely to applications over different transport services enhance overall network performance.

SD-WAN uses a centralized control function that can direct traffic securely and intelligently across the WAN to trusted SaaS and IaaS providers. This increases application performance and delivers a high-quality user experience, which increases business productivity and agility and reduces IT costs.

Unlike SD-WAN, the conventional router-centric model distributes the control function across all devices in the network and routes traffic solely based on TCP/IP addresses and ACLs. This traditional method is rigid, complex, inefficient, and labor-intensive. It is not cloud-friendly and results in a less positive user experience.

SD-WAN enables cloud-first enterprises to deliver superior quality of experience (QoEx) for application users. SD-WAN facilitates intelligent, automated, application-aware routing across the WAN.

With SD-WAN, applications are identified and classified to ensure they receive the appropriate level of service and security policy enforcement, in accordance with business needs.

Secure local Internet breakout of IaaS and SaaS application traffic from the branch provides the highest levels of cloud performance while securely protecting the enterprise from threats.

Key Value Propositions

Lower Cost

One of the major benefits of SD-WAN is the ability to shift traffic from high-cost circuits such as MPLS to an Internet circuit intelligently, and with less latency sensitivity. MPLS is much more expensive than a commodity Internet circuit, so relying less on MPLS can lead to significant cost savings over time. With SD-WAN, organizations do not need to increase MPLS circuit sizes to accommodate increasing traffic. In some cases, MPLS circuit sizes can be reduced and MPLS contracts can be renegotiated at a lower rate.

SD-WAN delivers optimal application performance under any network condition or change, including congestion and impairments. Because of continuous network monitoring and self-learning, the business-driven SD-WAN responds automatically in real-time to changes in the state of the network to address network congestion, brownouts, and transport outage conditions without requiring manual IT intervention, so users can always connect to applications. For example, if a WAN transport service or cloud security service experiences a performance impairment, the SD-WAN network automatically adapts to keep traffic flowing while maintaining compliance with business policies.

Centralized Management

Most SD-WAN solutions offer streamlined management via a central online portal. In traditional networking, administrators need to access routers remotely and configure the devices manually, requiring significant personnel overhead. The centralized management of an SD-WAN architecture allows organizations to make changes to thousands of devices with relative ease.

Modern Security Architecture

Secure Access Service Edge (SASE) architecture combines branch WAN edge functions, including SD-WAN, routing, segmentation, zone-based firewall, and WAN optimization, with comprehensive security services that are delivered and managed in the cloud.

SASE addresses the need to expand the network quickly as the number of remote users increases and enterprises continue to migrate applications to the cloud, while improving overall application performance and network security.

Traditionally, all application traffic from branch locations crossed over private MPLS services to the corporate data center for security inspection and verification. This architecture was appropriate when applications were hosted exclusively in the corporate data center. However, as applications and services migrate to the cloud, the traditional network architecture falls short. When applications are hosted in the corporate data center, all Internet-destined traffic must traverse the data center and corporate firewall before reaching its destination, causing diminished application performance and user experience.

As more remote workers connect directly to cloud applications, traditional perimeter-based security also has become insufficient. Transforming WAN and security architectures with SASE helps to ensure direct, secure access to applications and services across multi-cloud environments, regardless of location or the devices used to access them.

Components of SD-WAN

SD-WAN comprises four core components: centralized management, WAN virtualization, overlays and Edge devices.

- **Centralized Management** - Centralized management tools are used to define WAN topologies, orchestrate routing, and manage policy.
- **WAN Virtualization** - Devices can create Virtual WAN Point-to-Point IPsec tunnels using any underlying transport.

- **Overlays** - Overlay tunnels are malleable, enabling an organization to enforce policy orchestrated to all devices.
- **Edge Devices** - Management tools provide the capability to build overlays and onboard and manage devices.

Each component is detailed more fully in the EdgeConnect SD-WAN and EdgeConnect Branch sections.

Aruba EdgeConnect SD-WAN & SD-Branch

WAN Transport Design

Choosing the circuit types under which to run an SD-WAN deployment is a critical design step. Each circuit type has different strengths and weaknesses that can affect SD-WAN operation.

This section describes design decisions and best practices to help operators make the best circuit choice for Aruba's two SD-WAN solutions.

Commodity Internet

Commodity Internet is the most common transport type for SD-WAN deployments due to its low cost and high bandwidth.

The main benefit of choosing commodity Internet is full IP reachability to any other location using the Internet as a transport type. The downside is the limit or lack of guaranteed performance from the transports, especially when traffic must cross numerous backbone Internet providers.

MPLS

MPLS circuits provide private connectivity between locations across a service provider network and ensure more predictable SLA-backed performance. Although MPLS can serve as a viable transport for SD-WAN, MPLS is generally the most expensive form of transport. In most SD-WAN solutions, MPLS is phased out of the network and replaced with commodity Internet, relying on the SD-WAN technology to provide the SLA.

If the MPLS network is used to host services such as provider SIP circuits, firewalls, or cloud ingress points, running SD-WAN on the circuit may present routing challenges for accessing the services. Carefully consider the routing design to accommodate service requirements and access flows. Route-leaking between the underlay and overlay or hairpinning traffic through a hub may be required.

As SD-WAN is adopted, it is typical to phase out MPLS circuits in favor of the higher speeds and lower costs Internet circuits can offer.

3G/4G/5G

Mobile carrier data is traditionally used as a path of last resort at many locations to provide last-mile resiliency if the physical circuits at a location are unavailable.

Because mobile circuits have lower throughput and data caps, they generally are not recommended for use as primary circuits. However, some modern 5G providers now offer plans with no data caps and high speeds that may be appropriate to consider for use as a primary circuit.

Private Layer 2

Private Layer 2 circuit types include VPLS, Metro-Ethernet, and Pseudo Wire. These circuit types are commonly found in DCI environments or in interconnections in large campus environments. Though they are less common for large WAN transports they are still often deployed.

Both Aruba EdgeConnect SD-Branch and EdgeConnect SD-WAN support layer 2 WAN transports, but these designs require additional consideration. Consult with your local Aruba Solution Architect if deploying SD-WAN over this transport type.

Starlink / Low Orbit

Low Orbit WAN circuits, such as SpaceX Starlink, provide a new RF-based WAN solution. They are generally high-speed, low-latency broadband Internet circuits used in remote or rural locations.

For network operators in these areas, it can be a challenge obtaining other WAN circuit types with adequate bandwidth and SLA performance required for their business needs. The emerging low orbit WAN circuit is a worthy consideration.

Recommended Pairings

Many SD-WAN implementations aim to remove private circuits from their environment completely. Commodity Internet circuits provide large amounts of bandwidth cost effectively and SD-WAN technology now enables networks to steer traffic to the correct circuit, improving the performance of commodity Internet to better meet business requirements.

A combination of multiple high speed Internet circuits with a 4G/5G backup is the most desirable end state for circuit pairings. It is common to see dual, or even triple, Internet circuits at the data center, with dual Internet or single Internet with mobile backup at the branches.

Businesses with stricter transport SLAs or complex established MPLS networks may prefer keeping provider MPLS in their infrastructure. In this case, paring MPLS, commodity Internet, and mobile backup makes sense. The business can choose to reduce the bandwidth of their MPLS circuits gradually as more traffic is shifted to overlays built on commodity Internet.

Last-Mile Pairing Considerations

When pairing WAN circuits for a site, the need to provide last-mile resiliency cannot be overlooked. Avoid the common practice of bundling circuits together for the last mile so ingress occurs at the same physical location. The common ingress point can be susceptible to a layer 1 failure event, or backhoe event, taking out all WAN circuits at the location.

Some customers use 5G circuits as the primary WAN transport, along with commodity Internet. Service providers now offer unlimited 5G plans that work well with a wired commodity Internet circuit to provide last mile resiliency.

Zero Touch Provisioning

SD-WAN solutions enable the deployment of devices “out of the box” with automated connection to a single cloud platform for management. This process is known as Zero Touch Provisioning (ZTP).

To benefit from ZTP, branch networks must be standardized so templates can be applied for consistent, universal configuration across multiple locations.

ZTP can be beneficial even if some sites have unique topologies that cannot be configured remotely. However, the true value lies in the ability to use templates to deploy all topologies universally.

Hub Designs

SD-WAN fabrics generally have *hub* locations, typically data centers or other sites where applications reside. The hub location is responsible for:

- WAN site reachability to applications residing in the hub
- Aggregation of IPsec tunnels, for hub-and-spoke topologies
- Serving as a regional routing hub for large deployments
- Centralized security services, as needed
- Cloud onramp instead of brick-and-mortar hubs.

This section reviews key considerations when determining the set-up of hub locations for SD-WAN topologies.

Appliance Deployment Models

When installing SD-WAN appliances in the data center, two deployment methods are commonly used:

- Inline deployment
- Out-of-path deployment.

Inline Deployment

For inline deployments, the SD-WAN router serves as the edge router, terminating WAN circuits directly. This is the preferred deployment method with two gateways acting as an active/standby pair. This method works well in simple environments with one or two SD-WAN appliances that do not require horizontal scale. It also works well in greenfield environments when support for migrating a legacy WAN is not required.

The diagram below illustrates an in-path deployment.




Figure 1: Inline Deployment

Out-of-Path Deployments

For out-of-path deployments, the SD-WAN appliance does not physically terminate the WAN circuits. Instead, traffic is redirected to the SD-WAN appliances, generally using routing protocols or a redirect protocol. This method is not the recommended starting design, and should be used only if either of two specific requirements are needed:

- The edge router must be kept in place and cannot be replaced with a gateway. This may be desirable when the edge router is doing other complex functions, such as third party IPsec tunnel termination.
- More WAN bandwidth is needed than a simple two-box, in-path deployment can provide. The gateways can scale past two, providing horizontal scale, if deployed out-of-path.

The diagram below illustrates an out-of-path deployment.




Figure 2: Out of Path

Branch Designs

A branch is defined as a small remote location that does not host applications. Restaurant chains and retail outlets are examples of businesses with hundreds of branches spread over large geographical areas. For remote locations with larger footprints, refer to the Aruba [small campus reference architecture](#).

This section presents key considerations for designing branches, with a focus on SD-WAN appliance integration.

SD-WAN appliances at branch locations are almost always deployed as edge routers, in-path, replacing legacy routers and firewalls. They physically terminate the WAN circuits and handle all route peering with the providers. These branch routers may be deployed as the default gateway for the VLANs at the branch or they may participate in routing with the branch LAN.

LAN-side design for switching and wireless infrastructure is covered in the [Campus Design](#) section of the VSG. For specific branch designs, refer to the reference architectures for EdgeConnect SD-WAN and EdgeConnect SD-Branch.

Internet Egress

As more business resources are migrated to the Internet, designing Internet egress becomes more important. Many corporate resources are migrating to SaaS solutions, such as Office 365, accessed over the Internet. The change in application location requires network architects to rework traditional Internet egress designs.

The conversion of all traffic to SD-WAN enables more flexible and direct Internet egress designs, without compromising security. With application recognition and SD-LAN's enhanced policy application, different Internet egress policies can be applied to different traffic.

Trusted business applications can be sent directly to the Internet or sent through on-box firewall features of the SD-WAN appliance. Traffic also can be sent to a cloud firewall service for more detailed inspection and centralized firewalling. Finally, select traffic can still be backhauled to a data center for more specific in-house inspection as needed.

Regardless of the location in the network, it is critical for clients to use the geographical-appropriate DNS. Many Internet resources have geolocated onramps for their services. For example, Microsoft has numerous onramps to the core software services based on geographic locations. If you are located in Seattle and request Office 365 DNS entries, but your DNS server is in New York, you receive IP addresses for the New York on-ramp, which creates suboptimal traffic flows.

Overlay Components

The specific interworkings and features of an overlay are highly dependent on the SD-WAN solution. This section outlines how overlays work generically. More solution-specific detail is provided in the EdgeConnect SD-WAN and EdgeConnect SD-Branch sections.

Overlays, which are typically IPsec tunnels, normalize all transports within a WAN by building logical connections between the SD-WAN appliances. After the tunnels are established, the SD-WAN solution begins to monitor the tunnels to obtain performance metrics and transport health.

The goal of an overlay is to ensure application stability and reliability, using the following mechanisms:

- **Traffic Identification** - SD-WAN appliances are configured to differentiate traffic by type. The identification can be simple matching of DSCP markings or use of L3/L4 ACLs. More advanced application recognition engines can be employed, based on the capability of the SD-WAN solution.
- **Traffic Policy** - The operator creates policy for different types or groups of traffic and assigns a specific SLA. The SD-WAN solution then steers the traffic over specific transports to ensure the best path is selected based on the assigned policy and SLA.
- **Combat WAN Transport Instability** - SD-WAN appliances can apply forward error correction (FEC) to traffic, which enables the network to recover from packet loss caused by a variety of network layer conditions, such as queue overflows, constrained bandwidth links, or long equipment delays in the carrier network. Packet-level FEC adds a parity packet for every “N” number of packets in each block, with “N” being 2, 4, or 8. The parity packet enables the far-end SD-WAN appliance to reconstitute lost packets before the packets are delivered to TCP, UDP, or other transport layers. This avoids the loss of information in the case of UDP traffic, as well as the delays associated with multiple round-trip retransmissions or the performance impacts of congestion avoidance mechanisms in the case of TCP.
- **Topology Choice** - Topology dictates where IPsec tunnels are built and can affect the scalability of the SD-WAN solution. There are two main topology types: hub-and-spoke and full mesh. In hub-and-spoke, IPsec tunnels are built only between the spokes and the hubs. This minimizes the number of IPsec tunnels, allowing for higher scale, but also creates suboptimal traffic flows between the spokes, requiring them to transit the hub. With a full mesh topology, the spokes can all communicate directly, but this creates a very high number of IPsec tunnels which limits the scalability of the topology.

The best practice is to use hub-and-spoke topologies for overlays that do not require optimized spoke-to-spoke routing. The main traffic type that requires a full mesh topology is “real-time” to allow latency-sensitive applications to route directly between spokes.

Cloud Integration

SD-WAN has changed the way many architects connect their networks to cloud providers such as AWS, Azure, or GPC. EdgeConnect SD-WAN and EdgeConnect SD-Branch support numerous deployment models to connect to cloud resources. EdgeConnect deployed directly into the IaaS provider is most common and generally recommended, which is further discussed below. For more information on how to integrate SD-WAN with a multi-cloud solution such as Megaport or Alkira, contact your Aruba account team.

Aruba SD-WAN solutions help to automate much of the deployment and reduce complexity, as described in the deployment guides. An outline of the underlying solution architecture is illustrated below.

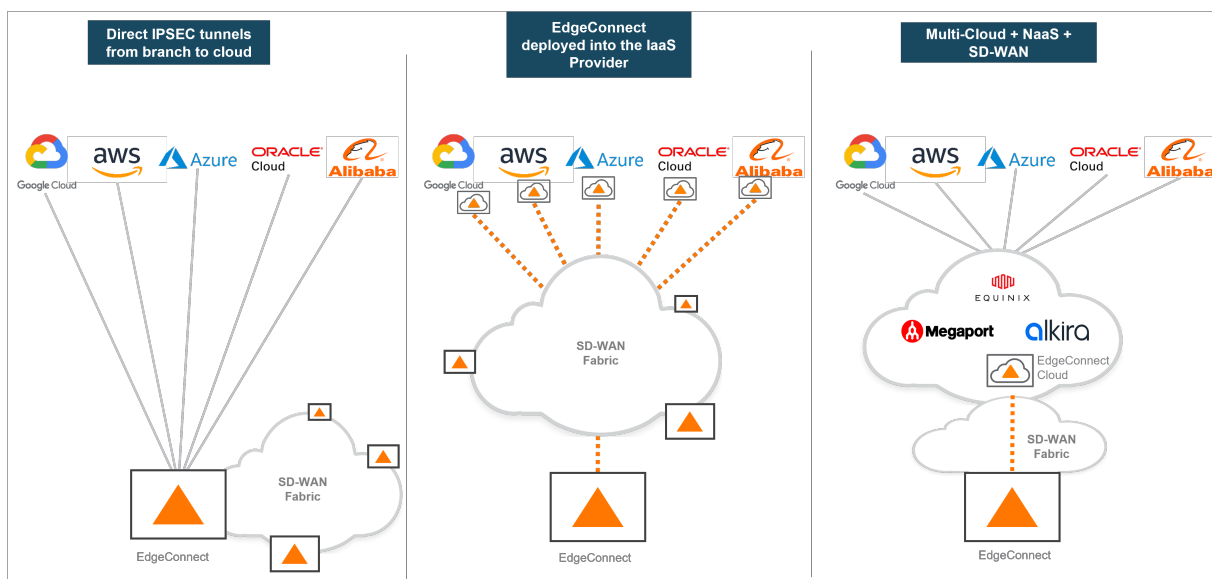


Figure 3: Cloud Integrations

Microsoft Azure

To integrate with Microsoft Azure Aruba recommends customers deploy a vWAN solution to interconnect their workload Virtual Networks (VNET). An SD-WAN VNET is then used to provide connectivity between the Azure VNETs and the SD-WAN fabric. More information on this method can be found in [Microsoft's documentation](#).

Within Azure a pair of redundant EdgeConnect appliances are deployed via either Aruba Orchestrator for SD-WAN or Aruba Central for SD-Branch. On the LAN side EdgeConnect appliances will peer BGP to the VNET hub. SD-WAN fabric routes will be advertised to the hub, and workload VNET routes will be learned. On the WAN side the interfaces will connect to the Azure Internet Gateway to obtain internet reachability. This will allow the devices to register with Orchestrator or Central and establish SD-WAN IPSEC tunnels and routing adjacencies.

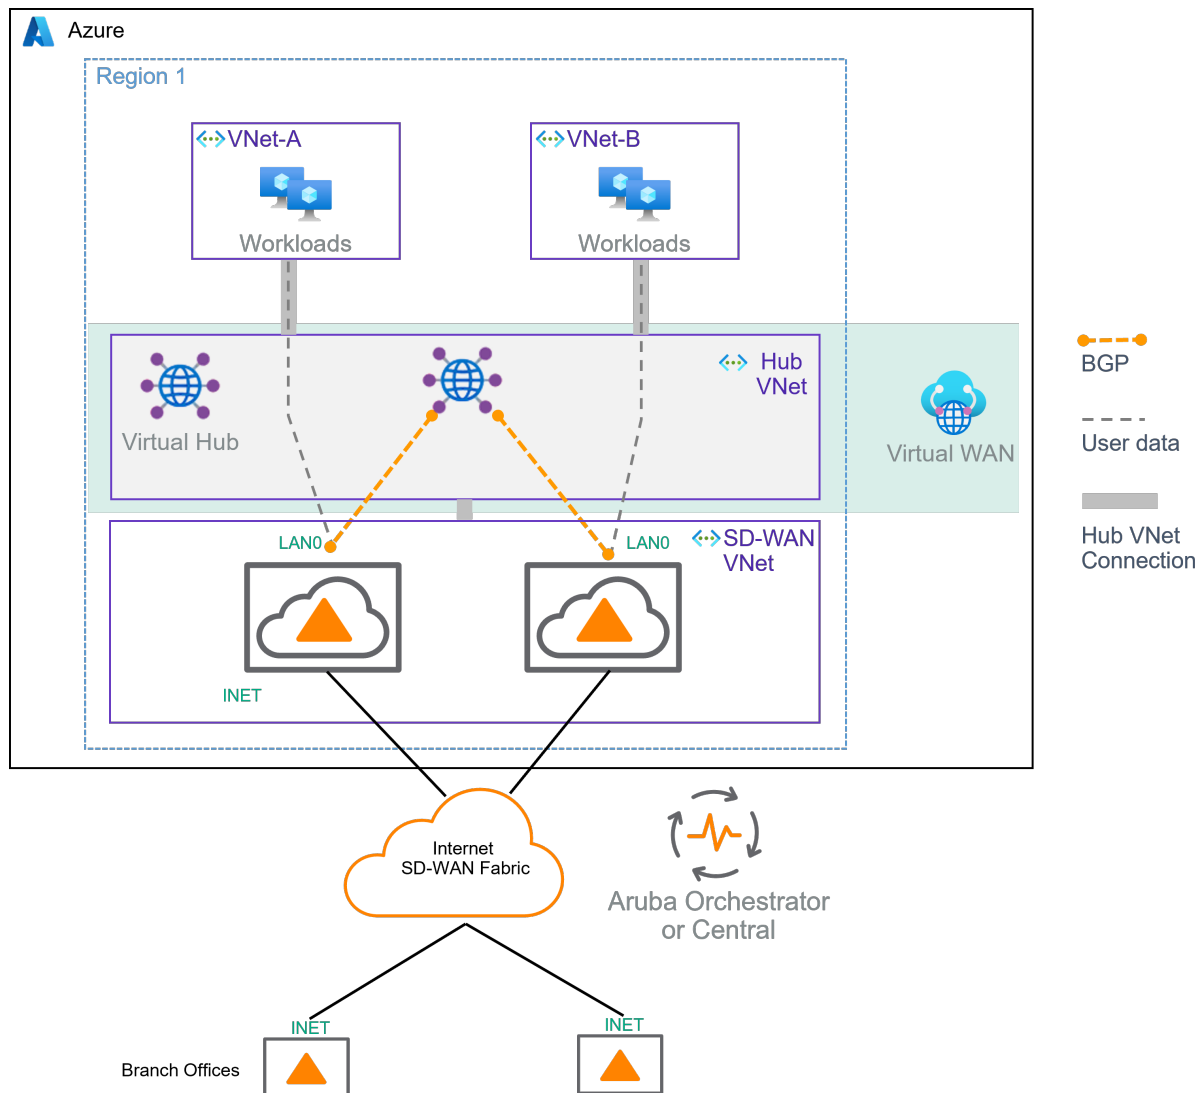


Figure 4: Cloud Integrations

NOTE:

While the EdgeConnect can be deployed directly into vWAN, it is not the recommended approach due to current caveats, including lack of support for Express Route connections.

Amazon Web Services

To integrate with Amazon Web Services (AWS) Aruba recommends customers deploy a transit gateway solution to interconnect their workload Virtual Private Cloud (VPC). An SD-WAN VPC is then used to provide connectivity between the AWS VPCs and the SD-WAN fabric. More information on this method can be found in [Amazon's documentation](#).

Within AWS a pair of redundant EdgeConnect appliances are deployed via either Aruba Orchestrator for SD-WAN or Aruba Central for SD-Branch. On the LAN side EdgeConnect appliances will peer BGP to the Transit Gateway (TGW). SD-WAN fabric routes will be advertised to the TGW, and workload VPC routes will be learned. On the WAN side the interfaces will connect to the AWS Internet Gateway to obtain internet reachability. This will allow the devices to register with Orchestrator or Central and establish SD-WAN IPSEC tunnels and routing adjacencies.

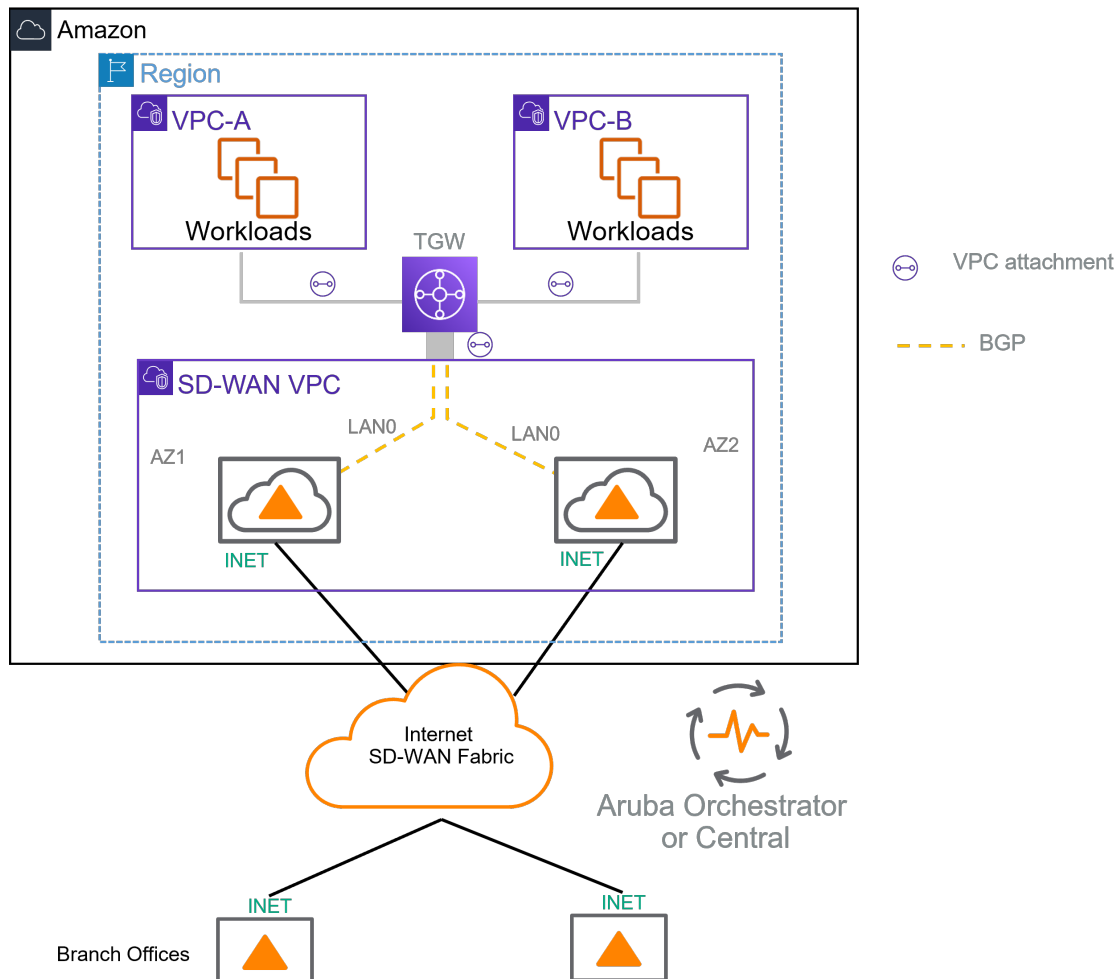


Figure 5: Cloud Integrations

EdgeConnect SD-WAN Solution

Fundamentals

This section introduces the Aruba EdgeConnect SD-WAN solution with descriptions of the Aruba Orchestrator, EdgeConnect SD-WAN appliances, and some key features.

Aruba Orchestrator

Aruba Orchestrator provides centralized policy management, monitoring, and reporting capabilities for the SD-WAN platform.

Orchestrator has three flexible models for deployment:

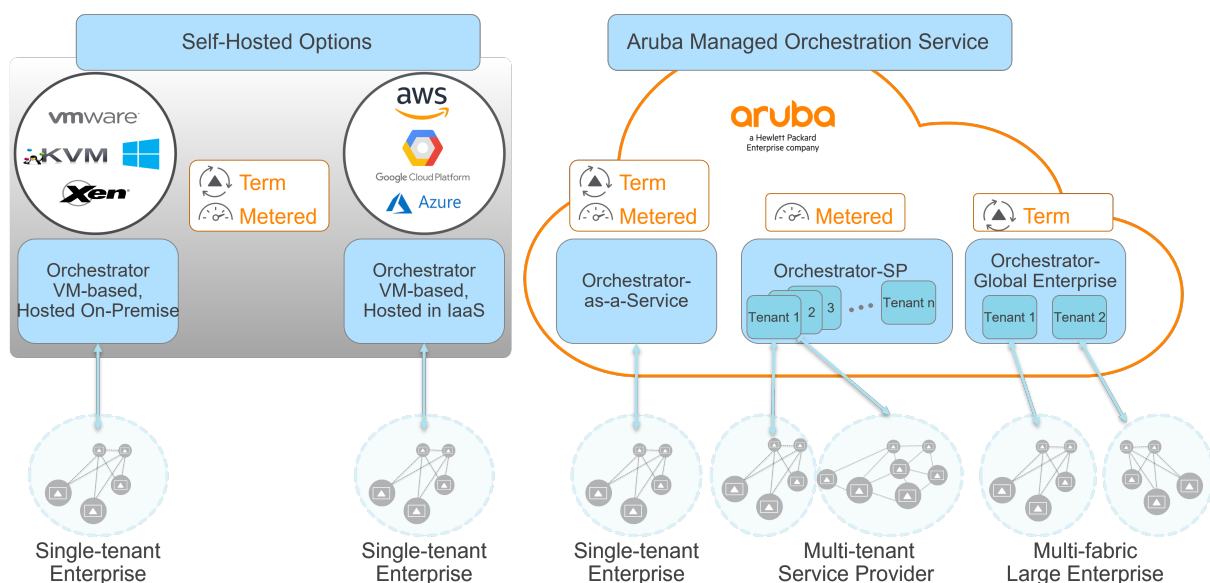
- On-premise VM deployment
- Customer-managed cloud deployment
- Aruba-hosted SaaS (software-as-a-service) deployment.

When EdgeConnect SD-WAN is hosted as-a-service, Aruba manages and maintains the platform, eliminating the need for additional capital expenditures. This option offers maximum flexibility, ease of deployment, full use and customization of Orchestrator features, and long-term sustainability.

Use Orchestrator to configure and monitor application quality of service and security policies for thousands of sites rapidly from a centralized dashboard with single-screen administration.

With Orchestrator's centralized configuration, users have a single-pane dashboard for real-time monitoring, alerting, and visibility into the network, as well as access to a detailed historical log of reporting and analytics for better understanding of business needs related to the SD-WAN fabric.

The dashboard can be customized in several formats to provide high-level geographic health overviews, granular analysis of live traffic flows, an overview of appliances connected to the network, and more.



Best practice is to use the SaaS option because it decreases the operational complexity of the deployment. For information on the on-premise options, refer to the Aruba Edge Connect Enterprise [User Guides](#).

EdgeConnect SD-WAN Appliances

EdgeConnect SD-WAN physical or virtual SD-WAN appliances (supporting any common hypervisors and public clouds) are deployed in branch offices and data centers to create a secure, virtual WAN overlay. This enables organizations to migrate to broadband WAN at their own pace, whether site-by-site or using a hybrid WAN approach that leverages MPLS and broadband Internet connectivity.

	EdgeConnect SD-WAN 10104	EdgeConnect SD-WAN XS	EdgeConnect SD-WAN S-P	EdgeConnect SD-WAN M-H	EdgeConnect SD-WAN L-H	EdgeConnect SD-WAN XL-H
Model	EC-10104	EC-XS	EC-S-P	EC-M-H	EC-L-H	EC-XL-H
Typical Deploy- ment	Small Branch / Home Office	Small Branch	Large Branch	Head Office/DC Large Hub	Data Center Large Hub	Data Center Large Hub
Typical WAN Band- width	2-500 Mbps	2-1000 Mbps	10-3000 Mbps	50-5000 Mbps	2-10 Gbps	2-10 Gbps
Simultaneous Conne- ctions	256,000	256,000	256,000	2,000,000	2,000,000	2,000,000
Recommen- WAN Boost up to	200 Mbps	250 Mbps	500 Mbps	1 Gbps	1 Gbps	5 Gbps
Redundant / FRUs	No	No	SSD and Power (AC or DC)	SSD and Power	SSD and Power	SSD, NVMe, Power
Data Path Interfaces	4 x RJ45 10/100/1000	4 x RJ45 10/100/100	8 x RJ45 4 x 1/10G Optical	8 x RJ45 4 x 1/10G Optical	6 x 1/10G Optical	6 x 1/10/25G Optical

NOTE:

WAN bandwidth assumes bidirectional traffic (symmetric up-link and down-link). For total WAN throughput (Rx+Tx), multiple these numbers by 2. For the most up-to-date SKU information, refer to the [EdgeConnect SD-WAN SD-WAN Data Sheet](#).

First-Packet iQ

First Packet iQ provides robust capabilities for Application Visibility and Control (AVC) that simplify establishment of route policies by application or domain. Rapid classification is critical for making traffic forwarding decisions.

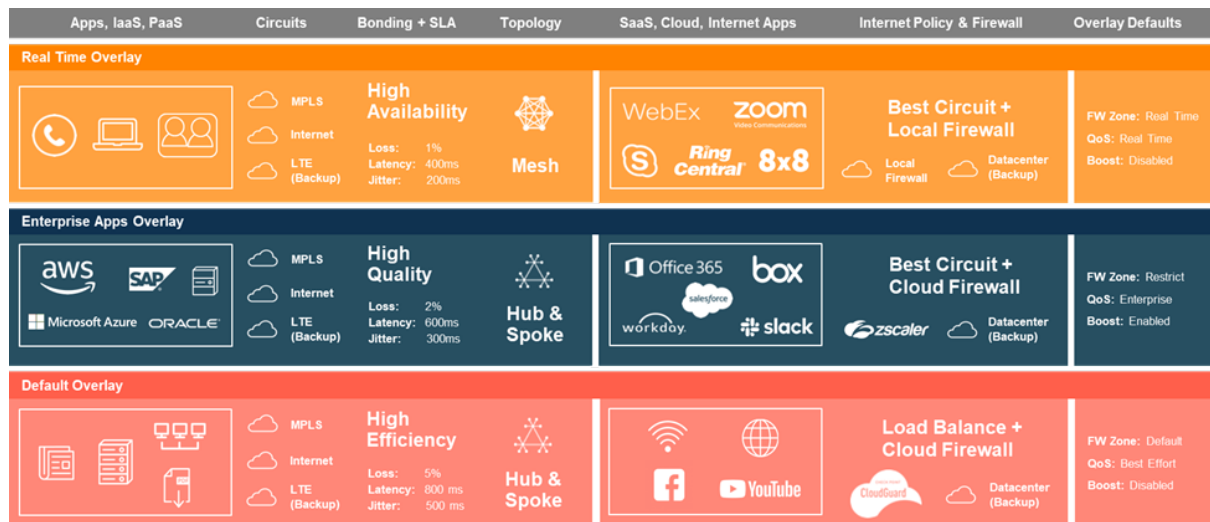
When different groups of applications are mapped to Business Intent Overlays, the decision on which overlay to place the flow must be correct on the first packet, or application performance will suffer.

Unique in the industry, First-Packet iQ goes above and beyond traditional Deep Packet Inspection techniques that typically require several packets of HTTP or HTTPS to identify applications accurately. Using First-Packet iQ, EdgeConnect can immediately and efficiently steer flows to the best route and avoid the need for after-the-fact route remediation. With Silver Peak's unique technology to "map the Internet," EdgeConnect can provide the granular Internet breakout policies as shown above.

Business Intent Overlays

The Aruba EdgeConnect SD-WAN platform enables enterprises to create multiple application-specific WAN overlays. Each overlay, or Business Intent Overlay (BIO), specifies priority and quality of service requirements for application groups based on business requirements or intent.

With these policy definitions in place, EdgeConnect automates traffic steering on an end-to-end basis across all underlying WAN transport services including MPLS, broadband Internet and 4G/LTE, providing the ability to deliver an application Quality of Experience significantly better than the underlying transport services can deliver individually. Each BIO has its own link bonding policy that specifies the underlay transports the BIO will use, the service level including path conditioning, and topology (full mesh, hub-and-spoke, or regional hub-and-spoke).



Each BIO has settings that include Traffic Class/QoS, Firewall Zone, and the option to enable Boost (WAN Optimization).

Link Bonding, Path Conditioning, and Dynamic Path Control

SD-WAN uses multiple underlay transport networks to provide applications the best possible virtual overlay network experience. The configuration of each BIO contains two primary sections:

- SD-WAN Traffic to Internal Subnets (i.e., EdgeConnect-to-EdgeConnect)
- Breakout Traffic to Internet and Cloud Services (i.e., EdgeConnect-to-Internet)


Each section includes a Link Bonding Policy that specifies the underlay transports to use (specified as a label), how to bond the transports, and service level quality metric that includes the amount of Forward Error Correction to be applied.

As an IPsec-based overlay, performance of the hybrid WAN is optimized while maintaining complete independence of the underlying infrastructure. Optimal path choices are based on application requirements, geolocation, and packet-level determination of link quality, including line characteristics such as delay, loss, and jitter.

EdgeConnect's Path Conditioning includes both Forward Error Correction (FEC) and Packet Order Correction (POC).


Forward Error Correction (FEC)

Packets lost in transmit across the WAN are rebuilt



Packet Order Correction (POC)

Packets delivered out-of-order across the WAN are reordered into the correct order



Forward Error Correction reconstructs lost packets to avoid the need for TCP retransmission, substantially increasing the performance of Internet links. The ratio of FEC packets to data packets is configurable depending on business criticality and real-time requirements for the application.

Packet Order Correction (POC) algorithms reorder packets that arrive at their destination out of order. This is a common occurrence when load balancing across different service providers' networks. With FEC and POC, EdgeConnect can make internet connections perform as well as or better than private lines.

Boost

Aruba EdgeConnect takes SD-WAN performance even further for latency-sensitive applications and applications that transfer large amounts of data across the WAN. With the optional Unity Boost software performance pack, EdgeConnect integrates Aruba's field-proven WAN Optimization features in a single SD-WAN solution. When Boost is integrated with SD-WAN, it can be provisioned after the initial SD-WAN roll-out for a given end-customer without the need to service-chain an additional physical appliance or Virtual Network Function (VNF) for the sole purpose of WAN Optimization.

Boost is enabled at the Business Intent Overlay level so it can be activated for critical applications without running on less-sensitive applications.

Boost includes:

- Application Acceleration (latency mitigation) to improve application response times over distance
- Data reduction (compression and deduplication) to eliminate the transmission of redundant data. This capability is also referred to as "Network Memory".

Latency Mitigation

TCP Acceleration uses techniques such as selective acknowledgments, window scaling, and message segment size adjustment to mitigate poor performance on high-latency links.

Data Reduction

Data reduction technology addresses limited bandwidth using advanced fingerprinting algorithms that examine all incoming and outgoing WAN traffic. Network memory localizes information and transmits only modifications between Boost-enabled SD-WAN nodes.

IP Header Compression is the process of compressing excess protocol headers before transmitting them on a link and uncompressing them to their original state at the other end.

Payload Compression uses algorithms to identify relatively short byte sequences that are repeated frequently. The sequences are replaced with shorter segments of code to reduce the size of transmitted data.

Licensing

EdgeConnect

EdgeConnect software, applied to EdgeConnect hardware, virtual, or cloud instances, can be purchased using one of two options:

- Subscription-term basis for enterprise end-customers
- Monthly-metered basis for service providers.

In both cases, the appliance license is based on the WAN-side, bidirectional bandwidth of the appliance instance (not LAN side).

EdgeConnect Subscription: Subscription licensing is term-based and renewable at 1, 2, 3, 4, and 5 years. The appliance license is based on the composite WAN-side bandwidth and has the following bandwidth tiers:

EdgeConnect License	EdgeConnect Aggregate WAN Provisioned Bandwidth (Mbps)	Description: Per EC Instance, Term (n) is 1,2,3,4 or 5 years
EC-BW-20-nY	20 Mbps	EC BW License, 20 Mbps Bandwidth, n Years
EC-BW-50-nY	50 Mbps	EC BW License, 50 Mbps Bandwidth, n Years
EC-BW-100-nY	100 Mbps	EC BW License, 100 Mbps Bandwidth, n Years
EC-BW-200-nY	200 Mbps	EC BW License, 200 Mbps Bandwidth, n Years
EC-BW-500-nY	500 Mbps	EC BW License, 500 Mbps Bandwidth, n Years
EC-BW-1G-nY	1 Gbps	EC BW License, 1 Gbps Bandwidth, n Years

EdgeConnect License	EdgeConnect Aggregate WAN Provisioned Bandwidth (Mbps)	Description: Per EC Instance, Term (n) is 1,2,3,4 or 5 years
EC-BW-2G-nY	2 Gbps	EC BW License, 2 Gbps Bandwidth, n Years
EC-BW-UL-nY	Unlimited Bandwidth	EC BW License, Unlimited Bandwidth, n Years

Boost

Boost is an optional WAN optimization performance pack for EdgeConnect.

NOTE:

Boost is integrated in the same software image and is enabled with a simple “checkbox” provisioning action.

With Enterprise, Boost is licensed in units of 100 Mbps of WAN optimization and can be deployed flexibly to sites that require application acceleration.

It is important to note that with the data reduction feature of Boost, a site’s application traffic may be several times the provisioned WAN bandwidth.

Advanced Security License

The EdgeConnect Platform supports IDS that requires a new software feature license called Advanced Security. This license type comes in two options: “standard” and “unlimited”. For IDS/IPS, the difference between standard and unlimited is the maximum throughput supported by the inspection engine, and is specific for each EdgeConnect model and software release. The Advanced Security feature license is optional. In the future, other security-related features may be tied to the Advanced Security License.

- Boost
- Optional
 - Per Mbps

Unity Boost (WAN Optimization)

- Optional add-on feature
- Shareable pool across EC’s
- Dynamically adjustable

Advanced Security

- Optional add-on feature
- Attached to appliance
- Includes IDS

Advanced Security options:

- Optional
- Standard
- Unlimited

Tiered Bandwidth License (required)

Routing
Zone-based FW
Advanced segmentation*
Path conditioning
Network analytics

Business Intent Overlays
Cloud Services Orchestration
Intelligent Internet Breakout
1st Packet iQ
Zero Touch Provisioning

Multiple BW tiers available:

- 20Mbps
- 50Mbps
- 100Mbps
- 200Mbps
- 500Mbps
- 1Gbps
- 2Gbps
- >2Gbps

EdgeConnect SD-WAN Overlay Design

Overlays in the EdgeConnect SD-WAN solution are built using IPsec tunnels. On top of the IPsec tunnels, traffic is placed into a Business Intent Overlay (BIO) that directs the traffic using specific guidelines and settings. The BIO includes many different components (interfaces, policies, SLAs, etc.). Design considerations are outlined below.

More information about BIOs and link bonding can be found in the [EdgeConnect SD-WAN BIO and Link Bonding Policy](#) technical white paper.

The above video summarizes the information in this chapter of the VSG.

Traffic Classification

Classifications are assigned in each BIO. The SD-WAN appliance can be assigned to recognize specific applications or use an overlay access list based on DSCP markings, L3/L4 ACLs, and the built-in DPI Engine referred to as First Packet IQ.

In the design phase, consider how to coordinate the set-up of BIOs for different types of traffic. Application identification is recommended, with DSCP or ACLs used as fallback.

Review the default overlay ACLs and modify as needed. Also use application groups, which bundle common applications together, when appropriate.

Topology

Topology dictates where IPsec tunnels are built and determines the scalability of the SD-WAN solution. Use of hub-and-spoke topologies is recommended for BIOs that do not require optimized spoke-to-spoke routing. Generally, the only BIO configured as mesh is real-time to allow latency sensitive applications to route directly between spokes.

Regional Routing

Regions can be defined within the SD-WAN topology to accommodate larger scale networks. This enables using a hub-and-spoke or full-mesh topology within a region with a full mesh of IPsec tunnels between the regions. This is an advanced design generally used only for larger scale networks. For regional routing design assistance, contact the Aruba account team.

Interface Selection

Primary and backup interfaces are defined in the BIO. Interfaces in the primary interfaces list are frequently probed for health and are all used to transport traffic, based on the bonding policy.

Backup interfaces are used only when all primary circuits are down or do not meet the configured SLAs. Backup interfaces pass only enough traffic to keep the IPsec tunnel running and do not have health probes.

Best practice is to place all commodity Internet and MPLS interfaces on the primary interface list and use the backup list for metered interfaces such as mobile circuits.

Service Level Objectives

BIOs can be defined with Service Level Objectives (SLOs) to remove them from the available interfaces if they fall out of a policy range. High Availability and High Quality link bonding policies exclude links based on overlay brownout. For this reason, setting SLOs is not needed for most environments. If a BIO-level SLO is required, it should be set at a high value to let the link bonding policy manage the overlay.

Suggested values for BIO-level SLOs are:

Real-Time BIOs:

- Latency - 250ms
- Loss - 10%
- Jitter - 50ms

TCP-Based BIOs

- Latency - 500ms
- Loss - 10%
- Jitter - 50ms

Link Bonding Policy

Link Bonding policies define how traffic is distributed across eligible links and how forward error correction (FEC) is applied to the traffic. Four default link bonding policies enable the creation of custom policies.

The default real-time BIO uses the High Availability link bonding policy and the remainder of the BIOs use the High Quality bonding policy. Characteristics of each bonding policy are provided below.

High Availability

High availability chooses the best performing path, uses the path until it is near full, then waterfalls traffic onto the next best performing path. All traffic receives 1:1 FEC when a copy of the packet is placed on another transport. High availability link bonding policy type should be used only for real-time traffic, since it renders the effective bandwidth to 50%.

High Quality

High quality policy chooses the best performing path, uses the path until it is near full, then waterfalls traffic onto the next best performing path. Adaptive FEC is used to provide parity packets only if there is degradation of the circuit. High quality link bonding policy should be used as the default selection for all non-real-time traffic types.

High Throughput

High throughput policy load-balances packets across all transports performing below the SLO defined in the BIO. Adaptive FEC is used to provide parity packets only if there is degradation of the circuit. This link bonding policy is used only in unique circumstances.

High Efficiency

High efficiency policy load-balances packets across all transports performing below the SLO defined in the BIO. No FEC is used in this bonding policy. This link bonding policy is used only in unique circumstances.

QOS, Security, and Optimization

QOS

Different applications have different quality of service (QOS) and end-user experience requirements. For example, voice and video traffic requires zero packet loss and extremely low delay while file transfers require large amounts of bandwidth, but some amount of delay is acceptable.

Silver Peak enables network designers to define logical or virtual WAN overlays that reflect application QOS requirements relevant to the business.

EdgeConnect maps applications to the appropriate overlay, which enables the SD-WAN to optimize routing decisions automatically. EdgeConnect continuously monitors bonded tunnels and physical WAN links, factoring real-time data about delay, jitter, and packet loss to make intelligent routing decisions. The Silver Peak SD-WAN learns and adapts to optimize and dynamically change paths if necessary, based on actual performance with no application disruption. When link conditions change, the SD-WAN can revert to the original path.

EdgeConnect performs both egress and ingress traffic shaping. IT staff can program minimum and maximum bandwidth limits on the egress traffic to shape the engine by traffic class to ensure that no single application consumes all the WAN bandwidth.

Ingress shaping can be programmed to ensure that low-priority traffic does not override higher priority traffic. For example, video streaming or social media applications can be prevented from compromising the performance of higher priority business applications.

Within the EdgeConnect SD-WAN solution, two main components manage QoS:

- The shaper provides a simplified way to configure QoS (Quality of Service) globally on the appliances. It shapes traffic by allocating bandwidth as a percentage of the system bandwidth. The shaper's parameters are organized into ten traffic classes. Four traffic classes are preconfigured and named *real-time*, *interactive*, *default*, and *best effort*.
- The BIO is used to trust or remark traffic. It determines how the packets are mapped to their assigned traffic class.

When designing the branch, it is critical to consider the overall QoS policy and determine the QoS markings to use and how to map the traffic classes.

Optimization

The EdgeConnect enterprise solution delivers robust WAN optimization, as described in the EdgeConnect SD-WAN Solution Overview section. Optimization should be applied individually for each BIO and applied only to BIOs that require or would benefit from using boost. Use optimization for BIOs when caching is needed or high latency (100ms+) is expected.

Breakout Traffic to Internet and Cloud Services

As enterprises migrate more applications to the cloud, changing traffic patterns drive the need to transform wide area network (WAN) and security architectures.

When applications were solely hosted in enterprise data centers, traffic from branch locations was backhauled to the data center over private circuits. In today's modern, cloud-first enterprise, applications are hosted everywhere – at the data center, in public and private clouds – and delivered by myriad Software-as-a-Service (SaaS) providers.

Users can access applications from anywhere, from any device, and across diverse WAN transports including broadband Internet. The increasing variety of access complicates IT development of effective security models. The dissolving enterprise security perimeter leads to an expanded and varied attack surface, significantly increasing the critical need for advanced data and threat protection services to mitigate exposure to threats.

Aruba responded to these challenges with the Secure Access Service Edge (SASE) architecture that combines the strength of SD-WAN with cloud-delivered security services.

The Aruba EdgeConnect Platform delivers industry-leading integration and automation with cloud-delivered security services.

Security Partner	Service	Integration
Checkpoint	Harmony Connect	API Integration
Forcepoint	Cloud Security	Service Orchestration

Security Partner	Service	Integration
iBoss	Cloud Security	Service Orchestration
McAfee	Unified Cloud Edge	Service Orchestration
Netskope	NewEdge SWG	Service Orchestration
Palo Alto	Prisma Access	Service Orchestration
Symantec	Web Security Service	Service Orchestration
Zscaler	Zscaler Internet Access	API Integration

Aruba Orchestrator automates the establishment of data plane IPsec tunnels between EdgeConnect appliances and their nearest security service cloud instances. This is performed with full API integration where connection establishment is single-ended and the Orchestrator arranges the connection to partner services, or with Service Orchestration where the Orchestrator ingests a list of partner-service access points then builds the connections.

Internet egress is configured on a per-BIO basis. General recommendations for designing Internet egress are outlined below. Recommendations may vary based on specific security requirements.

- Trusted traffic should be sent directly to the Internet, with minimal security inspection to meet corporate policy, using the built-in zone-based firewall if needed. Backhauling egress traffic over the data center should be used only if local Internet is not available.
- Untrusted, high-risk, or sensitive traffic should be sent to a cloud security solution for additional inspection. Backhauling this egress traffic over the data center should be performed only if local Internet is not available.
- Guest traffic should be sent directly to the Internet, with minimal security inspection. Due to the high cost of backhauling traffic, backhauling should not be performed if a local Internet outage occurs. Exceptions may be needed for important guest traffic in specific environments.
- Traditional cellular connections should be used only for backup Internet egress, and only for specific high-impact BIOs. If modern 5G cellular backups with high bandwidth and no data caps are used, consider using them as part of the primary Internet egress selection.

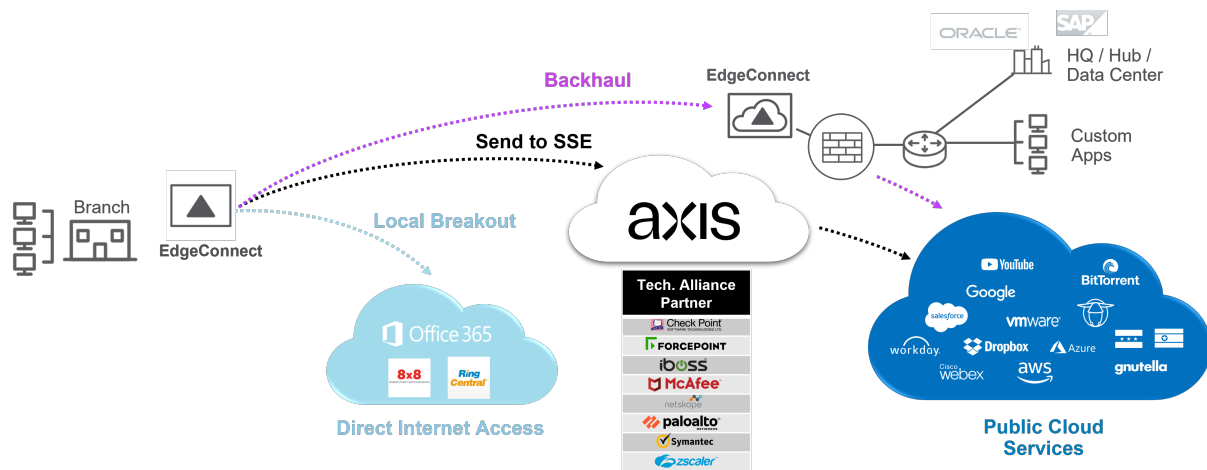


Figure 6: Internet Breakout

Zone Based Firewall

Zone based firewalling (ZBFW) can be applied between LAN interface pairs or between a LAN segment and the Internet. The ZBFW feature can support L3/L4 inspection based on port/protocol using First Packet IQ.

Intrusion Detection System (IDS) also is supported. It monitors traffic for potential threats and malicious activity and generates threat event notifications based on configured rules. Packets are inspected against signatures downloaded to Orchestrator from Cloud Portal. Orchestrator sends appliances the signature file along with rules added to the allow list. Traffic is designated for inspection using matching rules enabled in the firewall zones.

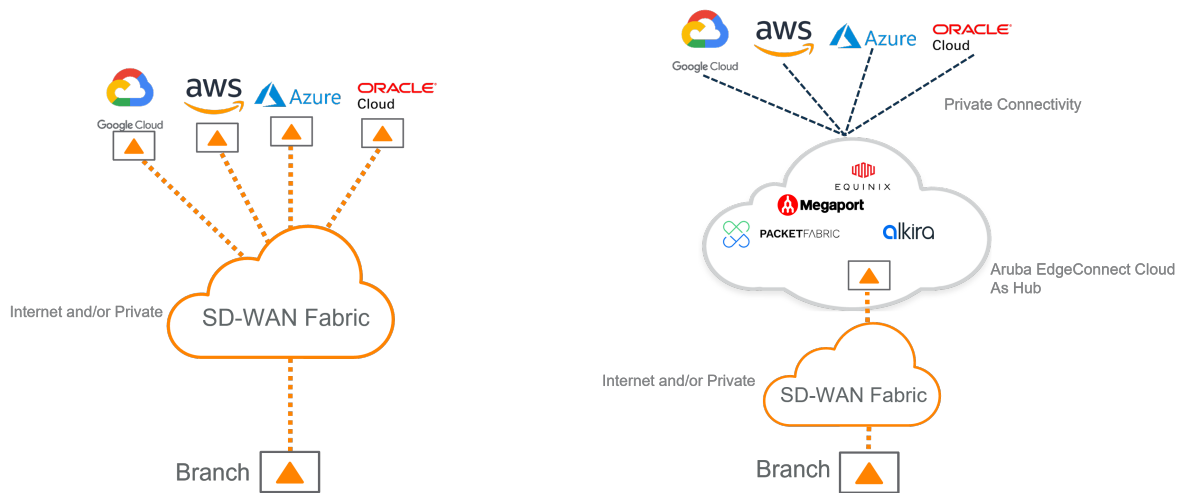
Cloud Integration

This guide is intended to provide a foundation for SD-WAN design; detailed design guidance for cloud integration is outside the scope of the content.

It is important to consider the following elements of cloud integration in the foundational design.

EdgeConnect can be deployed in virtual private clouds within Public Cloud IaaS providers.

Many large enterprises have workloads in and across multiple public IaaS providers. EdgeConnect can be deployed in “Multi-Cloud-Hub” providers such as Equinix, Megaport, and Alkira. The multi-cloud-hub provider mediates between the different cloud IaaS providers, and can connect using dedicated interconnection services for each provider.

**Figure 7:** cloud integration

EdgeConnect SD-WAN Hub Design

This section explores the basics for designing a SD-WAN deployment at hub locations and outlines key design points:

The above video summarizes the information in this chapter of the VSG.

NOTE:

Aruba recommends using the SaaS option of Orchestrator because it decreases the operational complexity of the deployment. For information on the on-premise options, refer to the Aruba Edge Connect Enterprise [User Guides](#).

WAN Transport Integration

Hub locations should accommodate every transport type that branches use to enable full hub-and-spoke connectivity. For example, if the environment includes both ISP1 and ISP2 MPLS circuits for underlay transport, both MPLS circuits should have a presence at the hub.

For private circuits, such as MPLS, it is critical to maintain underlay routing with those providers. This enables traffic flows during migrations and ensures reachability to underlay services, such as SIP trunks, that may continue to be used in the future.

Inline Deployment Methods

Inline deployment is the preferred method of deploying EdgeConnect gateways to connect two or more network segments.

The gateway is the device placed between the WAN and LAN network segments. Assigned gateways serve as routers for locally attached subnets and can provide pass-through capability for local traffic flow.

In addition to pass-through, the appliances can exchange route updates using OSPF or BGP (LAN side) and BGP (WAN side) with non-EdgeConnect SD-WAN devices to accommodate varying topologies.

The inline deployment method is generally selected when WAN bandwidth will be less than 10Gbps. Other design considerations include:

- The hubs must be installed during scheduled outage windows, since the edge router is being replaced.
- The gateway must be configured properly to support existing underlay flows during migration. This commonly entails peering BGP with existing MPLS providers to ensure reachability to data center resources from non-SD-WAN sites.
- The hub should operate with two gateways functioning as an active/standby pair.




Figure 8: inline

Out-of-Path Deployment Methods

Some situations may require deployment of an EdgeConnect gateway in a router mode configuration placed out-of-path from the existing flow of data within the network. Typically, this is due to an architectural constraint or pre-existing design.

To accommodate this requirement, deploying gateways in out-of-path mode supports dynamic routing capabilities to interface with the network infrastructure and assist with traffic redirection.

While not preferred, the out-of-path deployment method is typically used in data center hub-style environments and used as a concentration point so overlay tunnels can access data center resources.

The out-of-path design adds the gateways to the topology, leaving edge routers in place to terminate circuits. This can be useful when “scale out” architectures are required for WAN environments with extremely high throughput or when replacing the edge router is not possible.

When possible, routing protocols should be used to attract traffic to the gateways, avoiding redirection protocols such as WCCP and PBR.

LAN Routing Integration

At the hub, LAN integration should always be layer 3. The gateway should use transit P2P links (/30 or /31) to peer via BGP or OSPF with the downstream WAN Aggregation block.

BGP should be used in scenarios where there are 4+ hub locations and 1000+ gateways to provide enterprise class control of routing. BGP also is commonly deployed in highly complex environments where specific route filtering needs to be applied. This is commonly seen when there is a large amounts of mergers and acquisitions and for complex migrations from a legacy WAN to SD-WAN.

OSPF should be used in scenarios where simplicity is desired and the deployment size is under the above mentioned scale. For this example design OSPF is selected.

When advertising routes into the LAN, set the routing metric higher on the standby appliance to ensure that traffic is sent there only during an outage of the first appliance and avoid ECMP. This maintains flow symmetry to help ensure that features such as Boost function properly.

When multiple hubs are interconnected using a back-end interconnect, suboptimal routing must be avoided when introducing SD-WAN hubs. Set the routing metric sufficiently high or use more specific prefixes between the hubs to ensure the SD-WAN overlay is used only when the middle mile is unavailable. A tag should be applied to the routes as they are redistributed from subnet sharing into the IGP. Routes with this tag should then be filtered by other gateways connected by the middle mile. This avoids unwanted transit behavior.

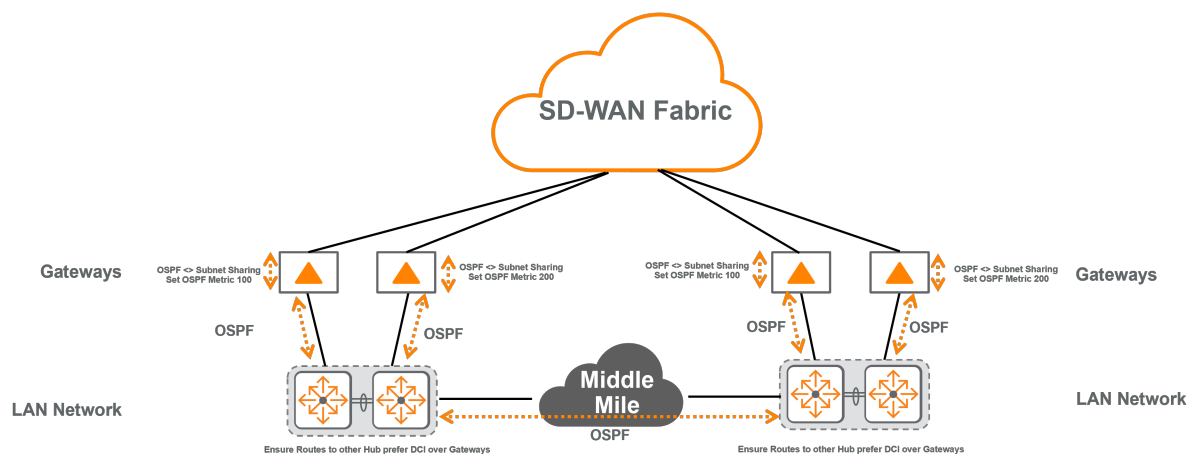


Figure 11: LAN Routing Integration

WAN Routing Integration - Overlay

The main goal when designing overlay WAN routing is to simplify the routing design while achieving the desired reachability. Proper IP scheme design is critical to network architectures and especially important in WAN designs.

Well-planned IP schemes allow for easy summarization of sites in the network. Summarization produces clean and minimal route tables which, in turn, enable easier troubleshooting and network scalability.

The recommended approach for summarization at the hubs is that each hub advertises a hub summary route and a default route. A summary route covers the prefixes at that hub while a default route attracts all other traffic as a path-of-last-resort. Branches should be configured, utilizing a template, to hub peer-priority to select a next-hop gateway should the same prefix, such as the default route, be received from two specific hubs. This simplifies hub selection and allows for flexibility in the future with regional designs, where different groupings of branches may have different hub priorities.

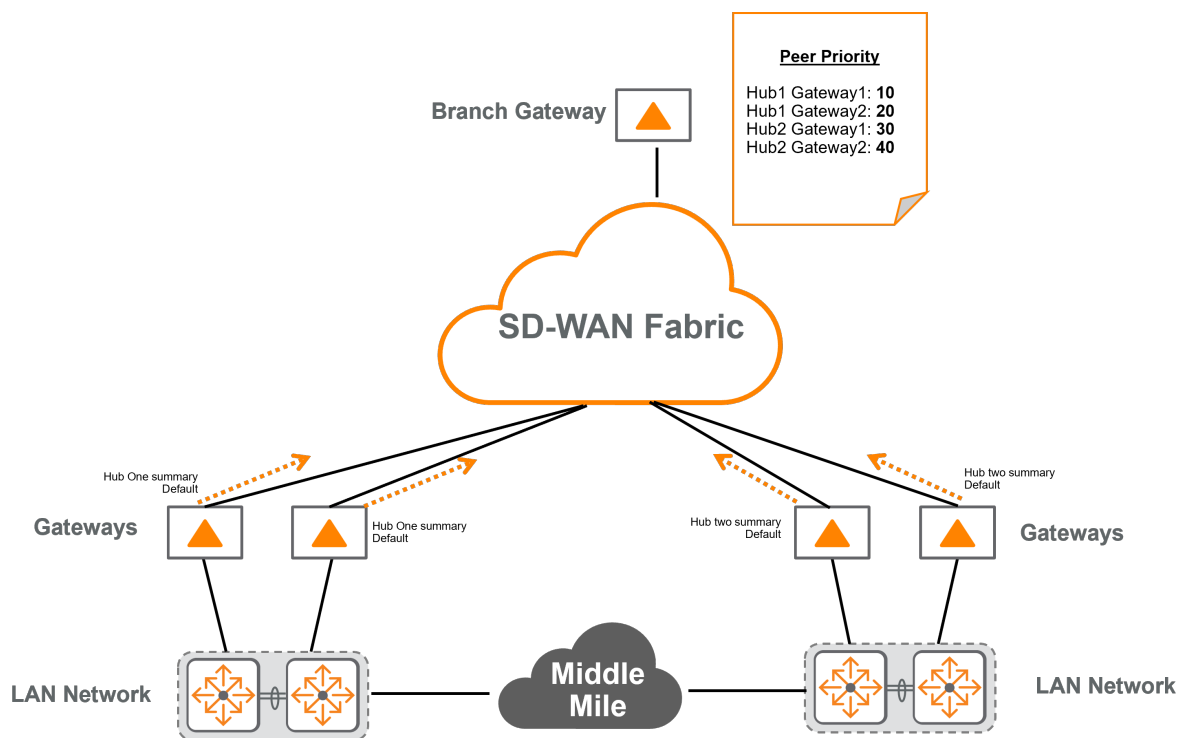


Figure 12: RA - WAN Overlay Routing

Numerous valid methods can be used to generate these two summary routes, based on the topology. Aruba recommends the method of using a static route set to null0 on the core for the hub summary route, which is redistributed into OSPF. The default route may come via OSPF from an Internet router or firewall or it can be a static route redistributed into OSPF with the core pointed to an upstream Internet device. Ensure that a method such as a route track is in place to remove the default route from the table. This avoids blackholing backhauled Internet traffic if an Internet outage occurs at the hub. When redistributing routes for subnet sharing, use a route map to limit the redistribution to only those two summary routes.

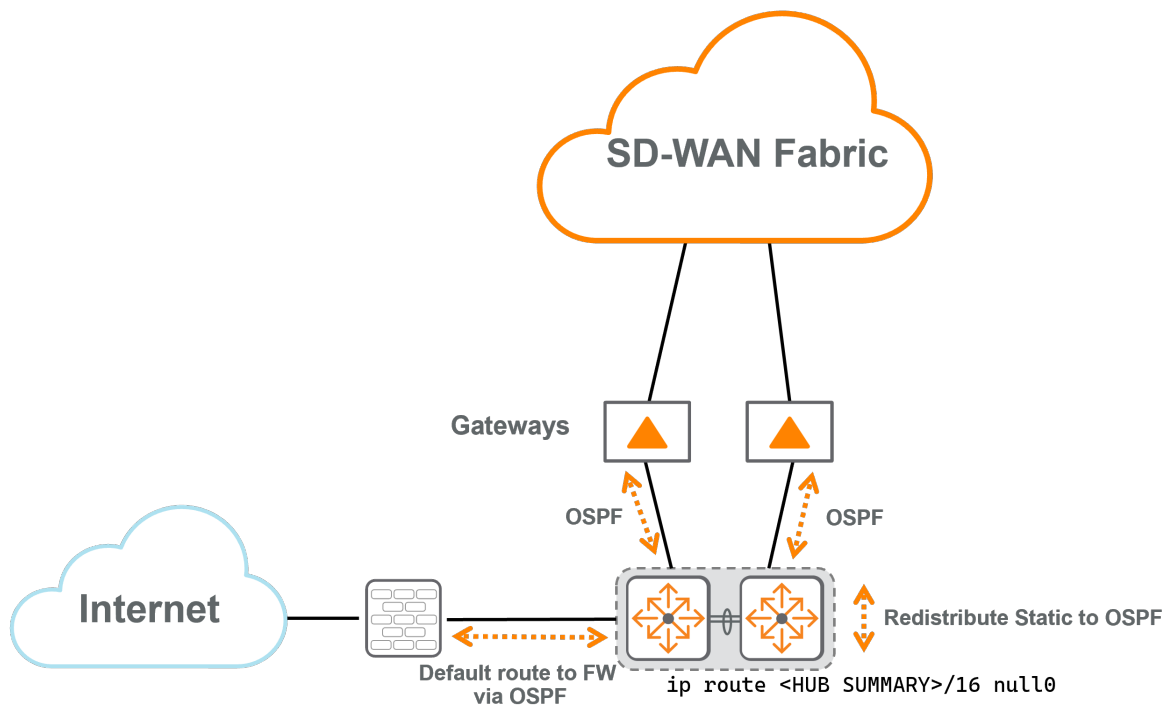


Figure 13: RA - Route Generation

WAN Routing Integration - Provider

Routing integration with the WAN providers, also known as the underlay, ensures route reachability between the gateways to establish IPsec tunnels.

For private transport types such as MPLS, eBGP is used to advertise reachability of the WAN interface for IPsec tunnel establishment

In most deployments, operators are migrating from a traditional WAN to an SD-WAN solution. As sites are converted from traditional to SD-WAN, branch-to-branch connectivity must be considered. When planning the migration, ensure that the hub advertises a default route into the underlay to attract traffic for converted to non-converted traffic flows.

For public transports such as Internet or cellular, a static default route is used.




Figure 14: RA - Provider Routing Routing

EdgeConnect SD-WAN Branch Design

This section provides the basics for designing and orchestrating an SD-WAN deployment for branches, with considerations specific to the branch network design.

The above video summarizes the information in this chapter of the VSG.

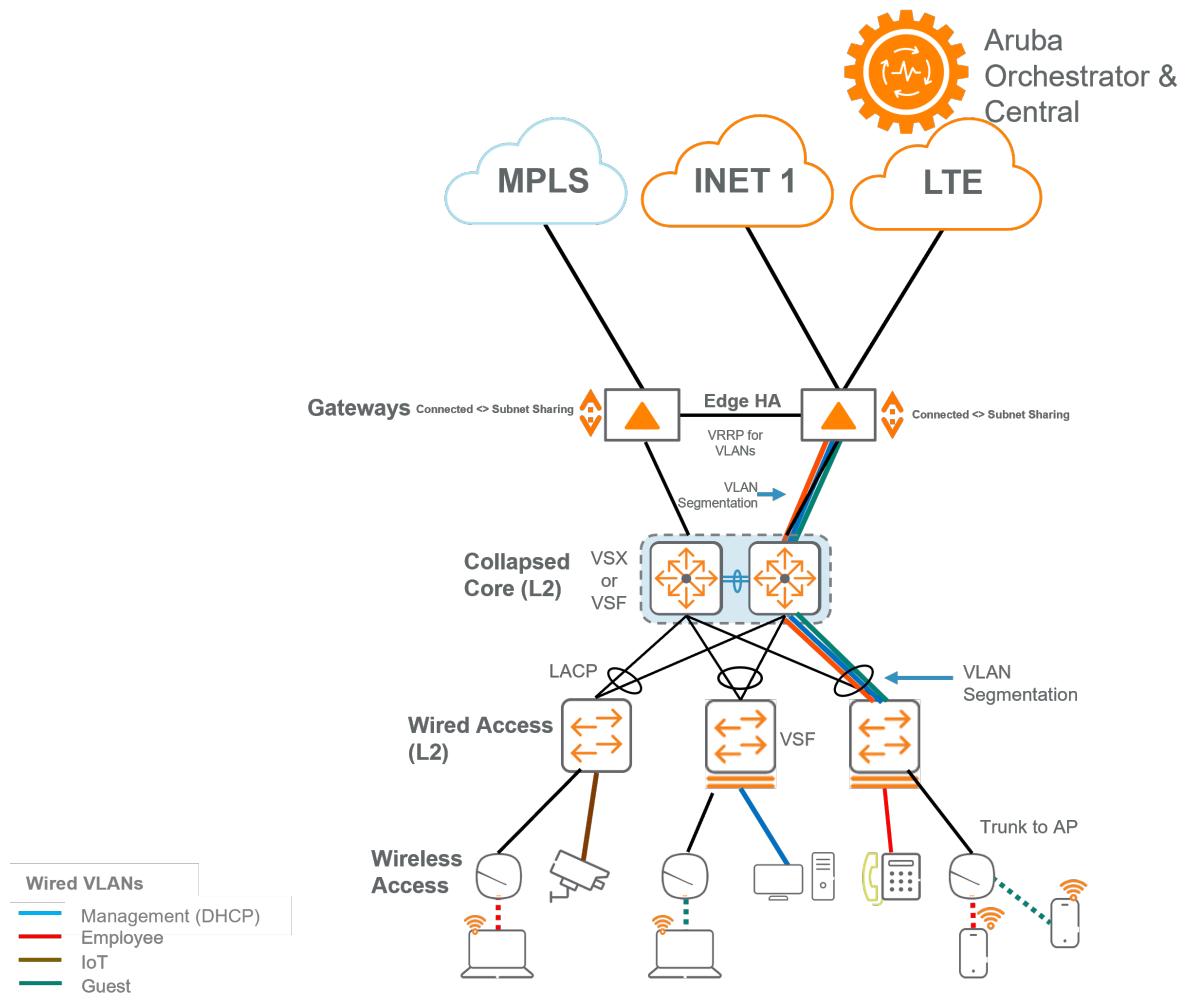
Design Overview

This section describes design details for the SD-WAN gateways, switches, and access points required to design branch networks. The design elements of the branch include:

- **EdgeConnect SD-WAN gateways** provide the default gateway for all VLANs, as well as connectivity to the SD-WAN fabric and Internet services.
- **Wireless APs** operate in bridge mode to eliminate the need for a wireless gateway at the branch, reducing the number of total infrastructure devices.
- **Wired access** is provided by stacked 6200/6300 switches for ease of management.
- **A collapsed core** is provided with VSX or VSF with multi-chassis LAGs to the access, providing a non-blocking Layer 2 domain for effective bandwidth usage in the large branch.
- **Segmentation between VLANs** is provided by the EdgeConnect SD-WAN gateways using role-based policies, application-aware policies, and IP address policies.

The topologies below illustrate large and small branch samples.

Large Branch Topology



Small Branch Topology

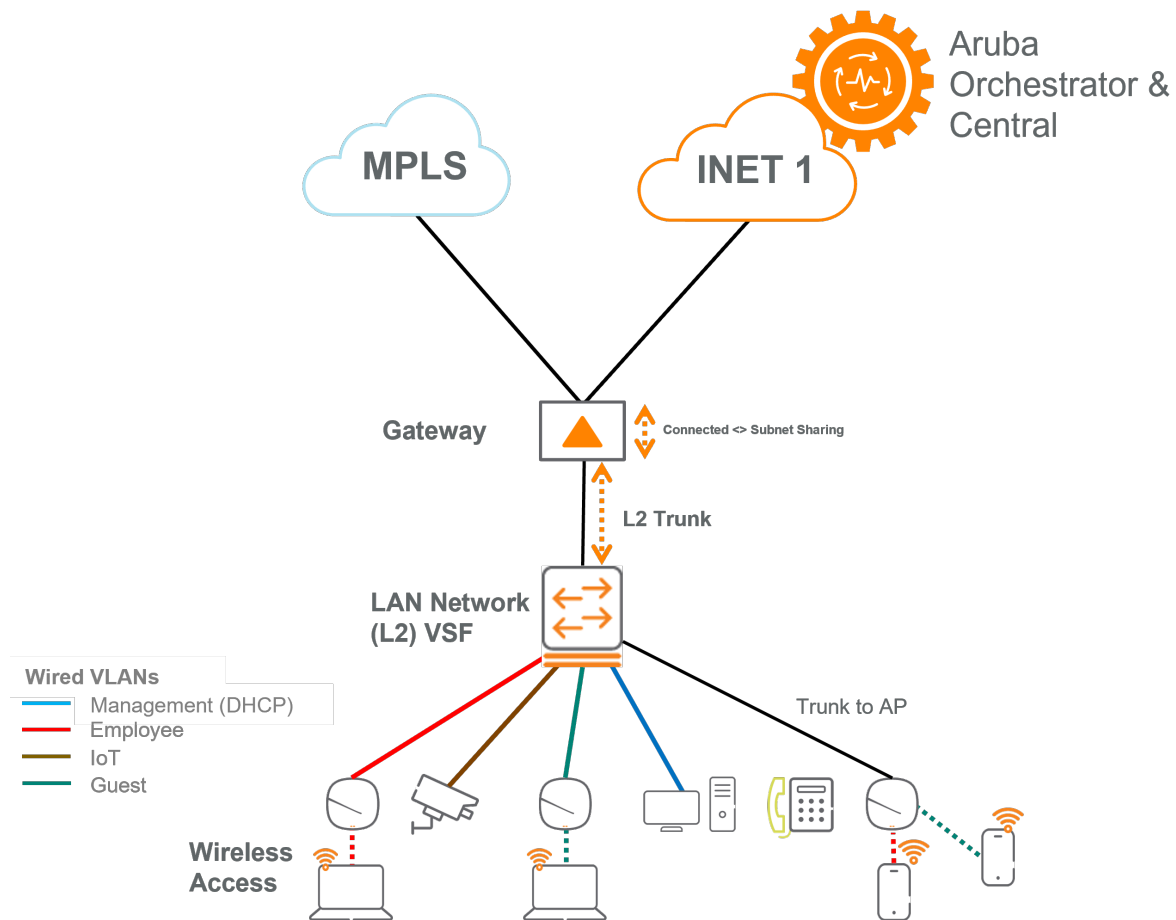


Figure 16: Small Branch Overview

SD-WAN Gateways

This section covers details of the SD-WAN gateway design at the branch.

Inline Topology

Inline is the preferred method for deploying EdgeConnect SD-WAN Appliances to connect two or more network segments at the branch.

The “appliance” is the device placed between the WAN and LAN network segments.

For inline deployment, the assigned appliances serve as the routers for locally attached subnets and can provide pass-through capability for local traffic flow. In addition to pass-through, the appliances can exchange route updates using OSPF or BGP (LAN side) and BGP (WAN side) with non-EdgeConnect SD-WAN devices to accommodate various topologies.

When deployed inline, appliances perform functions typically allocated to a traditional branch router, such as performing NAT translations, acting as a DHCP server, and providing a variety of interfaces for different connectivity types.

High Availability

High Availability (HA) at the branch is achieved using a feature called EdgeHA. EdgeHA provides fault tolerance for both WAN connectivity and equipment. The classic SD-WAN use of “hybrid WAN” enables the flow of customer traffic over multiple underlay networks (MPLS, Internet, 4G, etc.), a beneficial feature for maintaining customer traffic if a single underlay network experiences issues such as degradation or downtime.

In the diagram below, a branch deployment shows with two EdgeConnect instances, each connected with a single WAN link to two different underlay networks. Note that no WAN-side switches are required. The EdgeConnect appliances are interconnected with an EdgeHA link that enables tunnels for each underlay network to connect to both appliances. On the LAN side, protocols such as Virtual Router Redundancy Protocol (VRRP) or standard routing protocols are used to direct customer traffic to the EdgeConnect appliances.




Figure 17: EC HA

LAN Routing Integration

At the branch, LAN integration may be Layer 2 or Layer 3. Layer 2 LAN integration is recommended because it enables the SD-WAN Gateways to act as the policy enforcement point between VLANs.

For Layer 3, when advertising routes into the LAN, set the OSPF metric higher on the standby appliance to ensure that traffic is sent there only during an outage of the first appliance and to avoid ECMP. This maintains flows symmetry to help ensure that features such as Boost function properly.

For Layer 2, ensure that all redundant links are port-channelled. The EdgeConnect SD-WAN appliances do not participate in spanning tree, so connect their interfaces to switchports with spanning-tree loop prevention mechanisms such as BPDU-guard enabled. When using Layer 2, the gateways serve as the default-gateway for all VLANs. VRRP should be configured to provide gateway redundancy.

WAN Routing Integration: Provider

For private circuits such as MPLS, it is important to consider if underlay routing is required. Underlay routing enables direct branch-to-branch traffic flows between SD-WAN sites and traditional sites and enables reachability to underlay services, such as SIP trunks, possibly used in the future.

If underlay services are not required and there is limited branch-to-branch communication, consider removing the provider BGP connections at the branch and establish a default route. This causes all spoke-to-spoke traffic to traverse through the hub.

BGP adjacency can be maintained during migration, if required, but removing BGP is recommended after all sites are on the SD-WAN solution, when possible. This ensures that the MPLS circuit is used only for underlay tunnel establishment and pushes all traffic into the overlay.

WAN Routing Integration: Overlay

The main goal when designing overlay WAN routing is to simplify the design while achieving the desired outcome. Aruba recommends that each branch advertises a summary route and covers the prefixes at that branch. A static default route should be configured on Internet circuits for Internet egress traffic and establishment of overlay tunnels.

Switching

The size of most branch networks requires the use of a collapsed core topology. The sections below describe what to consider at each layer of the switching infrastructure. Many of these topics are covered in more depth in the [Campus Design](#) section of the VSG.

Collapsed Core

The collapsed core can function as a Layer-3 point in the network, terminating the VLANs, or act as a Layer-2 point. The recommended design is Layer-2, allowing the default-gateway to reside at the EdgeConnect SD-WAN gateways where advanced policy can be enforced. In a small branch, this is the only switching block. In a large branch, the collapsed core interconnects the access switches and the SD-WAN gateways.

Chassis Virtualization

Selecting the chassis virtualization technique at the collapsed core is critical for effective design. CX 8000 series switches run VSX, while CX 6000 series switches run VSF.

VSX provides increased resiliency by separating the control and management planes. VSX also enables in-service software upgrades (ISSU) that can help avoid downtime at the branch. The main challenge with VSX is that all ports on the CX 8000 series switches ship in the shutdown state and are configured as Layer 3 ports. Deploying VSX pairs requires a one-touch procedure, instead of zero-touch, to pre-configure the switch for reachability to Central.

VSF, which is the default on CX 6000 series switches, unifies the control, data, and management planes for slightly reduced resiliency, but easier management. Software upgrades generally require an outage. VSF stacks enable full zero-touch provisioning.

Access

Use Aruba CX switches that support VSF stacking for simplified growth in the network closet. For Layer 2 access designs, use uplink ports on different VSF stack members, one into each MC-LAG configured aggregation switch. This ensures efficient, fault-tolerant Layer 2 bandwidth up from the access layer.

Enable Aruba ESP Colorless Ports by configuring port policies to allow 802.1x dynamic authentication and network configuration.

Enable Layer 2 protection mechanisms such as Loop Protection, BPDU Filter, Root Guard, and BPDU Protection.

Wireless

Wireless access is provided using a bridge mode deployment. In this mode, the AP bridges traffic from the WLAN onto the correct user VLAN. AP-connected switch ports must trunk all wireless user VLANs. For more detail on this deployment mode, refer to the [WLAN Design](#) section of the VSG.

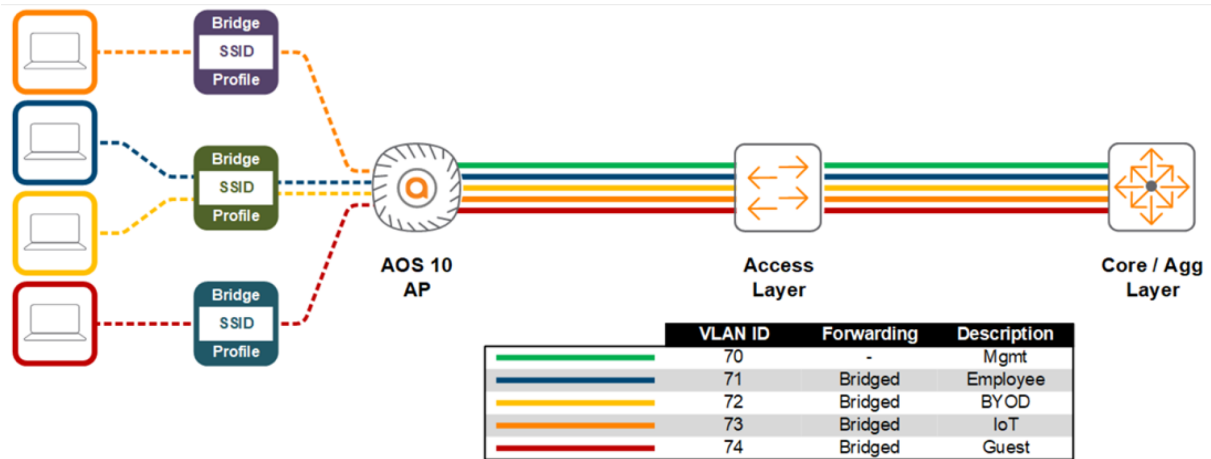


Figure 18: Bridge Mode Wireless

Zero Touch Provisioning

Zero Touch Provisioning (ZTP) is critical for the branch. For many deployments, prestaging all the equipment with configurations is not feasible and non-IT resources must be called to do the physical deployment.

EdgeConnect SD-WAN

ZTP enables the EdgeConnect SD-WAN appliance to communicate with the Aruba Cloud Portal. The Cloud Portal redirects the appliance to and registers it with the customer's Orchestrator, which enables remote configuration of the device by an operator or a configuration push from Orchestrator using a YAML file.

To accommodate ZTP, an Internet circuit with public DNS is recommended at the branch. This enables direct, unrestricted communication with the Cloud Portal and Orchestrator. Sites without a local Internet circuit likely are unable to use ZTP, and prestaging of the devices is required.

The EdgeConnect SD-WAN appliances should be configured to provide DHCP on a management VLAN to LAN network infrastructure to accommodate the ZTP process.

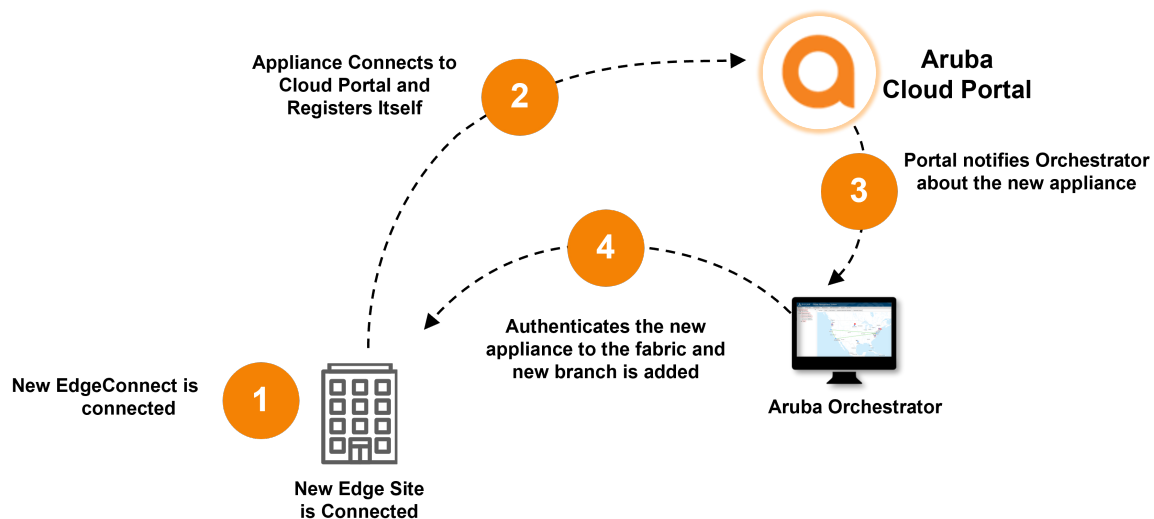


Figure 19: EC ZTP

Switching and Wireless

The APs and switches must receive IP addressing and domain name server (DNS) information from the SD-WAN gateways before they can communicate with Central. Switches and APs should be connected after the SD-WAN gateways are online and fully functional. The SD-WAN gateway must provide the following items to the APs and switches:

- DHCP addressing
- DHCP options 3 (Router) and 5 (Name Server)
- Internet access

Management

Management of APs and switches is handled by Aruba Central. Management of SD-WAN gateways is handled by Aruba Orchestrator. Although both solutions have an on-premises option, the cloud-hosted option is recommended.

These platforms offer integrations to ease the management and reporting between them, and will continue to be integrated more fully. For more information on the integrations, see the [Aruba SD-WAN Docs](#).

Segmentation and Policy

This section describes how to segment traffic within the branch, across the WAN, and to the Internet.

Within a Branch

Within the branch, Layer 2 VLANs are used to segment the traffic. This provides a coarse-grain level of segmentation. Devices within a VLAN can communicate directly, but communication outside a VLAN must be routed between the SD-WAN gateways.

The SD-WAN gateways act as the default-gateway for all VLANs. The SD-WAN gateways provide zone-based stateful firewalling, IPS/IDS, and role based policy to filter traffic between the VLANs.

Across the WAN

Across the VRFs, which are referred to as “segments” in Orchestrator, are used to provide layer 3 isolation and a dedicated route table. VLANs at the branch are mapped into segments. Common segments include Guest, IoT, and Corporate.




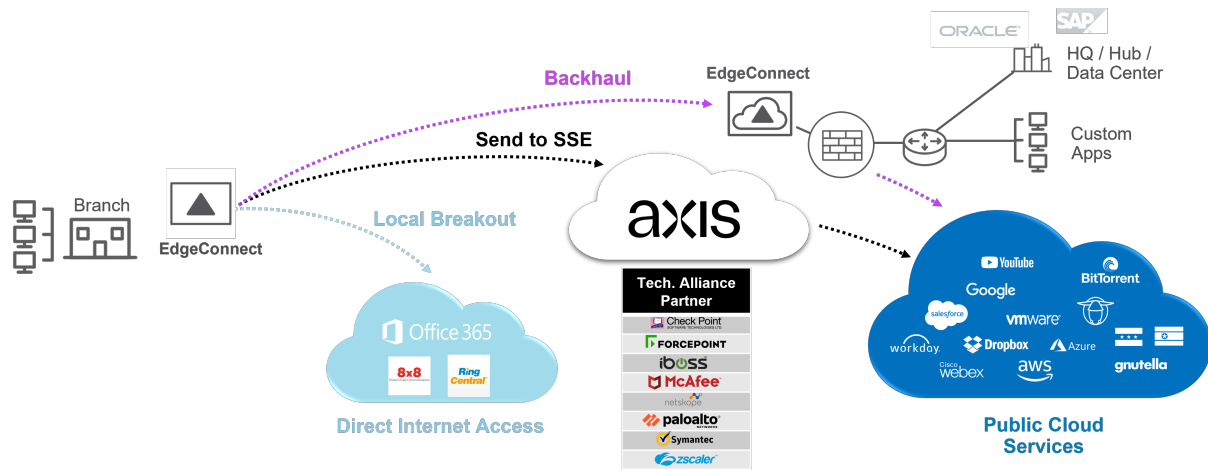
Figure 20: VRF Example

Internet

Internet egress design at the branch is very flexible. On-box IPS/IDS and stateful firewalling can be used for on-box policy of traffic destined to the Internet. For more advanced cloud-based filtering integrations, use Secure Serviced Edge (SSE) providers such as Axis and Zscaler.

Internet egress designs vary based on requirements. A common Internet egress design is shown below.

- Guest BIO has direct Internet access, with limited or no on-box policy applied to move the traffic out of the network as fast as possible.
- Business traffic is sent to an SSE for further logging and policy enforcement.
- Sensitive traffic, which may be subject to compliance, is back-hauled to a centralized DC for additional logging and inspection.
- Other traffic is subjected to IPS/IDS and basic on-box policy before being routed directly to the Internet.

**Figure 21:** VRF Example

Aruba EdgeConnect SD-Branch Solution

Fundamentals

This section reviews Aruba Central, Gateway devices, and Orchestration details to provide an overview of EdgeConnect SD-Branch components and offer guidance for device selection.

Aruba Central

Aruba Central is a powerful cloud networking solution. As the management and orchestration console for Aruba ESP (Edge Services Platform), Central provides centralized cloud-based management for devices, policy, and templates to enable deployment of branch sites quickly, with group-based configuration. For more detail on groups, see the “Orchestrator and Group Design” section.

Aruba Central offers key insights on WAN health and optimization to help organizations determine the best path for traffic on per-user, per-device, or per-application policies. Historical data reports, monitoring for PCI compliance, and troubleshooting for regional and global locations can all be viewed in the Central dashboard.

Aruba Central also hosts the cloud-based Orchestrators that allow dynamic building and scaling of IPsec tunnels as needed.

SD-WAN Orchestrator

The SD-WAN Orchestrator automates site-to-site tunnels and route propagation between Branch Gateways (BGW) and VPN Concentrators. The SD-WAN Orchestrator has two components: Overlay Tunnel Orchestrator (OTO) and Overlay Route Orchestrator (ORO).

The Aruba SD-WAN Orchestrator delivers the following features and functions:

- The IPsec overlay is automatically created using Overlay Tunnel Orchestration.
- Automatic route propagation is managed by Overlay Route Orchestrator. Route redistribution can be performed at the group configuration level.
- Hub preference enables an administrator to prefer one hub site to another site by setting routing costs that are translated into the data center’s dynamic routing process.
- Devices dynamically adopt the group overlay configuration, build the tunnels, and enforce route policy.
- Scalability is built into the orchestration, to assist organizations with building robust and cost-effective routing designs.

This section reviews the differences between each of the Orchestrator components and describes how the two components work together to automate tunnel orchestration. Additional details about the Orchestrator and other SD-Branch features can be found [here](#).

Overlay Tunnel Orchestrator

The Overlay Tunnel Orchestrator automatically generates and manages IPsec configuration between devices and identifies where devices should tunnel.

The Orchestrator identifies where to build tunnels using labels on interfaces. Labels include a combination of **Uplink Types** and **Link Names** configured on the WAN uplink page within Aruba Central. This enables the Orchestrator to grab interface IP addresses dynamically to build IPsec tunnels between branch gateways and VPN Concentrators.

Supported uplink types currently available include MetroEthernet, INET, MPLS, and LTE. The combination of both labels (Uplink Types and Link Names) applied to an interface identifies the interface as a WAN Uplink. If no WAN uplink is identified, tunnels are not built.

In order for a BGW to build a tunnel to a headend site, the uplink types should match, with some caveats. For uplink naming, MPLS uplink types must have the same uplink type and link name for the tunnel to be built. INET, MetroEthernet, and LTE uplink types prioritize building a tunnel to a device with the same link name; however, if the criteria do not match, a tunnel is built to the next available INET, MetroEthernet, or LTE uplink type regardless of link name. When deploying an EdgeConnect SD-Branch, it is important to maintain a consistent expandable naming convention so tunnels can be built between devices.

NOTE:

VPNCs only support MPLS and INET uplink types. If a branch gateway uses MetroEthernet or LTE uplink types, they connect to the INET uplink type of the VPNC using the uplink name as an additional matching criteria. If there is no matching name, the tunnel is built to any available INET uplink of the VPNC .

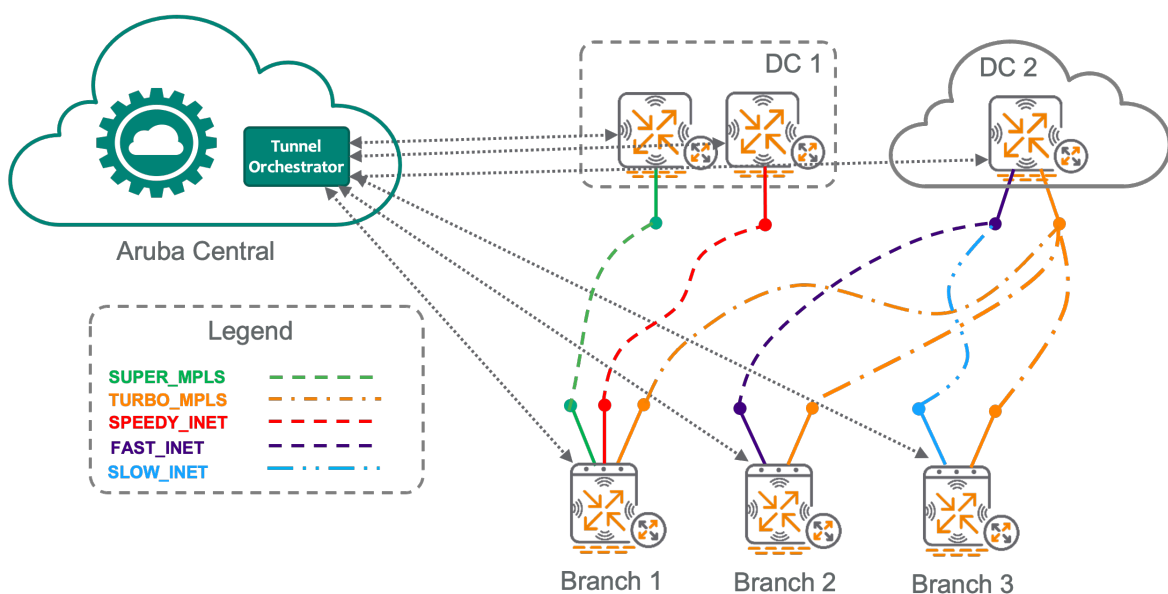


Figure 22: tunnel_orchestration

NOTE:

The slow_INET has an uplink name that does not match the headend, but due to INET uplink type, the tunnel is created between the two endpoints

Overlay Route Orchestrator

The Aruba Overlay Route Orchestrator enables the distribution of routing information between branches and VPN concentrators. It provides route distribution across sites using the Control Connection Overlay Agent Protocol (OAP). The OAP runs in a separate process on each device and interacts with the underlay routing stack to exchange and advertise prefixes with the Route Orchestrator that distribute routes to branches and VPN Concentrators.

The Route Orchestrator learns routing information by redistributing routes into the SD-WAN overlay. Administrators can select one or more of the routing sources (connected, static routes, OSPF, BGP) to determine the routes that enter the SD-WAN overlay.

The main functions of Aruba Route Orchestrator include:

- Learning routes from headend and branch sites
- Advertising routes across the SD-WAN network with appropriate costs
- Redistributing routes into the LAN side with appropriate costs.

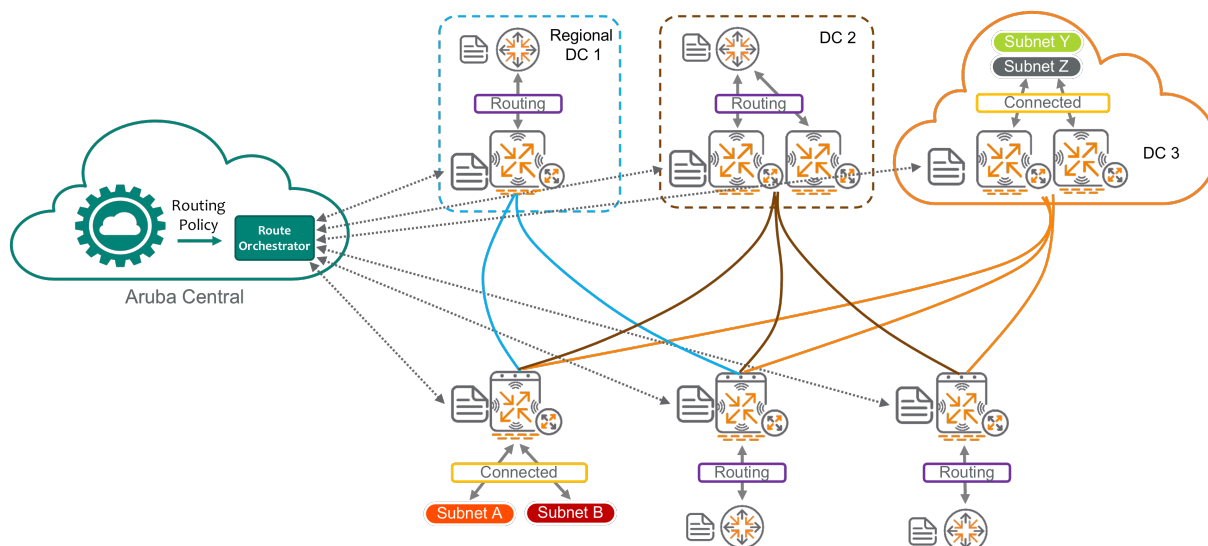


Figure 23: Route Orchestrator

Orchestrator and Group Design

The previous two sections described the Aruba Orchestrator components and the parameters required for user input. However, it is important to note that users do not need to input parameters for every device brought into the SD-WAN fabric. The Orchestrator works hand-in-hand with Aruba Central, which has a two-tiered hierarchy for group and device level configuration. Configurations for DNS, Route Redistribution, WAN uplinks, and other options are defined a single time at the group level. Specific configuration, such as IP addresses and hostname, is required at the device level.

Any devices brought into the group inherit the configuration parameters as soon as they are online. This enables administrators to configure a group a single time. When creating groups, it is important to decide the routes to be included in the fabric and the interfaces to be used.

Device Overview

The gateway can be deployed physically or virtually for WAN connectivity. Aruba uses two different types of gateways: a Headend Gateway and Branch Gateways.

The same model device can be a headend gateway or a branch gateway, depending on the scale of the SD-WAN fabric. The difference between a branch gateway and headend gateway depends on its function or deployment.

In addition to gateways, access points can participate in the SD-WAN fabric. They are referred to as Microbranch access points and fall under the branch category. The following list covers the definition of Microbranch, Headend, Branch, and Virtual gateways, as well as the role the device performs.

VPN Concentrator (VPNC) - The VPN Concentrator, also called a hub or headend gateway, acts as a VPN concentrator terminating IPsec tunnels, from branch gateways, Microbranch APs and VIA clients. The headend also advertises routes from the data center or campus environment to the branch gateways using redistribution of Connected, Static, OSPF, or BGP prefixes .

Virtual Gateways (VGW) - The virtual gateway extends the SD-WAN overlay services to the public cloud infrastructure. Virtual gateways function as VPN concentrators and terminate tunnels from branch gateways, Instant APs, and VIA clients. Like the hardware VPN concentrators, virtual gateways support routing, security, and tunneling features. Virtual gateways are supported in Amazon Web Services and Microsoft Azure.

NOTE:

The design considerations for using a physical headend gateway versus a virtual gateway are reviewed in the [Hub Design](#) section.

Branch Gateways (BGW) - The branch gateway is the appliance at each remote site that terminates its IPsec tunnels to a headend gateway. The branch gateway also can provide dynamic segmentation by serving as a policy enforcement point for wired, wireless, security, and WAN policies including routing. The gateway functions include stateful firewall, web content classification, hybrid WAN connectivity, IPsec VPN, QoS, and WAN path monitoring and selection.

Microbranch - Is a very small branch deployment that uses access points that support building an IPsec tunnel to a headend gateway.

The topology illustration below shows gateway and their placements in a WAN topology. LAN devices also are shown. The specific types of devices needed on the LAN side of a branch or headend site vary depending on the needs of the organization. Details are provided in the “Hub and Branch” design sections.

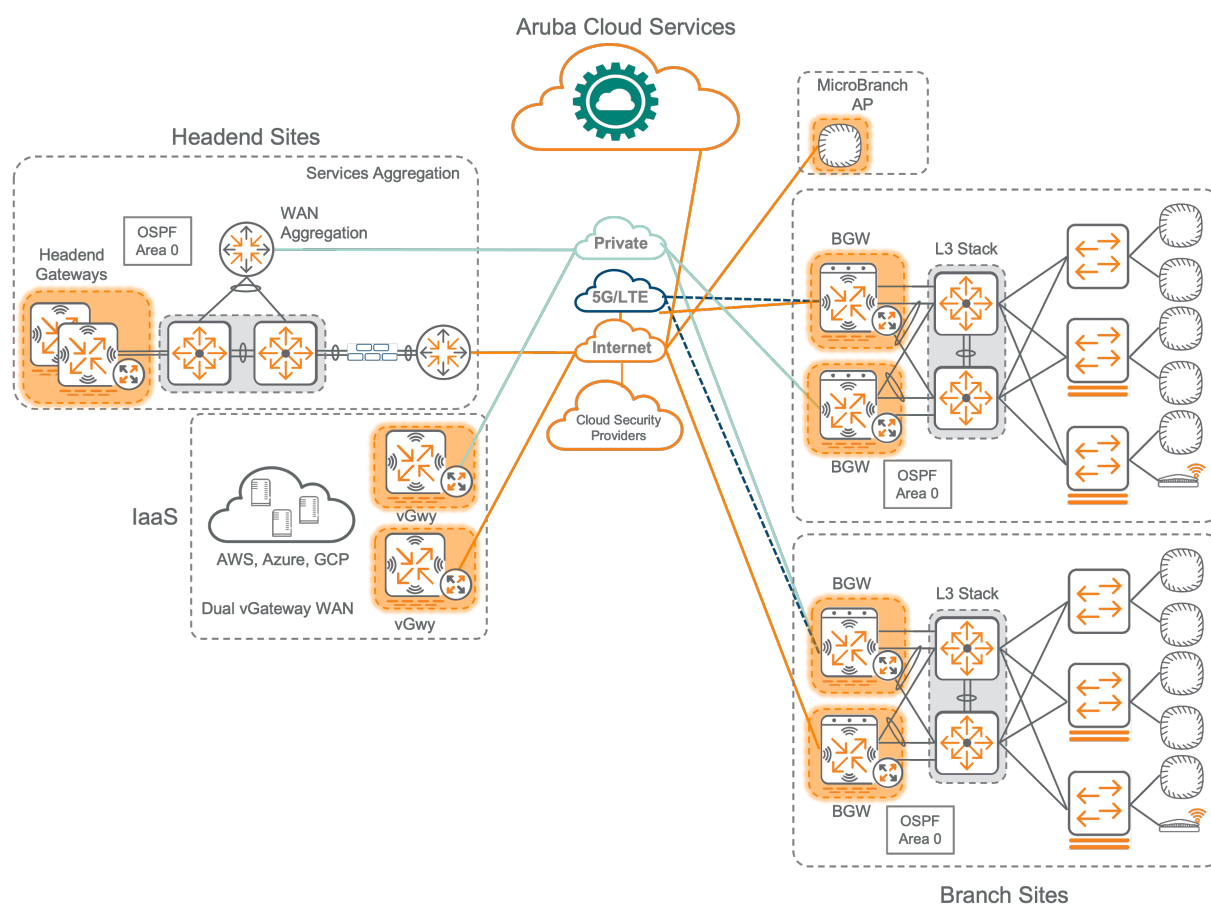


Figure 24: gateway_placement

The table below shows some of the scale numbers for gateway. For complete scale numbers, refer to [this](#) page.

Platform	Deployment	Max IPsec Tunnels	Max Routes	Firewall Sessions	WAN Throughput
7280	Headend	8192	32,768	2M	50 Gbps
7240XM	Headend/Branch	6144	32,768	2M	30 Gbps
9240	Headend/B	1024	32,769	—	30 Gbps
vGW-4G	Headend	8192	131,072	6M	4 Gbps

Platform	Deployment	Max IPsec Tunnels	Max Routes	Firewall Sessions	WAN Throughput
vGW-2G	Headend	4096	65,536	256k	2 Gbps
vGW-500M	Headend	1,600	2048	64k	500 Mbps
7220	Headend/B	24,576	16,384	2M	21 Gbps
7210	Headend/Branch	16,384	8192	2M	8 Gbps
7030	Headend/B	512	4096	65k	2.6 Gbps
7024	Headend/Branch	256	4096	64k	2.6 Gbps
7010	Headend/B	256	3840	64k	2.6 Gbps
9012	Headend/Branch	512	12K	64k	4 Gbps
9004	Branch	512	12K	64k (128k on 2.3+ with IDPS disabled)	4 Gbps
7005	Branch	—	3840	64k	1.2 Gbps
7008	Branch	—	3840	64k	1.2 Gbps

NOTE:

For details on [switching](#) and [access points](#), refer to the respective data sheets.

Device Onboarding

Users have three methods to onboard gateways: One Touch Provisioning (OTP), Zero Touch Provisioning (ZTP) and the installer application. All methods can be used in combination; however, using a consistent deployment procedure is recommended when setting up branch sites.

OTP and ZTP both require an admin user to place the gateways in the proper groups for configuration after onboarding in Central. The installer app automatically places onboard devices into the appropriate group during onboarding.

One Touch Provisioning is typically used for sites when DHCP is not available for connected uplinks, usually a headend site or a branch site. One touch provisioning requires using the CLI or Web UI to configure the uplinks of a gateway for internet reachability. The gateway is then directed to Central for the remaining configuration.

Zero Touch Provisioning is good to use for deployments where DHCP is available. After a gateway receives DHCP, it reaches out to Central for its configuration.

Installer Application allows an authorized installer to select a group and scan a brand new device to be added into the selected group. The admin must specify the device groups the installer is authorized to access.

System IP

The system IP (system-ip) is a critical configuration element for each gateway operating as a VPNC or BGW. When onboarding one VLAN interface as its system IP. By default, the Aruba Gateway uses this interface to communicate with network services such as RADIUS, syslog, TACACS+, and SNMP. The VLAN interface selected for the system-ip on each gateway must have an IPv4 address assigned for the gateway to be fully functional. A gateway cannot initialize fully unless the assigned VLAN interface is active and operational. Central does not allow a gateway to obtain addressing dynamically from Internet service providers using DHCP or PPPoE as the system IP.

Gateway pools can be used to allocate system IP addresses automatically to a dedicated VLAN interface, which is then designated as the System IP address. Each pool includes a unique name along with starting and ending IPv4 addresses. The range of addresses defined for each pool cannot overlap. Aruba recommends configuring one gateway pool for each group, since IP addresses are configured and applied to VLAN interfaces on a per-group basis. The gateway pool must include enough IPv4 addresses to support all the Aruba gateways assigned to the group. Although a group can support multiple gateway pools, specific IP addressing should not be applied dynamically.

Security Capabilities

Security in EdgeConnect SD-Branch is built in layers, from the hardening of the operating system to the integration with best-of-breed security partners. ArubaOS, running on gateways and microbranch, is a tightly hardened platform that includes:

- **Secure Boot** — Requires Aruba-issued certificate (TPM) to load ArubaOS.
- **Secure Zero Touch Provisioning (ZTP)** — ZTP leverages the TPM loaded in the Aruba gateways to secure communications with Aruba Central.
- **AES 256 encryption** — Encryption is used for SD-WAN overlay tunnels.
- **Aruba Role-Based Stateful Firewall** — The firewall supports scalable configuration using firewall aliases, ALGs, and role-based policies.
- **Deep Packet Inspection** — Qosmos's application engine and signatures provide the capacity to identify almost 3500 applications.
- **Web Content, Reputation and Geo-location Filtering** — WebRoot's machine learning technology classifies content, reputation, and geolocation for billions of URLs.
- **Aruba Threat Defense** — Powered by ProofPoint's Threat Intelligence, Aruba 9000 series gateways can perform IDS/IPS functions for all branch traffic.

The Aruba ESP solution can integrate with ClearPass (or any other AAA server) to form a true policy-driven branch. This model dynamically assigns policies based on users, devices and applications, as opposed to the traditional method of assigning these policies manually based on ports, VLANs, and IP addresses. The policy-driven branch can be enhanced further by integrating with more than 140 partners in the [ClearPass Exchange program](#), to take advantage of Aruba ESP's cumulative [AI/ML-driven Client Insights](#).

EdgeConnect SD-Branch WAN Overview

Designing the EdgeConnect SD-Branch WAN overlay requires consideration of three key elements: WAN Topology, WAN Monitoring and WAN Policy. The elements work hand-in-hand to provide a highly secure overlay that delivers optimal performance.

WAN Topologies

The following section reviews the three types of overlay topologies available to organizations using EdgeConnect SD-Branch. Any of the topologies can be used in combination with another topology.

Hub-and-Spoke

The Aruba EdgeConnect SD-Branch solution supports a hub-and-spoke topology in which SD-WAN overlay tunnels are established between headend gateways (hubs) and BGWs (spokes). The gateways at the headend sites provide routing and forwarding for hub-to-spoke and spoke-to-spoke traffic.

This is the default deployment, since most organizations' applications are centralized in a single data center and branch sites commonly exchange no data or minimal data with lower priority.

The figure below illustrates a hub-and-spoke topology with spoke-to-spoke traffic passing through the hub location.

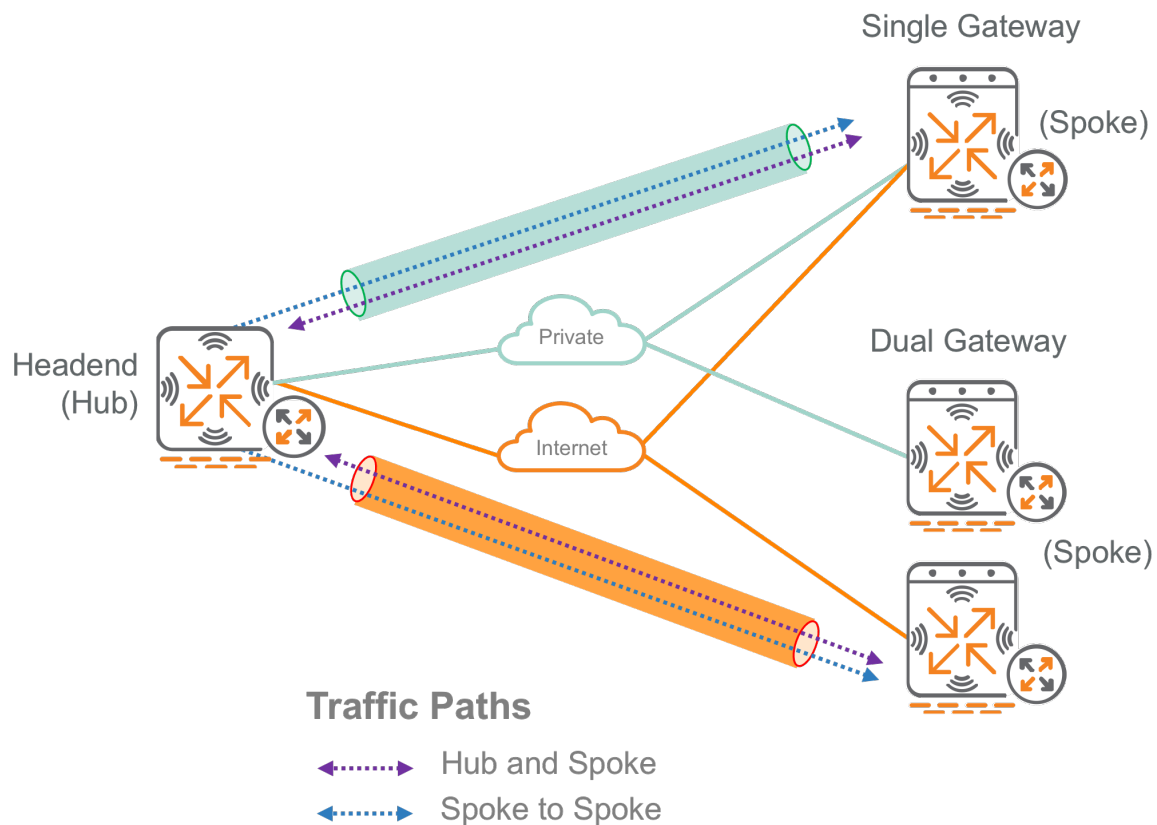


Figure 25: Hub-and-Spoke

Deployments require one headend site with one or more installed gateways that terminate VPN tunnels initiated from the BGWs installed at the branch sites. The number of gateways deployed in each headend site depends on overall deployment size and redundancy needs. Smaller deployments consist of one gateway installed at a headend site to service all the BGWs installed at branch sites.

Larger SD-Branch deployments can incorporate additional hub sites, providing redundancy in case of a primary hub failure. A typical large deployment includes a primary and secondary headend.

More complex topologies using additional hub sites also are supported. For example, a deployment might include a cloud-based data center hosting a specific application or service using virtual gateways.

Hub Mesh

Aruba supports mesh topologies between on-premise hubs (physical gateways) and/or cloud hubs (virtual gateways). This enables hubs sites to communicate directly with one another and is generally used for communications between regional hubs or between multiple cloud providers. This includes traffic coming from a branch site. For example, a branch site could have traffic destined to “AWS Cloud DC” and have a preference for the “On-Premises DC”. In this case the “On-Premises DC” would forward the traffic to the “AWS Cloud DC” using the hub mesh tunnels.


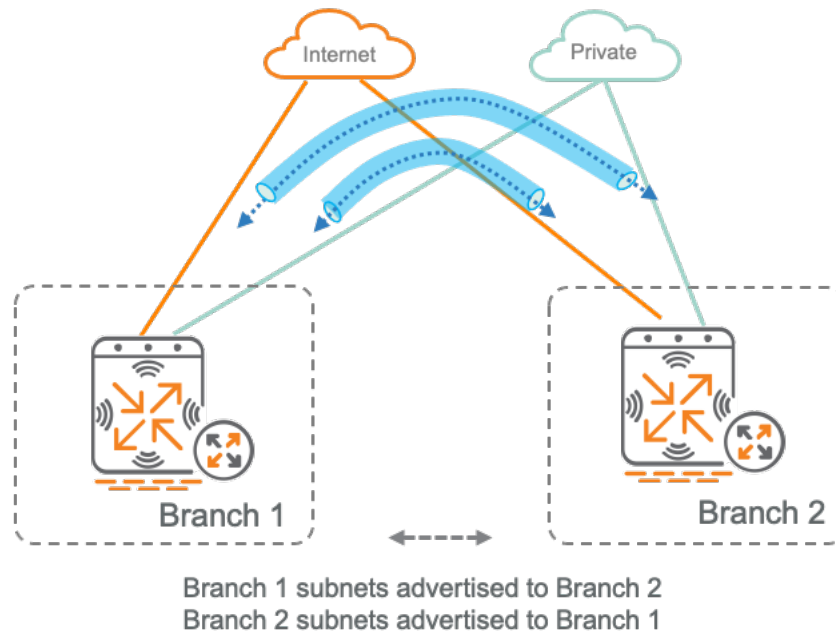


Figure 26: Hub Mesh

Branch Mesh

The branch mesh topology configuration enables branch gateways to establish secure overlay tunnels with other branch gateways in the same group or a different group. When a branch mesh topology is configured between two or more branch gateways, a branch mesh link is established to transport traffic securely between them. Branch mesh can be used for distributed enterprises or for organizations with branch sites that have multiple inter-branch communication that should not be hair pinned through a hub site. Starting with the AOS 10.5 release, designating a hub site is no longer required to establish a branch mesh.

AOS 10.5 or Greater**Traffic Paths**

←.....→ Branch Mesh

Figure 27: branch mesh 10-5

For deployments using AOS 10.4 or earlier, a hub site must be designated in the SD-WAN fabric to enable ORO between sites allowing branch gateways to exchange routes. However, this does not stop branch sites from communicating directly with other branch sites, the hub will be used as a backup path in case branch sites cannot communicate directly.

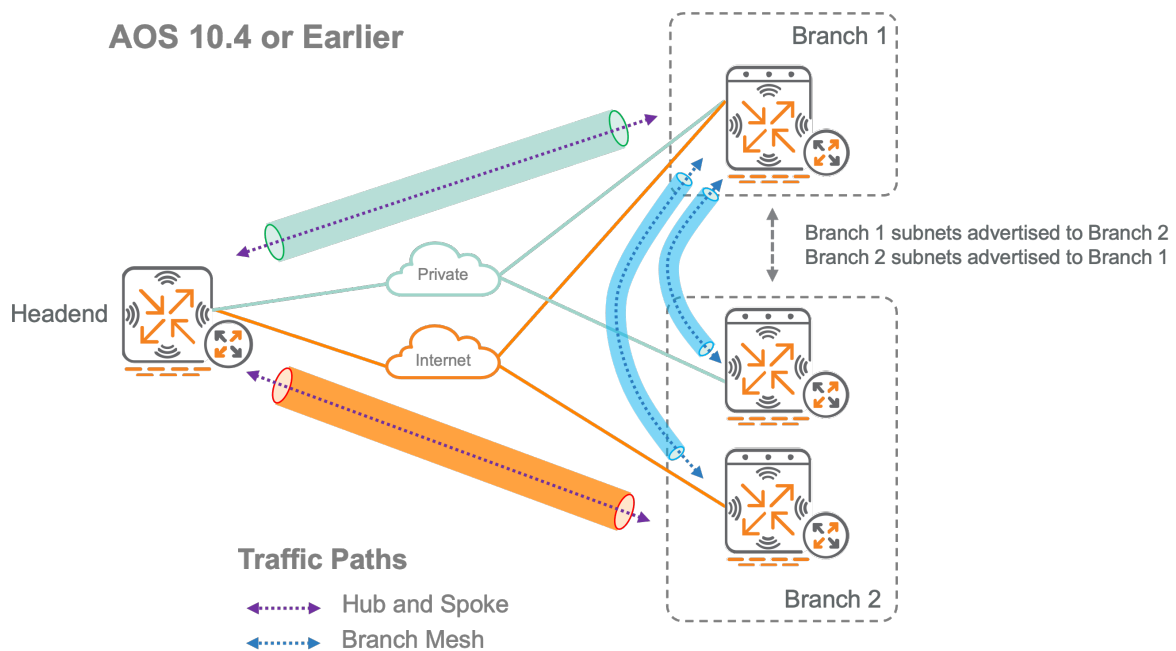


Figure 28: Branch Mesh

System IP Pool

The system IP (system-ip) is a critical configuration element for each gateway operating as a VPNC or BGW. In the three typologies illustrated above, each gateway uses one VLAN interface as its system IP. By default, the Aruba gateway uses this interface to communicate with network services such as RADIUS, syslog, TACACS+, and SNMP. The VLAN interface selected for the system IP on each gateway must have an IPv4 address assigned for the gateway to be fully functional. A gateway cannot initialize fully unless the assigned VLAN interface is active and operational. Central does not allow a gateway to obtain addressing dynamically from Internet service providers using DHCP or PPPoE as the system IP.

Gateway pools can be used to allocate system IP addresses automatically to a dedicated VLAN interface, which is then designated as the system IP address. Each pool includes a unique name along with starting and ending IPv4 addresses. The range of addresses defined for each pool cannot overlap. Aruba recommends configuring one gateway pool for each group, since IP addresses are configured and applied to VLAN interfaces on a per-group basis. The gateway pool must include enough IPv4 addresses to support all the Aruba gateways assigned to the group. Although a group can support multiple gateway pools, specific IP addressing should not be applied dynamically.

WAN Monitoring

The Aruba EdgeConnect SD-Branch solution relies on control-plane communication between gateways and Central, which enables the SD-WAN Orchestrator to negotiate tunnels and establish routes. At least two paths of communication are recommended between the gateways and Aruba Central. Aruba EdgeConnect SD-Branch actively monitors uplink availability to ensure connectivity.

This section presents design considerations for active and passive monitoring and WAN policy.

Active Monitoring

The gateway actively sends UDP or ICMP probes to determine that connectivity to underlay and overlay paths is available.

The gateway also actively monitors the WAN to identify the best path for applications, using one of three operations:

- **Default Gateway Monitoring** - Aruba gateways monitor the state of every WAN circuit by probing their default gateways. A default gateway must be configured on every WAN interface to be considered an uplink. Note that the default gateway does not need to respond to ICMP messages: as long as the WAN Health Check IP/FQDN responds to the probes, uplinks are considered valid.
- **VPNC Reachability** - Gateways send probes to all SD-WAN overlay destinations (through all uplinks) to measure health and state, as well as latency, jitter, and loss. Probes are sent every 2 seconds in batches of five. If packet loss is detected, the gateway switches to aggressive mode and sends 25 probes every 2 seconds to calculate packet loss accurately. UDP probes are managed by the BGW's data path and marked as DSCP 48 to receive priority over other traffic for a more timely response.
- **WAN Health Check** - Gateways send probes by default to the Aruba Path Quality Monitor (PQM), maintained by the Aruba Cloud Operations team. The PQM service is a set of distributed nodes that respond to ICMP/UDP probes. When using the PQM service, admins should set the PQM to UDP mode to measure latency, jitter, and packet loss (ICMP mode does not measure jitter). Admins can specify other health check locations by entering custom IP/FQDN locations. Failure to reach the Health Check responder over an uplink results in failing underlay traffic to a backup uplink. Overlay traffic is determined by probes destined to relevant VPNCs.
- **SaaS Express Optimization** - Branch gateways resolve a specific application using the application's FQDN to query the DNS servers configured on WAN uplinks (or learned through DHCP from the ISP) to determine the best uplink to use for an SaaS application. The probes provide a good measure of how the overlay communications are working, as well as the quality of the last mile for each WAN circuit. Without this monitoring, the gateway would not be able to provide the SaaS express optimization. SaaS express optimization is available only on branch gateways.

Note that business-critical SaaS applications may require a more dedicated method of enriching the user experience. Problems beyond the control of the enterprise network administrator, such as ISP-SaaS peering problems or DNS issues, may adversely affect critical business services.

NOTE:

Active monitoring is always turned on for default gateway monitoring and VPNC reachability.

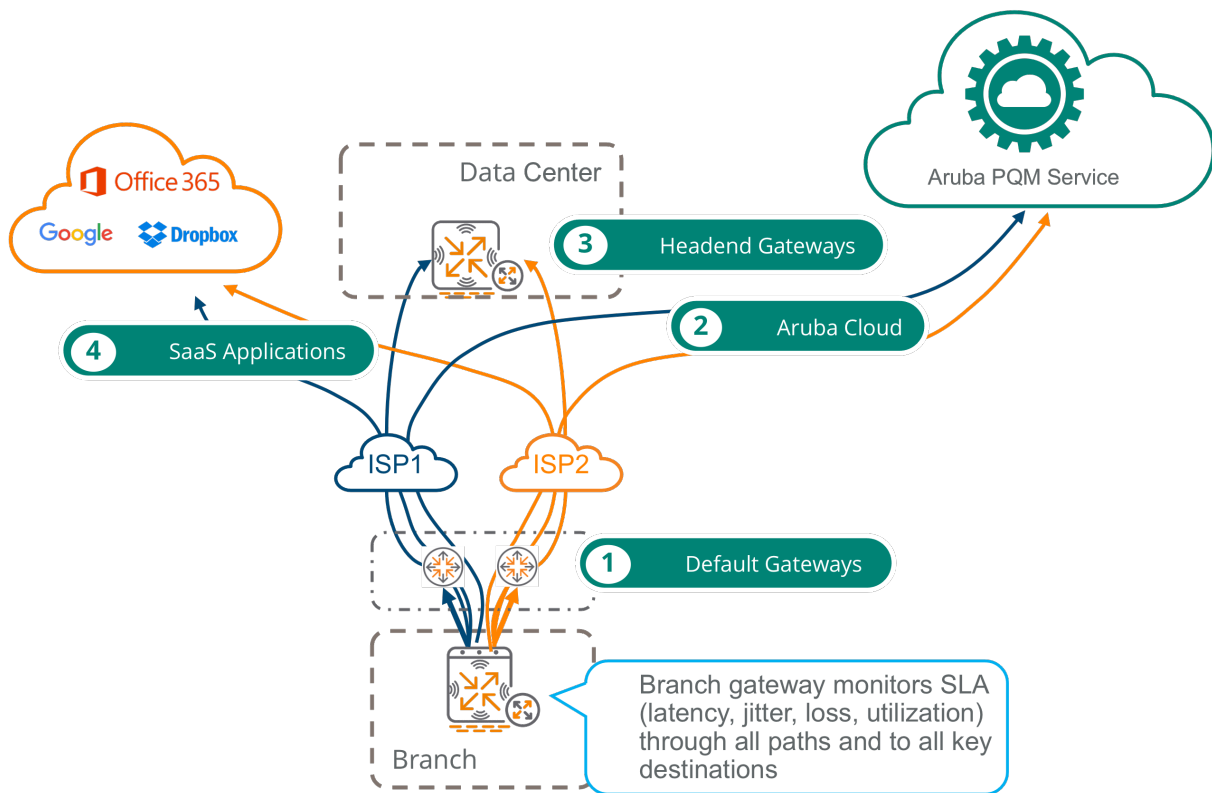


Figure 29: Active Monitoring

Passive Monitoring

The gateway passively monitors the bandwidth usage of the physical interfaces associated with each uplink. Usage is compared with the WAN speed configured on the interface to calculate utilization. If a gigabit interface has 600 Mb of traffic, the circuit is at 60% utilization. The uplink utilization and DPS policies factor in the amount of traffic on every interface when making path decisions.

The gateway monitors the TCP sessions for round-trip time and packet loss on traffic coming and going from clients to SaaS providers. This information is used to calculate a quality of experience (QoE) score for each application. The Central dashboard shows the bandwidth usage, QoE, loss, and latency for each application.

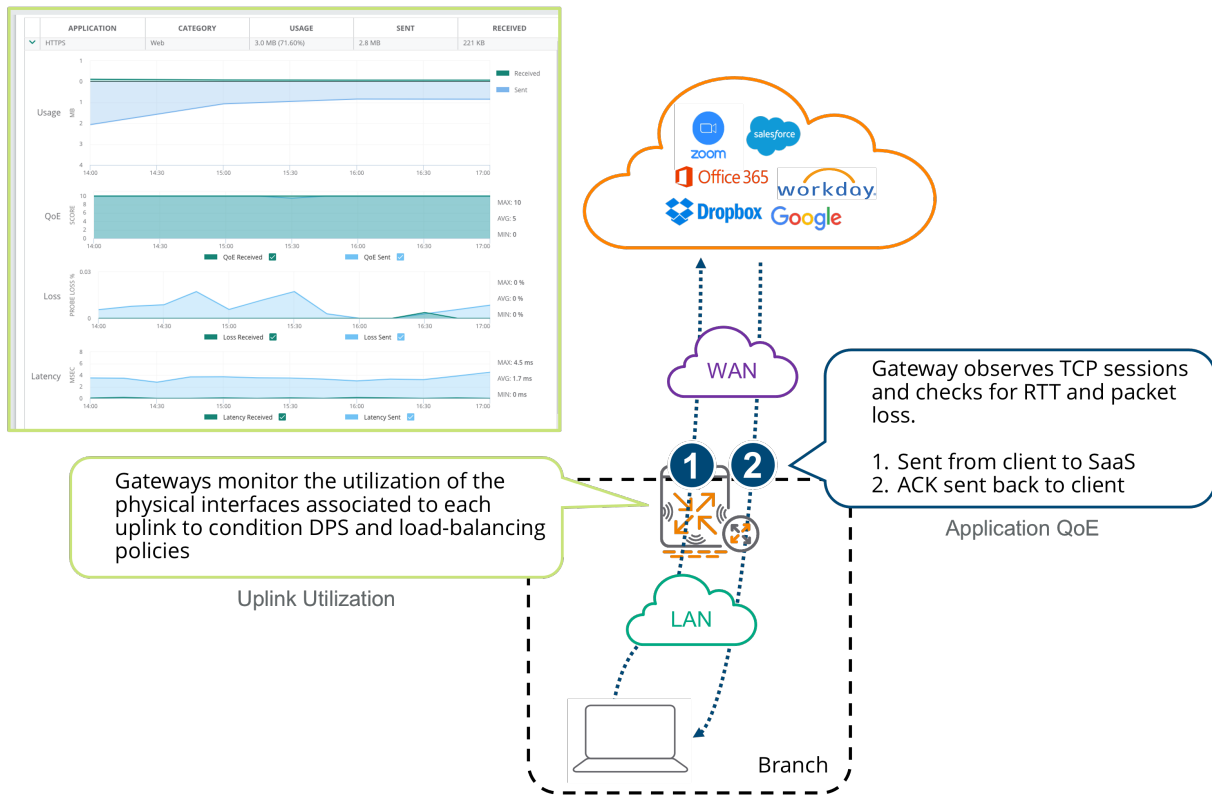


Figure 30: Passive Monitoring

WAN Policies

Aruba's SD-Branch has several WAN policies that help shape the traffic as it traverses the WAN transports at each location. The policies are configured with following features:

- **Policy-Based Routing (PBR):** PBR routes traffic across private or public WAN uplinks based on application and user role if the network destination is not found in the routing table.
- **Quality of Service (QoS):** Role and application-based 802.1p COS and DSCP marking on LAN ingress enables an organization to schedule traffic using a four-class queueing model on the outbound WAN interface. Strict priority queues support real-time applications and deficit round robin (DRR) queues with bandwidth percentages support business-critical applications leaving the gateway.
- **Dynamic Path Steering (DPS):** When multiple WAN links exist, DPS helps choose the best available path for an application based on characteristics such as throughput, latency, jitter, packet loss, and uplink utilization. **(Only available on branch gateways)**
- **Forward Error Correction (FEC):** FEC enables the network to recover easily from packet loss that may be caused by a variety of network layer conditions, such as queue overflows or constrained bandwidth links. FEC is applied on a DPS policy and is needed most when there is loss on the WAN. **(Only available on branch gateways)**

- **SaaS Express Optimization:** Specific applications can be monitored and steered to the best path available based on Observe SaaS traffic as it traverses the GW Firewall to gather latency, loss, and jitter measurements.

Policy-Based Routing

In most deployments, gateways follow the route table when making routing decisions, referred to as destination-based routing. If traffic must be forwarded to a specific overlay tunnel or Internet uplink, PBR enables admins to override the route table for both underlay and overlay traffic. PBR allows admins to use multiple paths by setting the same priority in the next-hop list, which is recommended for fault tolerance. If more than one active path is available, the gateway selects the path using a combination of DPS and load-balancing. A typical use for PBR is to force all traffic to a specific VPN Concentrator or a cloud firewall service. The figure below shows the traffic path when a PBR policy is defined on the LAN ingress.

The most common uses for which PBR policies are implemented include:

- All employee Internet traffic must be routed to the hub-site location to receive additional policy checks.
- Traffic from a specific subset of clients must be forwarded to a specific WAN path.
- Integration with third-party SaaS or unified threat management providers, such as Check Point, Palo Alto Networks, or Zscaler, requires steering certain traffic through a cloud-based security provider.

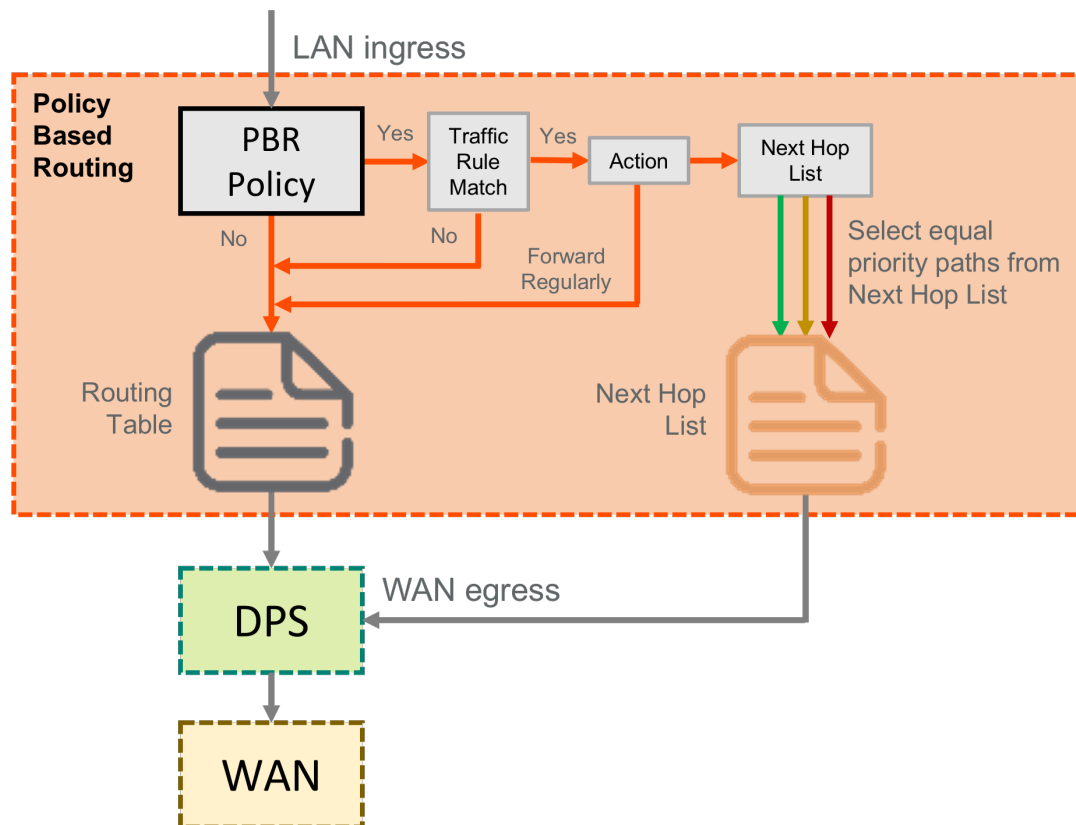


Figure 31: PBR on LAN Ingress

Cloud Security Integration

The branch gateways can redirect selected traffic through a cloud-based security platform such as the Zscaler, and Checkpoint. The integration between cloud security providers and SD-WAN fabrics is discovered automatically, and tunnels and routes are orchestrated based on business and topological requirements.

For information on cloud security providers, see the following guides: [Aruba SD-Branch and Zscaler Internet Access Integration guide](#).

[Aruba SD-Branch and Palo Alto Prisma Access Integration guide](#).

[Aruba SD-Branch and Symantec Endpoint Protection Integration guide](#)

[Aruba SD-Branch Integration with Check Point Intergration guide](#)

NOTE:

Cloud security integration applies to only branch gateways since VPNs generally send traffic from trusted sources.

Quality of Service

Quality of service (QoS) refers to the ability of a network to provide higher levels of service by identifying, marking, and prioritizing traffic. Applying the proper QoS policy is important when the network is congested and limited bandwidth is available. Real-time traffic such as Teams, video conferencing, or business-critical applications have specific latency requirements. When a network is congested, applications can be affected in several ways, including bit rate, throughput, path availability, delay, jitter, and loss. Delay, jitter, and loss can be improved by using the correct QoS policy on the egress interfaces of network devices so applications with higher priority are delivered before applications with lower priority.

Two main strategies can be considered when creating a QoS scheduling policy:

- The first strategy identifies applications that are important to the business and gives them a higher level of service using the QoS scheduling techniques described in this section. Remaining applications stay in the best-effort queue to minimize upfront configuration time and lower the daily effort needed to troubleshoot a more complex QoS policy. If new applications become important in the future, add them to the list of business-critical applications. Updating the QoS level can be repeated as needed without requiring a comprehensive policy change for all applications on the network. This strategy is normally used by organizations that do not have a corporate-wide QoS policy or that are troubleshooting application performance problems across the WAN.
- The second strategy creates a comprehensive QoS policy that identifies all traffic flows and applications using the Aruba deep packet inspection (DPI) engine. The engine can identify more than 3K applications using well-known signatures and protocols. The applications are placed in pre-defined categories in the DPI engine for convenience, but it may be necessary to create custom groupings if the categories do not align with specific organizational needs. This strategy is best suited for organizations that want to use an existing QoS policy with the SD-Branch solution.

Marking and Queueing

The first step in enabling a QoS policy is identifying and marking applications as they pass through a network device. Aruba recommends marking applications with class of service (CoS) for queueing.

Differentiated service code point (DSCP) also can be used for marking. However, DSCP values are not always honored. Check with service providers to ensure that markings are honored.

Applications should be marked using Access Control List (ACL) matching rules. It is recommended to match specific applications using a combination of aliases and TCP/UDP ports or a list of service applications when creating a matching ACL. The combination of both aliases and ports enables administrators to identify applications and mark them more accurately. For less specific applications that still may need some level of prioritization, admins can use other ACL matching methods such as Application Categories, TCP/UDP Ports, or Subnets.

When marking applications, it is important to categorize similar applications so they can be marked the same. For example, Zoom, GoToMeeting, and Teams are all chat/video/voice collaboration tools, so it makes sense to place them in the same category for marking.

After the ACL is defined, the ACL can be applied in two places: on the LAN side of the gateway or within a user role. When tunneling to the gateway, it is important to apply the ACL to user roles, since they are usually encapsulated in a GRE tunnel and a QoS policy cannot remark the incoming packet accurately.

All applications are marked at the ingress of the gateway. If applications are not identified, they are placed in the default queue, with a best-effort level of service. East-west traffic that remains in the location is identified and marked when it passes through the gateway between the VLANs.

After applications are marked, they are placed in a queue with a correlating marking to determine their priority levels. Aruba gateways support four QoS queues: one strict priority queue and three Deficit Round Robin (DRR) queues. Strict priority queues always have all traffic forwarded; other queues are serviced until the priority queue is empty. DRR is a scheduling algorithm that allocates a percentage of bandwidth allocated to each DRR queue for forwarding. Network administrators can define the DRR bandwidth percentages to allocate.

Real-time applications and network management traffic such as OSPF hello packets, etc., should always be placed in a strict priority queue. Business-critical applications should be serviced by one or two of the DRR queues to provide a higher level of service during congested times.

The last queue should be used as the default queue where all unmarked and low priority marked traffic is placed. This queue provides a lower level of service.

The figure below provides an example of marking and queuing.

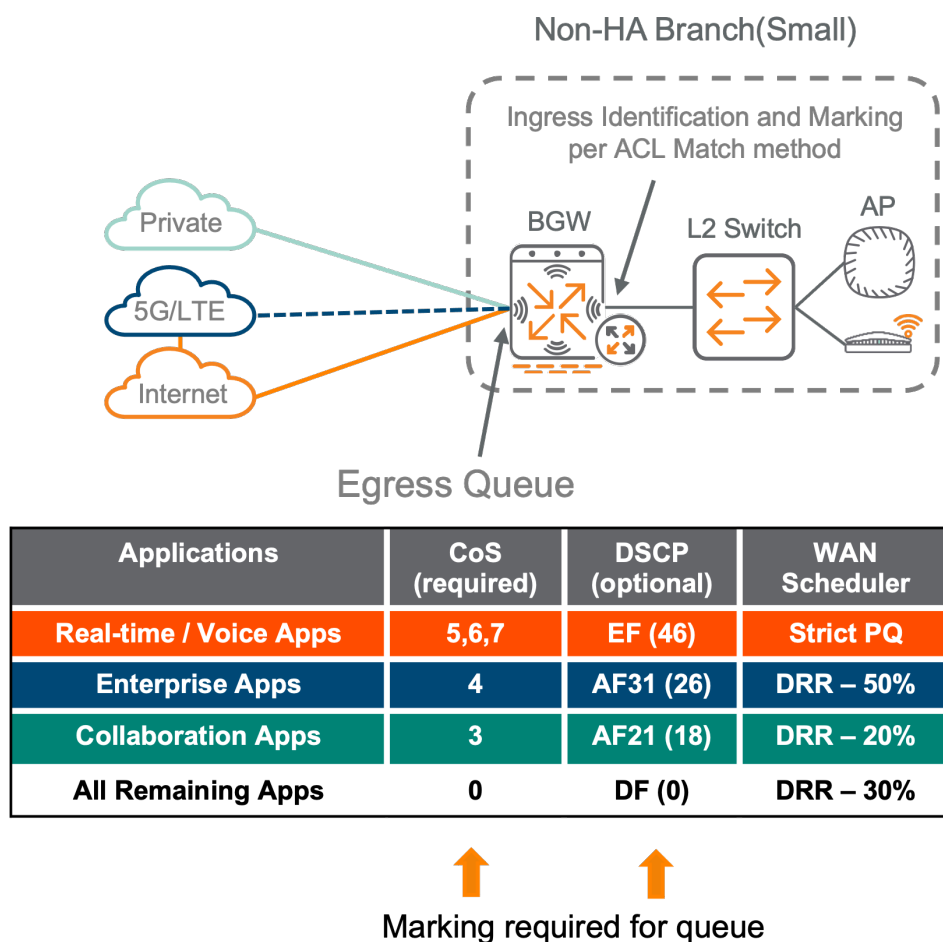


Figure 32: qos_explanation

Dynamic Path Steering (DPS) and Forward Error Correction (FEC)

Using the active and passive monitoring details above, DPS intelligently selects the best uplink for traffic. DPS ensures that applications are sent over the path most appropriate for their service level agreements (SLAs). For example, if a gateway has two paths, Uplink 1 and Uplink 2, and a Cloud SaaS application matches the DPS policy active monitoring criteria, DPS determines which uplink has the best SLA at the moment by comparing latency, jitter, and packet loss statistics from the WAN health check or VPNC reachability probes. In this example, DPS would use only the WAN health check information, since the SaaS application is not hosted at the VPNC site and the relevant path is the Internet. DPS policy uses only the relevant path statistics for each application to determine the uplink for sending traffic,

The network administrator can define SLAs, priority uplinks, and FEC thresholds for a DPS policy. The admin can set the SLA for an application based on traffic categorization, aliases, or IP/Subnet matching criteria. Admins can then use one of the built-in SLAs or adjust the latency, jitter, packet loss, and uplink utilization parameters. The FEC loss threshold is available to delay steering the application based on how much packet loss FEC can handle. For example, under normal circumstances, VoIP would be steered with more than 1% loss, but if FEC is enabled to protect it, steering can be delayed until the loss is 5%.

When configuring a SLA for a DPS policy, it is important to set the SLA thresholds to a point just before applications could register a negative user experience.

The following figure shows the traffic path when a DPS policy is matched on the WAN egress.

NOTE:

The gateway's routing table or PBR rules determine the next hop and the DPS policy selects an uplink.

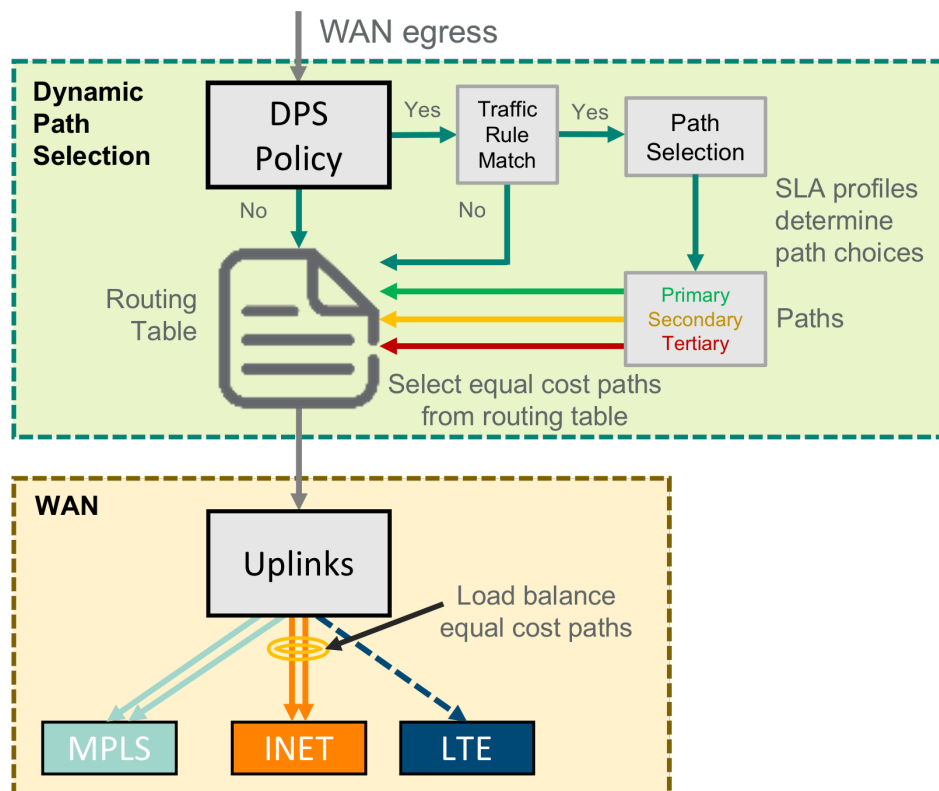


Figure 33: DPS on WAN Egress

After FEC is enabled in a DPS policy, all packets corresponding to the DPS policy are sent to the FEC engine for encoding. FEC enables admins to add a parity packet for every “N” number of packets in each block, where “N” equals 2, 4, or 8 packets. The number of parity packets to be sent per policy depends on the type of resiliency needed for the applications.

FEC parity packets are added to traffic only between a BGW and a VPNC. Even if FEC is enabled on a policy where traffic is destined to the Internet, the FEC parity encoding is not added. Packets sent to a VPNC or BGW are sent to the FEC engine as soon as they are decrypted from the IPsec tunnel.

In the FEC engine, the number of packets received for every “N” packets is checked. If there is no loss, the FEC parity packet is discarded. If a packet is lost or contains an error, the FEC parity packet is used to reconstruct the packet. If more than one packet for a given block is lost or erroneous, the FEC engine cannot reconstruct them.

NOTE:

There is a short wait time of 2 ms between packets in an FEC block to minimize the delay.

DPS policy should be configured to ensure that all traffic traversing the WAN hits the proper SLAs to ensure a smooth user experience. SLAs should be set so that similar applications can be grouped for easier configuration management. The following categories, FEC ratios, and SLAs are recommended.

Application type	FEC Ratio	DPS SLA	Path Priority
SaaS Applications (local-breakout)	n/a	SLA per app recommendations (SaaS Express to pick best exit)	Primary – ALL_INET Secondary – LTE*
VoIP	1:4 5-8 % loss	150 ms delay 30 ms jitter 1% packet-loss	Primary – ALL_INET Secondary - LTE
Other real-time business apps (telemetry)	1:8 5-8 % loss	150 ms delay 50 ms jitter 1% packet-loss BW utilization: 75%	Primary – ALL_INET Secondary - LTE
Business applications	Disabled	150 ms delay 50 ms jitter 2% packet-loss BW utilization: 75%	Primary – ALL_INET Secondary - LTE
Internet applications (local-breakout or exit through cloud security)	n/a	2% packet-loss BW utilization: 75%	Primary – ALL_INET

SaaS Express

As more businesses deploy SD-Branch to take advantage of inexpensive broadband Internet services, and as they adopt Software-as-a-Service (SaaS) applications such as Office 365, Box, Slack, and Zendesk, operations teams must ensure that users at a branch site can connect seamlessly and securely to cloud-hosted applications with the best possible performance. Cloud applications are hosted in multiple geographic locations, so different paths provide different levels of service.

SaaS Express is designed to optimize application performance by probing and steering SaaS applications to the path with the best connectivity. Probing is performed on every available path using the application's FQDN to query the DNS servers configured on the uplink interfaces (or learned via DHCP from the ISP) every 15 minutes. SaaS Express uses the FQDN as the match criteria for a proxy DNS request to the uplink DNS server to ensure that applications do not use non-local DNS servers that could forward traffic to another region and lower performance. The gateways then send the HTTP probes to the application every 10 seconds to measure loss, latency, and jitter for that particular application. Traffic steering is dependent on where the SaaS application exist. Most SaaS applications are broken out locally. If the applications are hosted at a hub site, the gateway follows the route table.

Unlike Dynamic Path Steering, SaaS Express uses the loss, latency, and jitter at the uplinks' exit point to determine the best path. SaaS Express considers the measurement of the full round-trip performance of a SaaS application by probing the FQDN of the application. SaaS Express policies take precedence over DPS policy due to the difference in monitoring. Admins should use SaaS Express with SaaS applications of special interest or when DPS must be used to organize and set a SLA for groups of applications.

NOTE:

In full tunnel situations or when the Internet traffic is sent through a cloud security service, exceptions must be introduced in the routing policies to prevent sending SaaS traffic in the overlay.

SaaS Application Profile Parameter

The gateway supports a set of applications and application categories in the DPI library. The built-in application profiles include a set of SaaS applications such as Adobe, Dropbox, Amazon, Google, Salesforce, Slack, Webex, etc. If a SaaS application is not available in the list, the network administrator can configure it.

Each SaaS application profile includes the following elements:

- **Name:** Name of the SaaS application
- **FQDN:** A list of domain URLs bound to the SaaS application
- **Exit profile:** Traffic steering policy to determine the optimal path exit
- **SLA:** Threshold profile for measuring path quality and performance
- **Health check probe URI:** URI to use for probes to determine the best available path.

NOTE:

For more information on the SaaS Express feature, see the [SaaS Express Feature Guide](#).

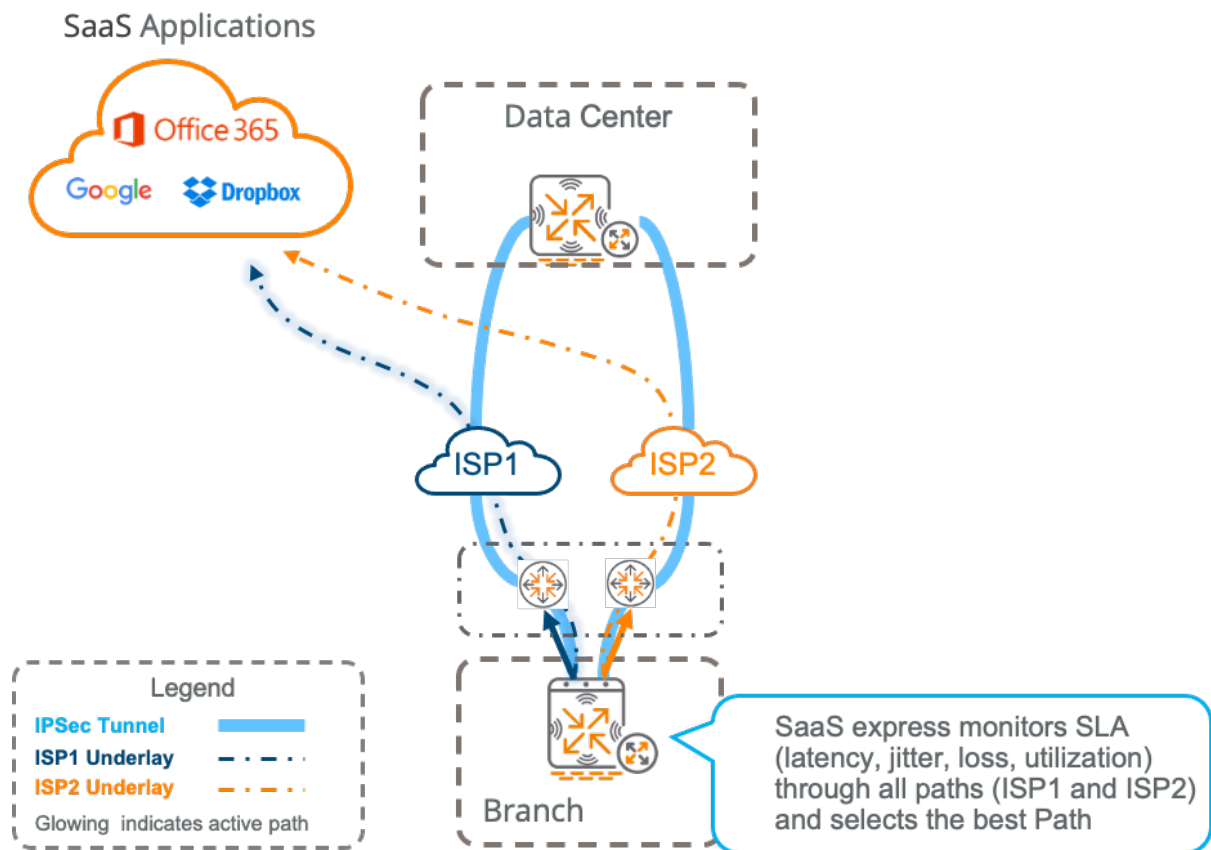


Figure 34: image-20220729170513582

Load-Balancing

The load-balancing algorithm determines how sessions are distributed among the active WAN uplinks. The algorithm kicks in only if route preferences are equal. For DPS and SaaS Express, load-balancing is activated only if the SLAs are the same.

Gateways support the following load-balancing algorithms:

- **Round robin:** Outbound traffic is distributed sequentially between each active WAN uplink. This is the simplest algorithm to configure and implement but it may result in uneven traffic distribution over time.
- **Session count:** Outbound traffic is distributed between active WAN uplinks based on the number of sessions managed by each link. This algorithm attempts to ensure that the session count on each active WAN uplink is within 5% of the other active WAN uplinks.
- **Uplink utilization:** Traffic is distributed between active WAN uplinks based on each uplink's utilization percentage. Uplink utilization considers the link speed to calculate the utilization for a given link and allows the definition of a maximum bandwidth percentage threshold. When the bandwidth threshold percentage is exceeded, that WAN uplink is no longer considered available.

The following figure illustrates the different load-balancing algorithms.

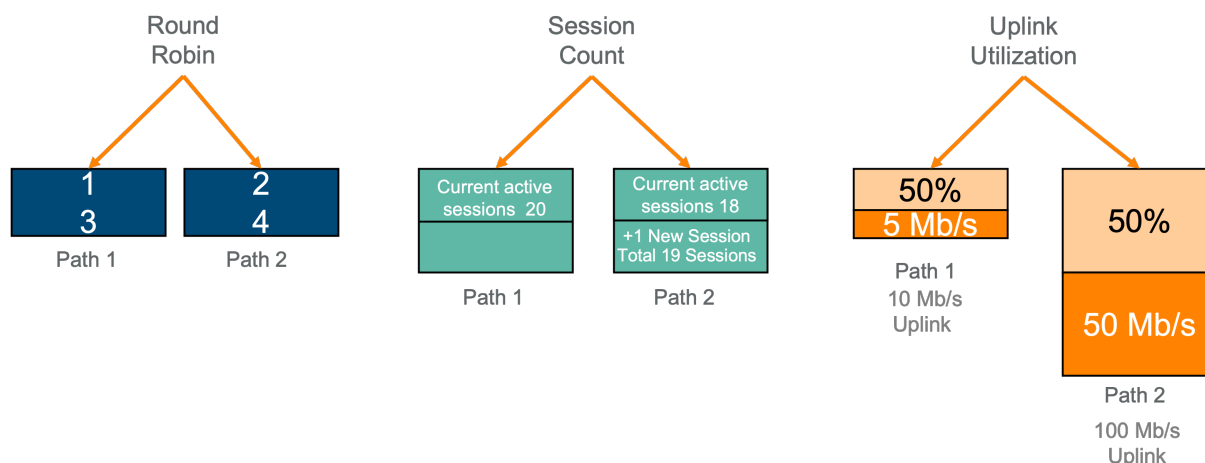


Figure 35: Load Balancing Algorithms

NOTE:

Aruba recommends the uplink utilization algorithm because it accounts for the service speed when making path selections.

Reverse-Path Pinning

When a path is selected for sessions destined for the corporate network through a VPN tunnel, the reverse traffic must take the same WAN path to prevent connectivity problems caused by asymmetric routing. Reverse-path pinning allows the headend gateway to choose the same WAN path for each active session to and from the branch. This is important because the branch gateway selects paths based on performance and SLAs. Reverse-path pinning is performed for corporate sessions originating from the branch destined to the data center, as well as sessions originating from the data center toward the branches.

When traffic originates from the data center, the headend gateway chooses the path based on equal-cost, multi-path algorithms. As soon as the traffic returns from the branch, the BGW steers the 5-tuple session to the correct path based on the DPS policy. When the headend gateway sees the return traffic, the session is updated to use the chosen path for the duration of the flow.

- The headend gateway selects an available WAN path using equal-cost, multi-path routing.
- If the WAN path matches the preferred path defined in the BGW's DPS policy, no additional steering is required.
- If the WAN path does not match the preferred path defined in the DPS policy, the branch gateway sends the return session over the preferred path. After receiving traffic from the new path, the VPNC steers the outbound session to the preferred path to maintain symmetry.

The figure below shows traffic from a branch location over the private WAN overlay tunnel and the reverse path pinning feature on the VPNC that returns the traffic on the same path to enforce symmetry.




Figure 36: Reverse-Path Pinning

EdgeConnect SD-Branch Hub Design

This section covers the following aspects of hub design:

Physical or Virtual Gateways

When designing the hub site, the first major decision is determining if the hub is physical or virtual. The organization must consider the following factors:

- **Physical gateways** - On-premise gateways are generally required for organizations with an on-premise data center site.
- **Virtual gateways** - Virtual environments can be used by organizations that use a cloud provider for Infrastructure as a Service (IaaS) or other private cloud workloads.

Cloud Provider Integration

Integrating with a cloud provider is primarily deployed in two different ways: by using a cloud provider or by using an MSP that peers with a cloud provider, as follows:

- **Cloud Integration** - Aruba gateways can be deployed directly within Azure, Google Cloud, or AWS for cloud connectivity that provides SD-WAN capabilities such as Route Orchestration, FEC, DPS, VIA, and Microbranch connectivity.
- **MSP Integration** - Aruba gateways can be deployed within a MSP/Colocation where the MSP has direct low-latency connectivity to a cloud provider. This hybrid approach provides the benefit of having the physical infrastructure and cloud connectivity in the same location.

Aruba gateways can be used to peer directly with a transit gateway to establish connectivity between VPC/VNETs. This is **not** recommended because organizations will **lose** the following SD-WAN capabilities:

- **Reverse Path Pinning** ensures that traffic always returns through the path of origin, enabling Branch Gateways (BGWs) to perform uplink load-balancing and Dynamic Path Steering.
- **Forward Error Correction** protects critical traffic flows from potential network issues between the branches and the cloud, especially when traversing the Internet.
- **Tunnel Orchestration** automates the process of establishing IPsec tunnels from all BGWs to all relevant VPNCs (including the vGW).
- **Orchestrated Routing** automates the exchange of routes across the SD-WAN.
- **End-to-End Visibility** enables single-source visualization and monitoring of the entire SD-WAN network using a single application (Aruba Central).

When integrating with a cloud provider, it is important to deploy a virtual gateway to ensure the high performance and stability that SD-WAN provides.

For more details on cloud integration see the following guides:

[Aruba SD-Branch Integration with AWS Public Cloud](#)

[Aruba SD-Branch Integration with Azure Public Cloud](#)

[Aruba SD-Branch Integration with Google Cloud](#)

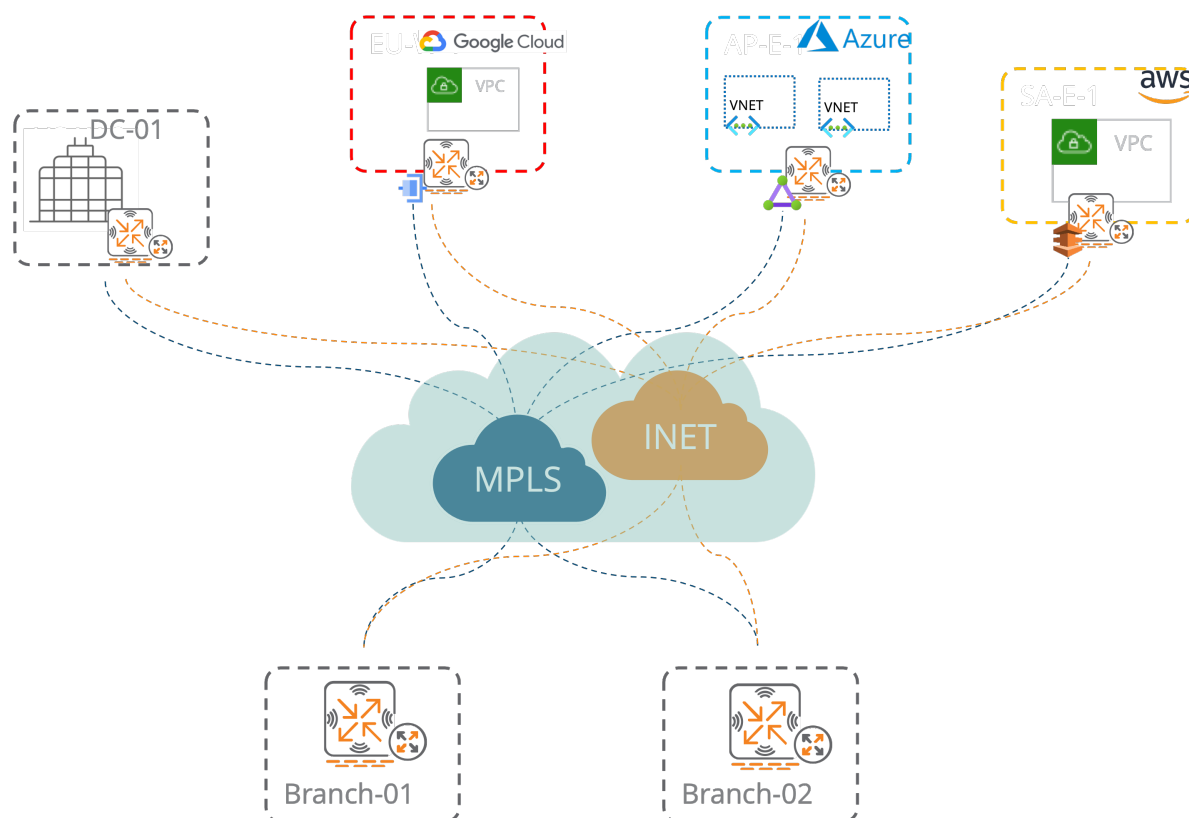


Figure 37: Cloud_Integration

One Touch Provisioning

Zero Touch Provisioning (ZTP) is the preferred method to deploy gateways that obtain IP addressing dynamically from an Internet Service Provider (ISP). The gateway connects to the Internet service, obtains an IPv4 address, then communicates with Aruba Central to obtain the configuration.

There may be circumstances when a gateway requires additional configuration before it can communicate with Central. The gateway may require:

- Static addressing
- Point-to-Point Protocol over Ethernet (PPPoE) credentials to initiate the Internet service
- Specific VLAN configuration.

For these deployments, Aruba offers a One Touch Provisioning (OTP) feature for gateways. OTP can use the serial console or the web user interface.

This method is recommended when a VPNC or BGW requires more advanced configurations that may require a specific VLAN ID or trunk configuration. Generally, when deploying a hub site, OTP is needed due to its placement. The OTP feature is available only for gateways in their factory default state and cannot be accessed after a gateway has received its configuration from Central.

Headend Redundancy

The SD-WAN Orchestrator brings up active tunnels to multiple headend gateways. The routes from the BGWs are advertised to the gateways. The overlay routes use weighted costs to select one gateway over another. Weighted costs are set on a group that enables traffic load-balancing by creating two branch groups with inverse priorities.

On the northbound LAN interface side, the overlay cost is automatically translated into the dynamic routing protocols as follows:

- OSPF: Direct translation into External 1 and External 2 cost
- BGP: Direct translation into Multi-Exit Discriminator
- BGP: Automatic prepending of Autonomous System numbers to ensure routing symmetry.

The figure below shows a headend site with redundancy and two gateways.



Figure 38: Layer 3 Redundancy

Multiple Active Data Centers

Aruba supports multiple active data centers to provide branch locations with easy access to resources in different locations. Tunnels are built to all data centers using the method described above.

To minimize the number of routes in the data center gateways, Aruba recommends summarizing the branch site routes advertised into the data centers and data center routes into the branch.

Proper IP address planning is required across the entire organization, so the number of subnets at each branch falls within an easily summarized bit boundary. If three or four subnets are currently in use at a location, plan for a minimum of eight summarized subnets to allow for future expansion without adding new summaries. A good rule of thumb is using a network range with a 255.255.248.0 mask or /21, for each branch location. This allows for eight /24 subnets at each location.

At the data centers, the routes also should be summarized to route table constraints on the branch gateways. In most cases, it is preferable to use a single supernet route for each data center location. If this is not possible, use as few summary routes as needed. Also consider creating a single summary route to cover all of branch locations.

Another recommendation is setting the DC preference for each branch to the closest hub location. The secondary and tertiary locations use lower DC preferences, so the closest one is always preferred. This enables the branches closest to a particular data center to use it as a regional hop to other branches in the area.

Aruba always recommends allowing branch-to-branch communications through the data center, even when planning to use branch mesh, as discussed below. This enables the closest data center to act as a backup path between branches if the branch mesh tunnel is not available.

The figure below shows the summary routes and DC preference with multiple active data centers.




Figure 39: Multiple Active DCs

To maintain traffic symmetry when using multiple data centers, the SD-WAN Orchestrator automatically sets different routing costs for the different VPNCs in increments of 10. The smart redistribution feature between VPNCs in different data centers works the same as two redundant VPNCs in the same data center, as discussed in the previous section.

EdgeConnect SD-Branch Branch Design

This section covers the following aspects of branch design:

Zero Touch Provisioning

Zero Touch Provisioning (ZTP) is the preferred method to deploy a branch gateway. ZTP automatically communicates with Central and downloads the branch's group configuration without requiring user interaction. Group assignments for gateways are made in Central before provisioning, so the group configuration can be applied without the need for admin personnel to move the gateway from the default group or perform configuration functions.

When deploying a new branch site, remember that the gateway(s) must be provisioned and online before APs and switches at the branch can complete the ZTP process and receive their group configuration. The APs and switches must receive IP addressing and domain name server (DNS) information from their respective gateways before they can communicate with Central. In order for a gateway to perform ZTP successfully, one or more ZTP ports must be connected to a WAN service that provides the following:

- DHCP addressing
- DHCP options 3 (Router) and 5 (Name Server)
- Internet access.

Physical ports should be consistent across BGWs to ensure consistency between branches and reduce the number of groups required. The objective is to select a common set of ports that work for as many branch configurations as possible.

WAN uplinks can be connected to any switchport on a BGW except for Gigabit Ethernet 0/0/1 which is reserved for One Touch Provisioning (OTP).

If the gateway is connected to multiple WAN services, it will attempt to perform ZTP over each service until it is successful. The first WAN service to respond to the DHCP's discover message is selected first. The figure below illustrates the interfaces reserved for OTP for some of the different form factors.

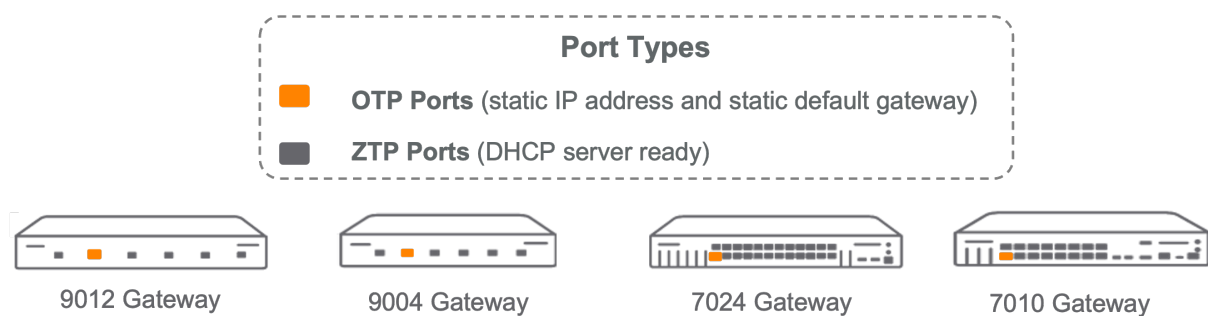


Figure 40: ZTP_Ports

The video below explains the topics covered in the Zero Touch provisioning section.

Sites, Clustering and System IPs

Branch gateways are generally deployed as a single device or as a redundant pair. Both deployment methods are placed automatically in a cluster based on the site. The cluster can contain one or two gateways. Automatic clustering works only for a maximum of two branch gateways at site. For both redundant and non-redundant sites, if gateways are not placed in a cluster, the access points are not able to tunnel to the gateway. It is always important to place both devices in a site, cluster automatically.

CAUTION:

Do not use auto group clustering, this will require each site to have a unique group. This will make managing a standard configuration very cumbersome.

Non-redundant sites and redundant sites vary slightly in the manner the system IP address is used.

- In a non-redundant site, the gateway uses the system IP address as the source for functions such as Radius, SNMP, Tunneling, etc. See [EdgeConnect SD-Branch Overview](#) for more details.
- In a redundant deployment, a VRRP Virtual IP (VIP) is required and serves as the source IP address for Radius and TACACS functions. The gateway proxies all COA requests sourced from the VIP. The system IP address is still used for functions such as tunneling and cluster formation.

“Split brain” situations can occur if LAN connectivity is interrupted and the gateways are unable to communicate with one another. Both gateways assume primary functions and use VIP. After LAN connectivity is restored, the secondary cluster member no longer assumes the primary VIP role and cedes the primary role after 280s. Gateways should have layer 2 connectivity, either directly or through an intermediary switch. Gateways also should use LACP in both cases for fault tolerance.

The video below explains the topics covered in the Sites, Clustering and System IPs section.

WAN Redundancy

It is very common for branch sites to have limited WAN drops on-site in deployments with more than one branch gateway and limited WAN transports.

Administrators can take advantage of uplink sharing. This enables the branch gateways to sync DHCP state, and share uplinks between each device over the LAN. Each branch gateway can share selective uplinks, no uplinks, or all available uplinks. When sharing uplinks, it is critical to use a unique VLAN for each share uplink. For example, the INET uplink should be VLAN 4085 and the MPLS uplink should be VLAN 4086. The VLANs can be mutually shared, as illustrated below.

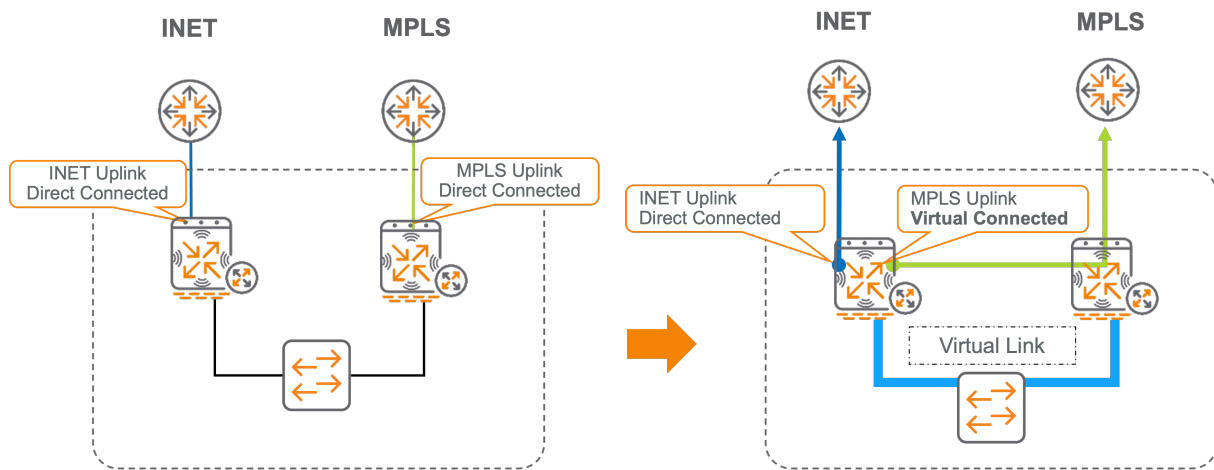


Figure 41: Uplink_Sharing

Administrators identify the uplink interfaces to share based on the VLAN ID. If administrators want to take advantage of DHCP state sync between gateways without sharing uplinks, use the same VLAN ID for WAN uplinks on both gateways. In the illustration below, the VLAN 4085 is in use for both BGW1 and BGW2. That uplink is also connected. In this example, the WAN uplink is not shared.




Figure 42: Uplink_Sharing

Branch Security

Branch gateways can serve as a site's first line of defense against external and internal threats when implemented with the following features:

- **Aruba Role-Based Stateful Firewall** - The firewall supports scalable configuration using firewall aliases, ALGs, and role-based policies.
- **Deep Packet Inspection** - Qosmos's application engine and signatures provide the capability to identify nearly 3500 applications.
- **Web Content, Reputation and Geo-location Filtering** - WebRoot's machine learning technology classifies content, reputation, and geolocation for billions of URLs.

- **Aruba Threat Defense (IPS/IDS)** - Powered by ProofPoint's Threat Intelligence, Aruba 9000 Series gateways can perform IDS/IPS functions for all branch traffic. See the [IPDS Guide](#) for more detail.

Before employing the security features available for the Aruba branch gateway, it is especially important to identify critical business assets, such as groups of users, point-of-sale systems, and applications (cloud based or on-premise). Assets must be clearly identified before comprehensive protection requirements and security protocols can be defined.

On a branch gateway, assets are identified clearly by their destinations, services, and individual roles:

- **Network Aliases** are used to identify and differentiate a specific host, a specific network, or a combination.
- **Service Aliases** are used to identify and differentiate services, such as DHCP, FTP, SIP, etc.
- **Applications** are identified clearly so they can be targeted individually or as a group for implementation of specialized security operations. For example, **Deep Packet Inspection** can be performed on a list of more than 3000 common applications, such as YouTube, Facebook, and Twitter.
- **Roles** are used to identify and categorize end user devices. Access control lists and policies are applied using assigned roles.

After assets are clearly identified, the organization must map which assets communicate with one another. For example, employees likely need to communicate with Office 365 or Google Workspace, to complete daily tasks. Create flows to outline the connections between assets.

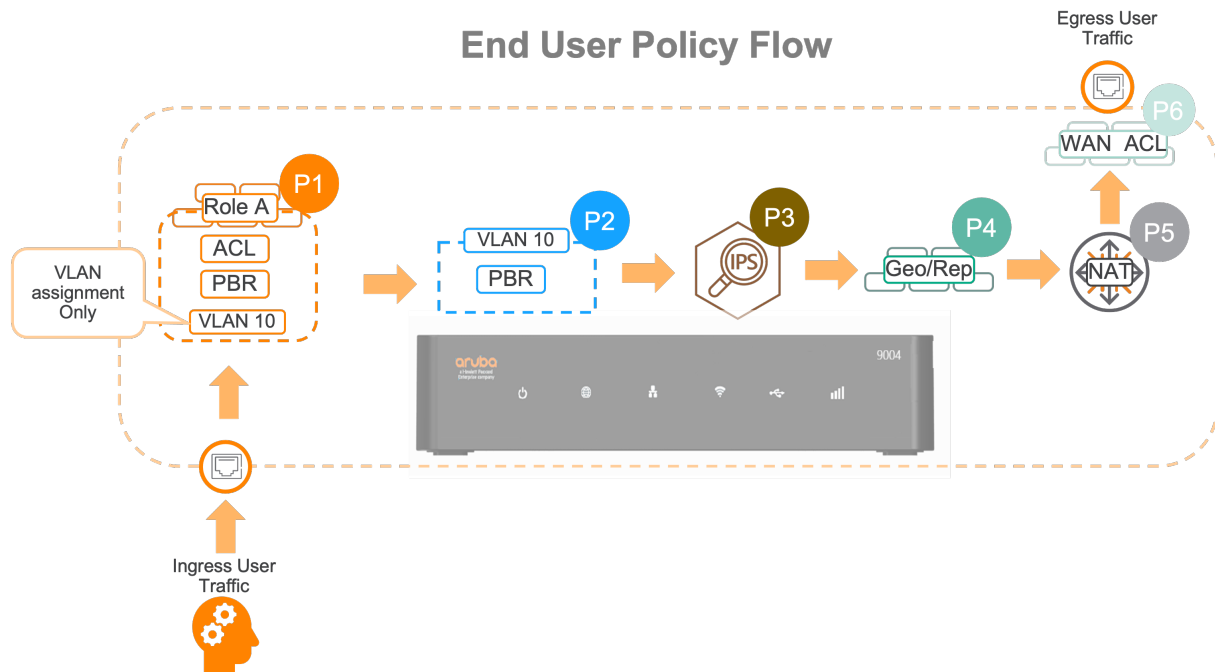
The next step is developing and enforcing policies based on the assets and their communication flows. The branch Gateway can use four methods to enforce policy: Geolocation Filtering, Web Reputation Rating, Intrusion Prevention Systems, and Access Control Lists (ACLs).

To enforce policy using Geolocation Filtering, Web Reputation, and Intrusion Prevention Systems, the gateway must have the Firewall Visibility, Deep Packet Inspection and Web Content Classification (Webcc) functions enabled. Also note that Geolocation Filtering has an implicit default "allow all" by default, so additional configuration is required to block unwanted countries.

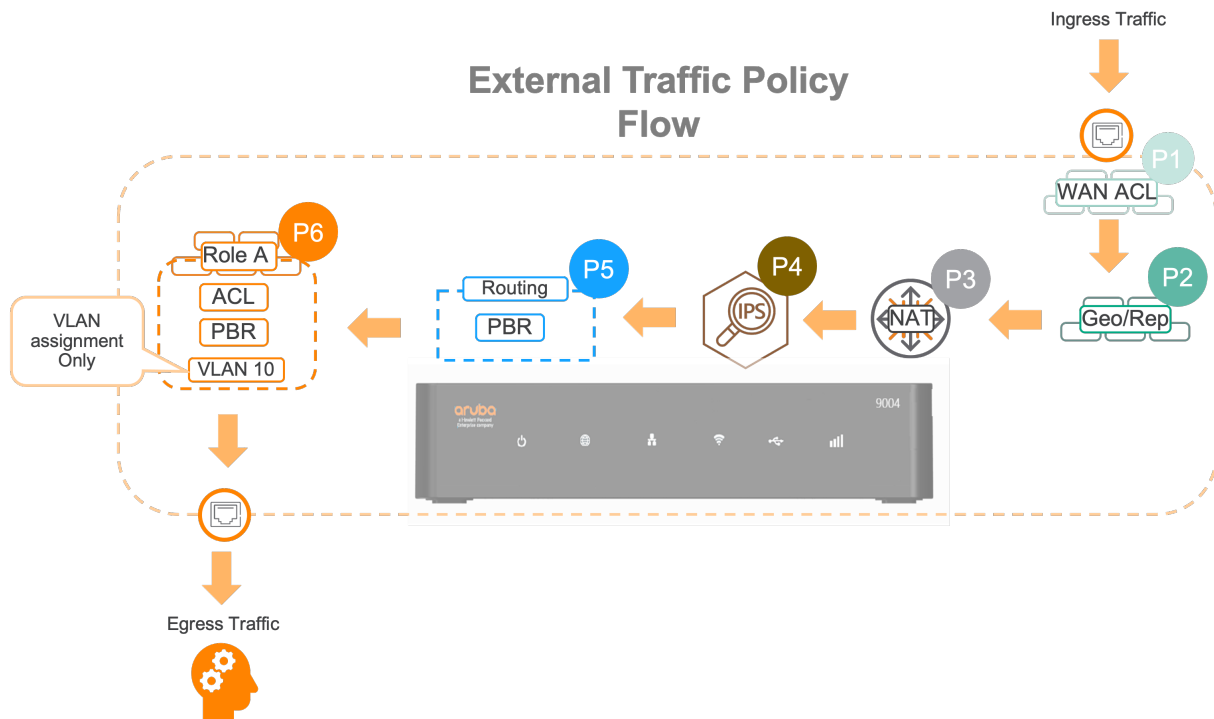
Web Reputation settings should be established to block malicious URLs with both inbound and outbound policies. If a legitimate website is accidentally blocked by the settings, it can be added individually to the list of allowed URLs. Logs are available for administrators to review the blocked URLs based on Web Reputation and identify URLs to unblock when required.

Depending on the direction of the traffic different enforcement rules take precedence. ACLs can be applied specifically to interfaces and roles, each with their own enforcement precedence.

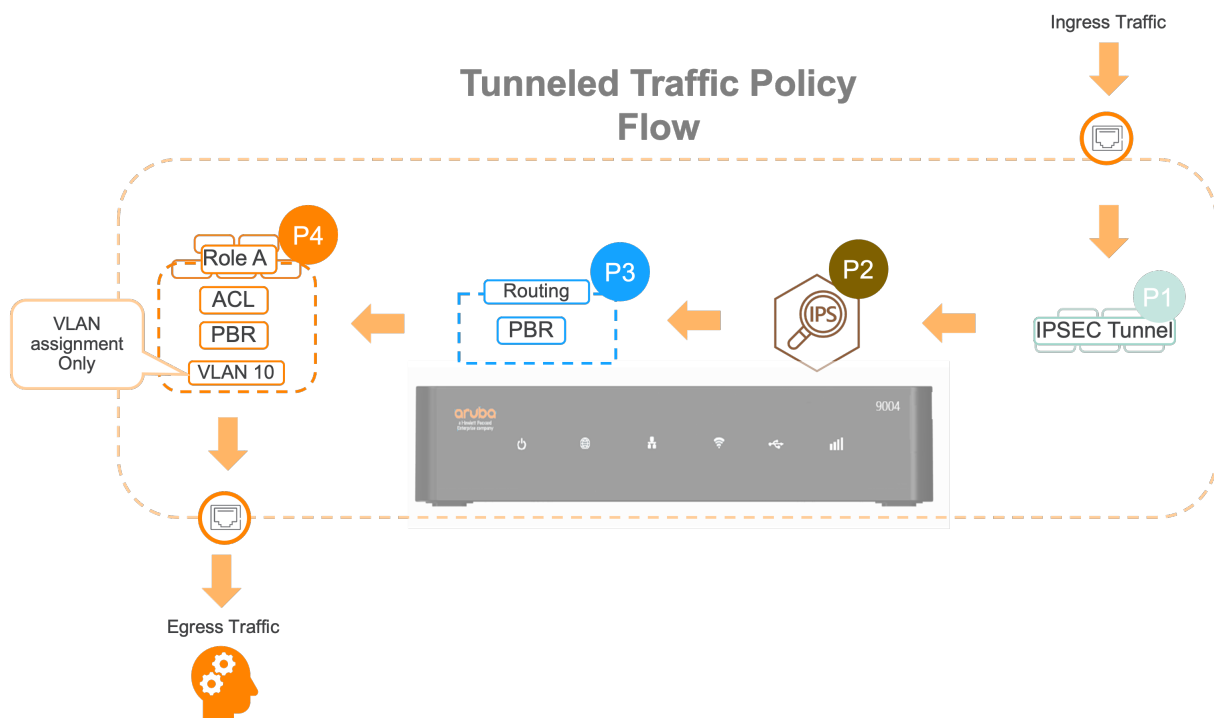
In the image below, Role A has a VLAN assignment ACL and routing policy applied. The route policy within the role takes precedence over the route policy applied to the VLAN. The next policy hit is the IPS/IDS which inspects traffic to verify it is not malicious. Depending on the organization's policy requirements, the gateway's IPS or IDS can be enabled. The IPS (prevention) blocks malicious traffic and the IDS (detection) logs instances of malicious traffic. If traffic is not determined malicious, it continues to the Geo Filtering/Web Reputation system. Finally, the packet is NATed or PATed through the firewall where it hits the WAN ACL.



When traffic is coming into the Gateway from its WAN interface, the first policy hit is the WAN ACL. From there, incoming traffic is routed through the Geo Filtering/Web Reputation system. Traffic then hits the NAT (1:1), then moves to IPS/IDS for inspection. The traffic then hits a global routing policy which is slightly different from the example above. Finally, the traffic hits the user role for its final ACL and policy checks.

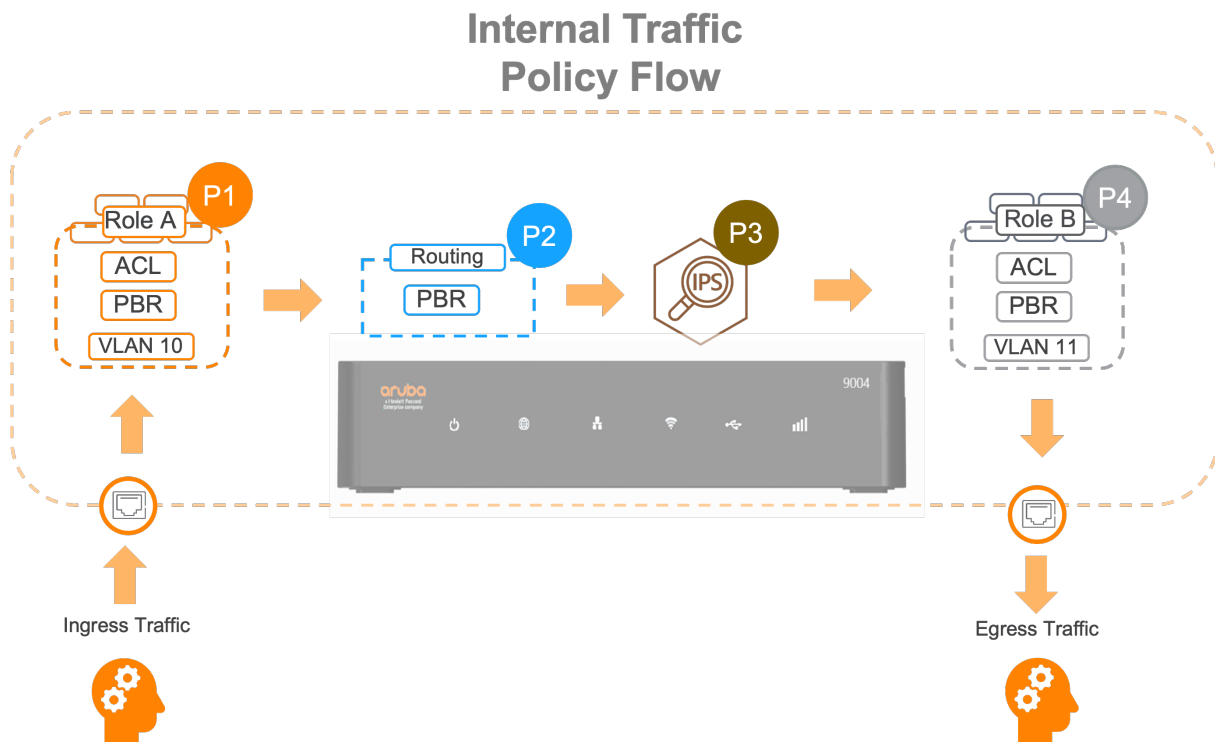


When traffic is received from a tunneled endpoint this could be a VPNC, or another branch site. The enforcement points for traffic is slightly different. Flowing directly to the IPS/IDS, then to the global routing policy and finally to the user role.



Within a gateway, each user role is tracked as an instance in the stateful firewall. Any traffic going between roles within the same gateway is inspected on ingress even when on the same gateway. Additional security for untrusted ports is described in detail in the “Trusted Ports” section.

In the example below, the role initiating traffic (Role A) passes through the stateful firewall. ACLs are enforced for the role and IPS/IDS inspection is performed. The traffic is routed to the destination role (Role B), where it must pass through another stateful firewall, with ACL enforcement for Role B.




Trusted Ports

Each port on a gateway can be configured as Trusted or Untrusted. The trust parameter determines how the gateway processes incoming user traffic. When a port or VLAN is configured as untrusted, the gateway tracks user sessions for each IPv4 address. User devices are assigned one of the predefined roles that determines the session ACLs applied to incoming and outgoing user traffic. The trust configuration for each port and VLAN depends on the role of the gateway and what is connected to that port. Aruba recommends the following trust configuration for VPNC and BGW ports:

- **VPNCs** – Configure all ports and VLANs as trusted.
- **BGWs** – Configure all WAN ports and VLANs as trusted.
- **BGWs** – Configure all LAN ports as untrusted.

There is no reason to track user sessions on VPNCs, so all ports and VLANs are configured as trusted. This also applies to BGW and VPNC ports that are directly connected to WAN services. If a VPNC or BGW is directly connected to a public WAN service, Aruba recommends configuring and assigning a restrictive session ACL to the port and configuring all BGW LAN ports as untrusted. This results in BGW tracking of all user sessions for internal IPv4 addresses. Each LAN/VLAN requires its own AAA profile. Each AAA profile can trigger MAC, 802.1X, or captive portal authentication and determine an initial role assignment. Device and user authentication is completely optional.



Aruba SD-LAN

After the SD-WAN is deployed, LAN side, also called the SD-LAN, the branch gateway must be designed. Generally, the SD-LAN design aligns with the campus **Two-Tier** and **Three-Tier** designs, except that applications are not hosted on-site.

This section covers the same architectures with a focus LAN interaction with the SD-WAN side and how deployment is affected at a branch site.

L2 and L3 Boundaries

When deploying a branch site, the LAN can go two directions, depending on the size of the branch. Organizations can use layer 2 or layer 3 from the branch gateways down, or a combination of both. In either case Centralized overlays can be leveraged or traditional segmentation methodologies.

Two/Three Tier Layer 2

- Allows for simple deployment of sites
- Increases broadcast domains
- Requires more configuration on BGW (the BGW would be the default gateway for all VLANs)
- Involves spanning tree complexity

Three Tier Layer 3 (routed to the BGW Layer 2 to access)

- Increases scale
- Reduces broadcast domain
- Minimizes spanning tree complexity
- Requires more switch-level configuration (the switch would be the default gateway for all VLANs)
- Requires more complex deployment of sites (point-to-point routes)

Each deployment type has benefits and drawbacks. Generally, for a smaller branch site with Two-Tier architecture, layer 2 facilitates bringing up a site quickly with minimal scale issues.

Organizations should consider using layer 3 if they have a Three-Tier architecture. Layer 3 increases the deployment complexity; however, it reduces the spanning tree complexity that can arise in a larger L2 domain.

Layer 2 Wired Access

In this design, the BGW provides layer 2 services for the site. The switches use VLANs for segmentation, enabling identical configuration of access switches to reduce design complexity.

Design Considerations

- Cluster gateways using (Site Cluster).
- Enable Default-GW mode.
- System IP is required for cluster formation, AP and switch tunneling.
- Both gateways must be able to reach each other through all configured VLANs for VRRP communication .
- Radius will be sourced from the VRRP, Virtual IP address.
- Use Default gateway mode to prevent asymmetrical routing. Traffic will be pinned to VRRP preferred member, will failover to the backup member.
- The BGW can be the DHCP server or point to a centralized DHCP server.
- Easy ZTP for switching infrastructure
- Connected VLANs should be redistributed into the SD-WAN overlay.

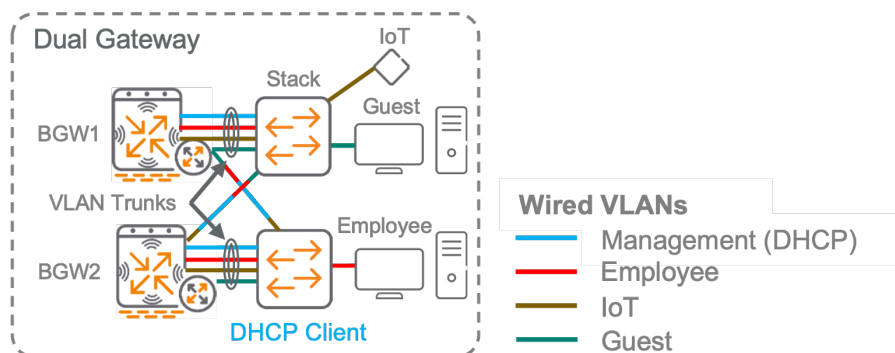


Figure 43: Non-Tunneled L2 Wired Access

Layer 2 Centralized Overlay

The following design uses user-based tunneling (UBT) also known as a centralized fabric. Access switches, and AP's tunnel clients to the BGW to apply user roles and for layer 3 termination. Administrators can take advantage of the user roles applied to clients to enforce role to role policy and designate WAN policies using Deep Packet Inspection (DPI). For example, guest traffic can be completely segmented and broken out directly to the Internet egress or to a cloud security provider.

Design Considerations

- Cluster gateways using (Site Cluster).
- Enable Default-GW mode.

- System IP is required for cluster formation, AP and switch tunneling.
- Radius is sourced from the VRRP, Virtual IP address.
- Use Default gateway mode to prevent asymmetrical routing. Tunnels go to the VRRP preferred member, tunnels will failover to the backup member.
- Policy is enforced at the BGW
- Gateways should run LACP trunks to the L2 aggregation block and trunk all tunneled VLANs. Gateways must have L2 adjacency for VRRP and Cluster formation.
- VLANs between the aggregation and access must have different VLAN IDs and subnets from the tunneled VLANs. (These VLANs can **NOT** be tunneled)
- All devices must be on the same management VLAN, do not tunnel this VLAN.
- Run LACP trunks between the Aggregation and BGWs.
- MTU/Jumbo frames **MUST** be enabled on all VLAN's

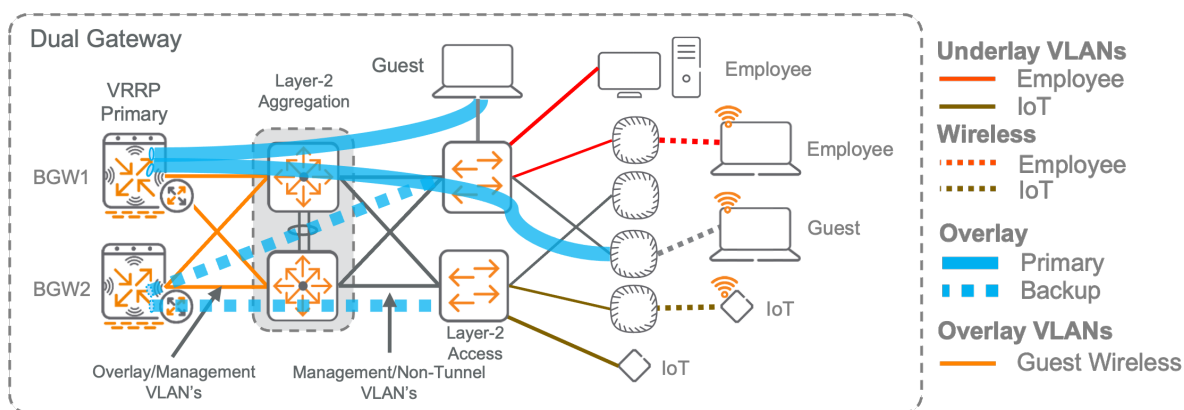


Figure 44: L2_tunneled_ubt

Layer 3 Wired Access

In this design, the aggregation switches provide layer 3 services for the site. The layer 2 access switches use multiple VLANs trunked to the aggregation switches to map the VLANs between them. The aggregation switches act as the IP default gateway for each IP subnet and provide DHCP services to the end devices. DHCP also can be centralized at the headend location or in the data center. The layer 2 access switch obtains its IP address using a DHCP client on the management VLAN. The aggregation switches are routed to the BGWs using layer 3 ports.

Design Considerations

- Clustering is not needed in this design.
- Gateways should act only as edge routers/firewalls.
- OSPF should be used to redistribute routes into the SD-WAN overlay.

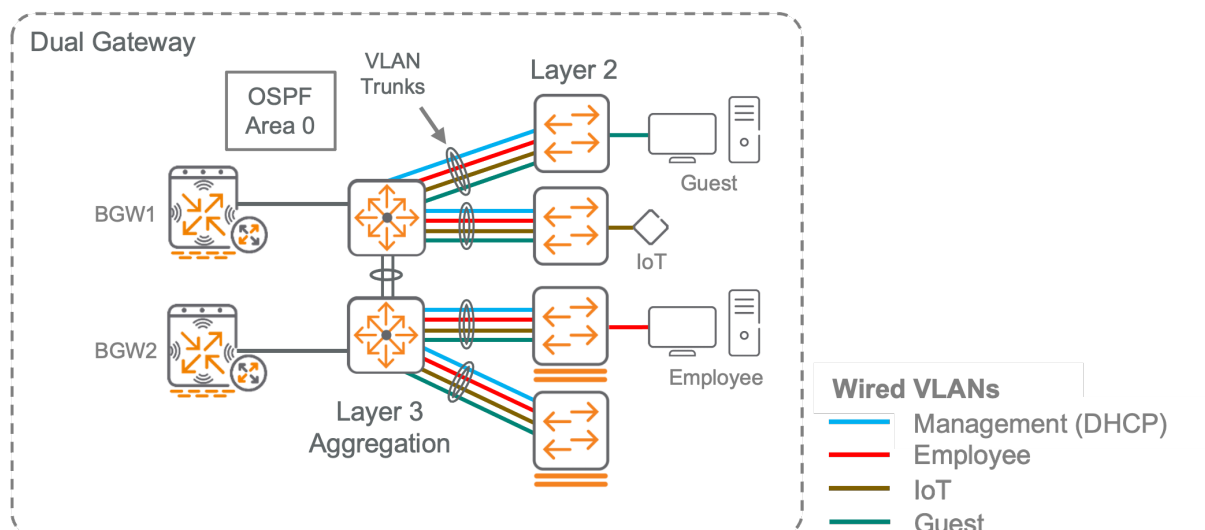


Figure 45: Non-Tunneled L3 Wired Access

Centralized Overlay

The following design uses user-based tunneling (UBT) also known as a centralized fabric. Access switches, and AP's tunnel clients to the BGW to apply user roles and for layer 3 termination. Administrators can take advantage of the user roles applied to clients to enforce role to role policy and designate WAN policies using Deep Packet Inspection (DPI). For example, guest traffic can be completely segmented and broken out directly to the Internet egress or to a cloud security provider.

Design Considerations

- Cluster gateways using (Site Cluster).
- Enable Default-GW mode.
- System IP is required for cluster formation, AP and switch tunneling.
- Radius is sourced from the VRRP, Virtual IP address.
- Use Default gateway mode to prevent asymmetrical routing. Tunnels go to the VRRP preferred member, tunnels will failover to the backup member.
- Policy is enforced at the BGW
- Gateways should run LACP trunks to the L3 aggregation block and trunk all tunneled VLANs. Gateways must have L2 adjacency for VRRP and Cluster formation.
- VLANs between the aggregation and access should have different IDs and subnets from the tunneled VLANs. (These VLANs can **NOT** be tunneled)
- OSPF should run between the LACP trunks to Aggregation and BGWs to learn about non-tunneled subnets.
- OSPF should **NOT** advertise the tunneled VLANs to the aggregation switch.
- MTU/Jumbo frames **MUST** be enabled on all VLAN's and routed links

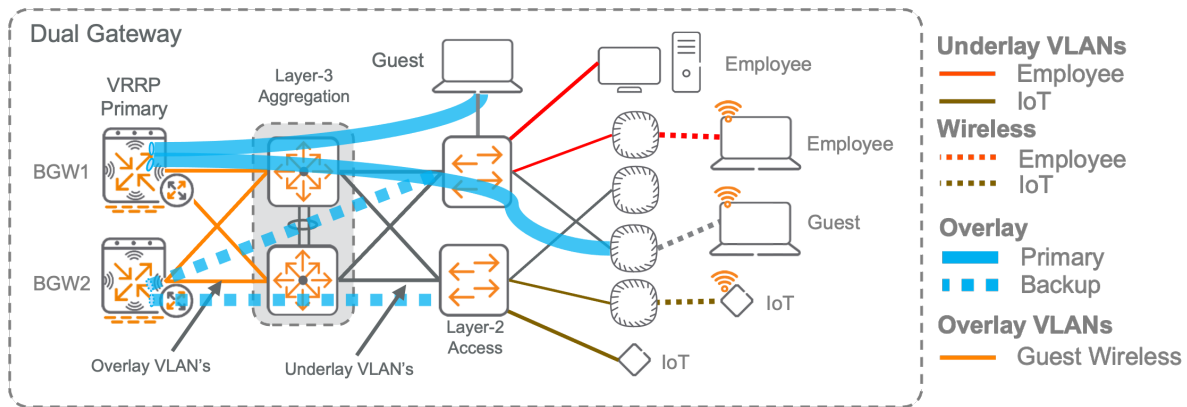


Figure 46: Tunneled Access with Dynamic Segmentation

NOTE:

User-based tunneling can be enabled in a layer 2 or layer 3 LAN deployment. This example uses only layer 3 architecture.

Centralized Multi-Site Fabric

The design below illustrates two branch deployments, both branches implemented user roles. Both branches have enabled role propagation to enforce policy across the WAN. Role propagation is done by mapping the user role to VXLAN Group-Based Policy (GBP) tag then encapsulating the tag in IPSEC. The fabric only uses the VXLAN GBP header and does not transport layer 2 payload.

For successful role propagation, two factors must be considered: the implementation of user roles and the gateway's awareness of user roles. There are several methods to apply user roles, including static user role to VLAN mapping or 802.1x. Additionally, it is essential for the gateway to accurately recognize user roles. To achieve consistency across all locations utilize the Global Policy Manager.

The method of applying user roles can determine whether Microsegmentation or VLAN segmentation is enabled across various sites. For instance, a static mapping of VLANs to user roles would appear as follows: The user role "MGMT" is mapped to VLAN 100. Consequently, whenever the gateway receives traffic on VLAN 100, it associates it with the "MGMT" user role. By creating a WAN policy, communication between MGMT user roles across the WAN can be permitted, instead of specifying particular subnets. Similarly, using 802.1x requires the gateway to have knowledge of the VLANs associated with each user role.

Another example is the utilization of Microsegmentation with User-Based Tunneling (UBT). In this scenario, each user authenticates using 802.1x and receives a distinct user role. For instance, User-Role-A and User-Role-B. Instead of relying on VLAN mapping, the gateway obtains awareness of each user role through the tunnel. This concept extends to tunneled SSIDs as well. As each user role possesses a distinct name, each user role can now be assigned a different Group-Based Policy (GBP) tags while still being part of the same VLAN. This enables Microsegmentation at the local branch, with the user role enforcing policies. Once the traffic traverses the WAN, the GBP tag ensures policy enforcement at the destination site.

Design Considerations

- Gateway Must be aware of the user roles of devices.
- Intra-branch policy is enforced at the BGW
- Branch to Branch policy is enforced at the destination branch.
- Use Global policy manager and NetConductor to define roles, this is also where the GBP-ID is set.
- VXLAN-GBP packets are fragmented and reassembled at each gateway. (No MTU changes required)
- Policy is enforced within enabled groups and across enabled groups.

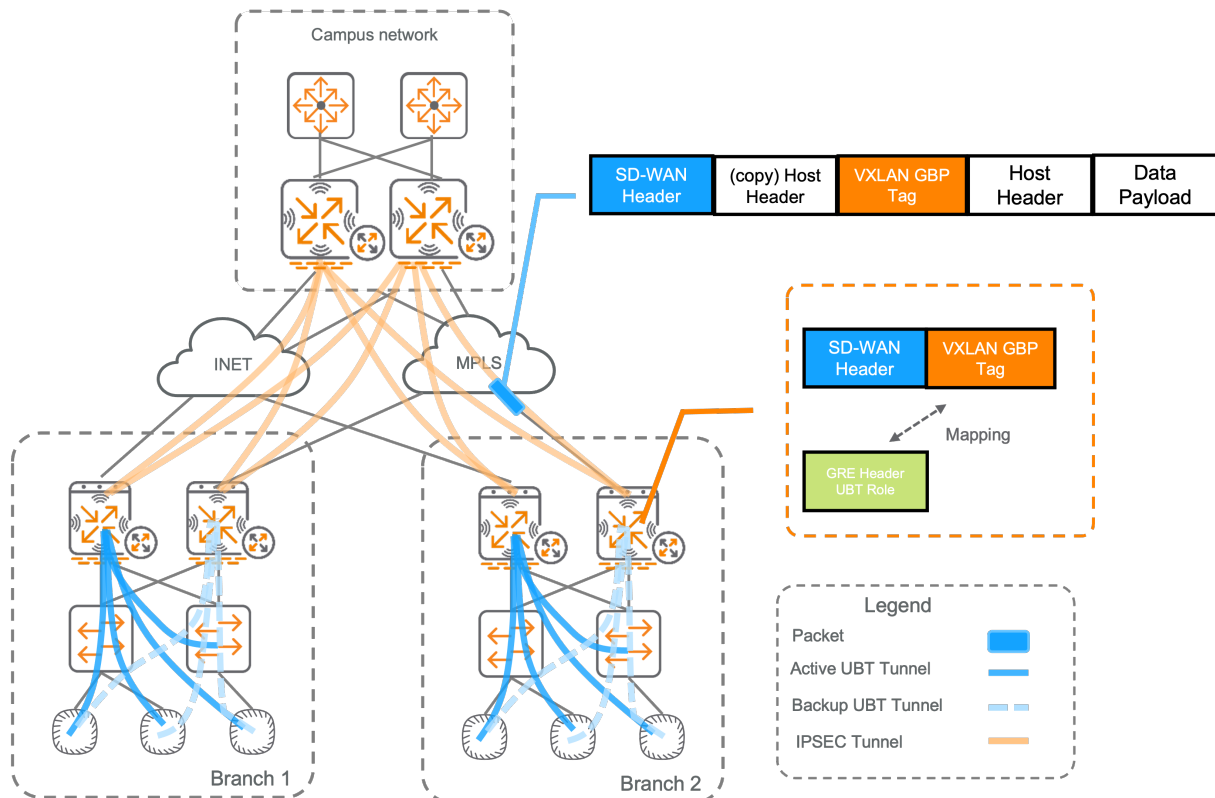


Figure 47: Centralized_Multi_fabric

NOTE:

This example shows role propagation with user based tunneling, however UBT is not required to allow for role propagation.

Microbranch

Microbranch enables an access point to act as a scaled-down branch gateway, providing remote connectivity via IPsec to a hub site. Microbranch is great for remote teleworks and small branches, such as a small suite. Microbranch shares much of the fundamental functionality as the branch gateway, including backup tunnels to other cluster members and data center preference. However, Microbranch has a maximum of five active hub sites.

Microbranch also supports features such as cloud security integration, Policy Based Routing, and WAN health checks. Microbranch does not support Dynamic Path Steering.

Microbranch supports three different tunneling modes which are linked to an SSID: Centralized Layer 2, Layer 3 NATed, and Layer 3 Routed.

Each tunneling type can be run in parallel on the same Microbranch AP. The table below shows capabilities for each mode.

Layer 3 Routed	Layer 3 NATed	Centralized Layer 2
Route orchestration Local Subnets/DHCP services Subnets shared with VPNC (Fully Routable in DC) Internet traffic broken out locallyLoad balancing based on route cost	Route orchestration Local Subnets/DHCP servicesSubnets not shared with VPNC Internet traffic broken out locallyLoad balancing based on route cost	No Route orchestration Subnets exist on VPNC Full tunneling to VPNC (Split tunneling can be done with PBR) Load balancing evenly distributed between VPNC's

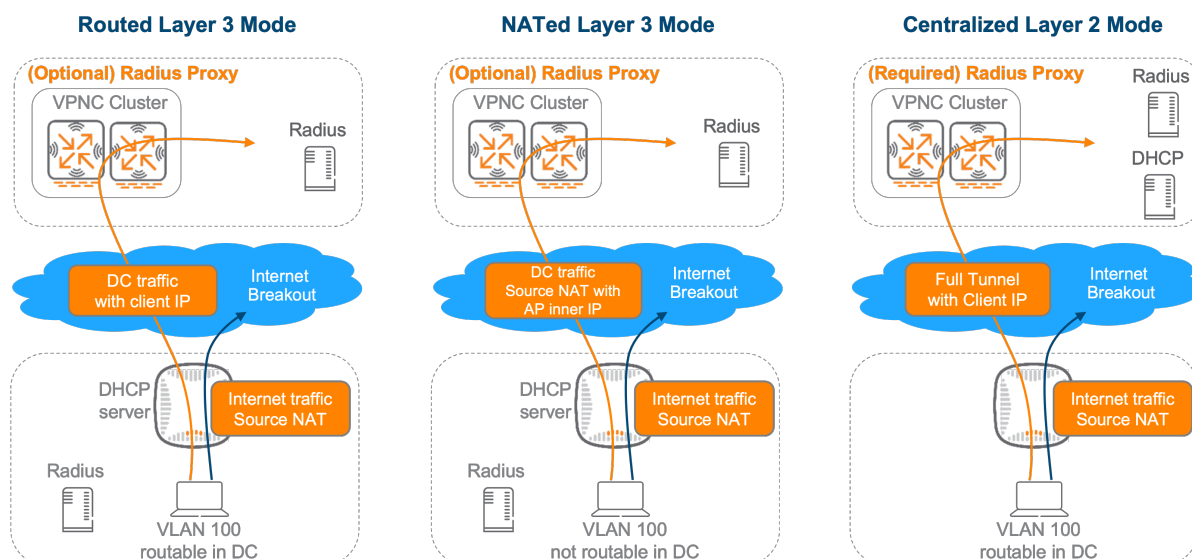


Figure 48: Microbranch_comparison

In the topology above, Radius proxy is required for layer 2 mode, since the VPNC is the default gateway for clients. In Centralized Layer 2 mode, the VPNC requires a VRRP Virtual IP address. The Virtual IP address is used as the Radius proxy IP address. If an organization is using layer 3 mode and needs to enable Radius proxy, the VPNC requires a layer 2 connection to establish a VRRP session. It is recommended to use an LACP connection for redundancy.

When deciding the tunneling method, consider:

Layer 3 Routed - The AP makes routing decisions based on its route table. There is a max of 512 routes, so it is important to use branch and hub route summarization. This method is optimal because there is less configuration overhead for the admin and it enables the AP to make optimal routing decisions (without requiring hairpinning through the VPNC). This is the preferred method to deploy Microbranch APs.

Layer 3 NATed - APs use the routing table to make routing decisions; however, traffic initiated by a client in a NATed SSID cannot reach resources located at the hub site. NATed SSIDs do not share their subnet with the SD-WAN fabric. This is a good method to use for guest SSIDs.

Centralized Layer 2 (CL2) - APs tunnel all traffic to the VPNC by default. The VLAN for each branch site resides on the VPNC. APs can break out locally to the Internet based on PBR policy; however, if local breakout is used, it is generally recommended to send all internal IP address space to the VPNC and send all other routing address space to the Internet. This method is generally recommended when devices require layer 2 adjacency to work.

Microbranch APs can use both layer 2 and layer 3 modes in parallel. The VPNCs also can support both modes in parallel; however, if layer 2 is in use, the VPNCs must be L2 adjacent to enable VRRP as described in the previous section.

Layer 2 tunnel orchestration differs from that of layer 3. Layer 2 tunneling builds tunnels to all members within a cluster and load-balances based on the client load. For example, if a hub site has two gateways, the Microbranch AP builds a tunnel to both VPNCs. If client traffic must be forwarded, the AP sends it over a tunnel to VPNC1 and the next client is sent to VPNC2.

Layer 3 tunneling also builds a tunnel to every VPNC; however, traffic is forwarded based on route cost which is dynamically adjusted based on the Overlay Route Orchestrator. This works in the same manner as a branch gateway.




Figure 49: clustering and load balancing

Customer Profile

This section provides an The requirements WAN and Branch requirements of Orange Widge Logistics.

In the two following sections, reference architectures are provided to illustrate:

- A solution design to meet the hypothetical customer's requirements using EdgeConnect SD-WAN.
- A solution design to meet the hypothetical customer's requirements using EdgeConnect SD-Branch.

Customer Overview

The customer is a US-based business operating hubs in Seattle and New York.

The business has 100 branch locations of varying size and importance spread throughout the US.

The customer expects an average of 10% year-over-year growth of branch locations.

The solution is scaled out for at least five years of growth.

Requirements

The customer seeks to:

- Decrease reliance on MPLS to reduce operational expense, with the goal to phase it out completely over time.
- Improve the experience for users using IaaS (Infrastructure-as-a-Service) and SaaS (Software-as-a-Service) applications as the business migrates to the cloud.
- Protect certain sensitive corporate data going to a SaaS service by traversing IPS/DLP.
- Provide critical guest Wi-Fi service at branch locations.
- Gather guest Wi-Fi usage data to assess customer behavior and apply a “digital-first” approach to improve the customer experience and attract customers to stores.
- Use a cloud-first approach for all IT infrastructure, minimizing the on-premise footprint as much as possible.
- Begin using IaaS providers such as AWS and Azure and employ SD-WAN optimization for the IaaS locations.

Site Profiles

As the customer gains comfort with the SD-WAN solution, the MPLS connections will be phased out in favor of additional Internet circuits.

Traffic between the traditional hubs and spokes will continue to decrease over time as more workloads are moved to IaaS environments and SaaS solutions.

- Seattle and New York both have existing 200 mbps MPLS connections. They will use two 1000/1000 mbps business class internet circuits.
- Although the virtual hubs in Seattle and New York are not currently in place, the customer plans for deployments in the cloud provider for the US West region and the US East region.
- 20 large sites, defined as:
 - The business has no tolerance for unscheduled downtime
 - Uptime is provided by the gateway HA and cellular backup.
 - Up to 200 users
 - Use existing 40 Mbps connection and add 200/50 mbps commodity Internet circuits and a 5G LTE backup
- 75 medium sites, defined as:
 - The business has less tolerance for downtime
 - More uptime is provided by the gateway HA, but with no cellular backup.
 - Up to 100 users
 - Use existing 30 Mbps MPLS connection and add 100/10 commodity internet circuit
- 5 small sites, defined as:
 - The business can tolerate downtime
 - Up to 10 Users
 - Require only a single gateway, with no device-level HA or cellular backup
 - Use existing 5mbps MPLS connections and the customer plans to add 50/10 mbps commodity Internet circuits.

Traffic Profile

The company's traffic predominantly includes:

- Use of Zoom and Microsoft Teams for real-time communications
- Real-time inventory data queries on in-house SQL systems, hosted at data centers
- Bulk FTP file transfers used throughout the environment to process transactions hosted at data centers
- SaaS applications, such as Sales Force, are used to provide optimal Internet egress.

EdgeConnect SD-WAN Reference Design

This section of the guide demonstrates how to design an EdgeConnect SD-WAN deployment based on the **customer profile**.

Component Selection

This section describes the components of a reference architecture using HPE Aruba's EdgeConnect SD-WAN solution.

NOTE:

This reference architecture is based on the 9.4 release of Aruba Orchestrator and EdgeConnect SD-WAN gateways.

EdgeConnect Gateways

Follow the steps below to determine which gateway is appropriate at each location:

- Determine bandwidth requirements.
 - If asymmetric circuits are in place, use the higher number for bandwidth calculations.
- Determine fiber port, power supply, and HA requirements.
- Identify if advanced features are required, such as Dynamic Threat Defense (DTD) or Boost.
- Select the appliance that supports the requirements:

	EC-						
Model	10104	EC-10106	EC-XS	EC-S-P	EC-M-H	EC-L-H	EC-XL-H
Typical Deploy-ment	Small Branch / Home Office	Small Branch	Small Branch	Large Branch	Head Office/DC Large Hub	Data Center Large Hub	Data Center Large Hub
Typical WAN Band-width	2-500 Mbps	2-1000 Mbps	2-1000 Mbps	10-3000 Mbps	50-5000 Mbps	2-10 Gbps	2-10 Gbps
Simultaneous Connections	256,000	256,000	256,000	256,000	2,000,000	2,000,000	2,000,000

Model	EC-10104	EC-10106	EC-XS	EC-S-P	EC-M-H	EC-L-H	EC-XL-H
Recommended WAN Boost up to:	100 mbps	250 mbps	250 mbps	500 mbps	1 gbps	1 gbps	5 gbps
Redundant / FRUs	No	No	No	SSD and Power (AC or DC)	SSD and Power	SSD and Power	SSD, NVMe, Power
Data Path Interfaces	4 x RJ45 10/100/1000	2 x RJ45 2 x Combo 1/10G SFP+	4 x RJ45 10/100/1000	8 x RJ45 4 x 1/10G Optical	8 x RJ45 4 x 1/10G Optical	6 x 1/10G Optical	6 x 1/10/25G Optical

NOTE:

WAN bandwidth assumes bidirectional traffic (symmetric uplink and downlink). For total WAN throughput (Rx+Tx), multiple these numbers times 2.

Transceivers

EdgeConnect SD-WAN gateways have been certified to operate with the following transceivers. Use the table to verify that the selected transceiver(s) supports physical connectivity.

Transceiver SKU	Description
EC-SFP-SR	SFP+, 1/10G, SR
EC-SFP-LR	SFP+, 1/10G, LR
EC-SFP28-25G-SR	SFR28 Transceiver, 10/25G, SR
EC-SFP28-25G-LR	SFR28 Transceiver, 10/25G, LR

Licensing

Follow these steps to license an EdgeConnect SD-WAN deployment. Licensing is purchased on a per-gateway basis.

Step 1 Select a Term based subscription license in intervals of 1, 3, 5, or 7 years. Choose **Foundation**, **Advanced**, or **On-Premise**. Choose the same tier for all gateways in the environment. More information on tiers is provided below.

Step 2 Choose the EdgeConnect gateway (physical or virtual) software bandwidth. The license is based on the aggregate WAN-side bandwidth of the given node and has options up to 8 bandwidth tiers, depending on the subscription tier. More information on these tiers is provided below.

Step 3 (Optional) Add Boost. licensed in units of 100 Mbps and 10Gbps of WAN optimization, can be deployed flexibly to sites that require application acceleration.

Step 4 (Optional) With the Add Dynamic Threat Defense (DTD) license, the EdgeConnect Platform supports IDS/IPS. In the future, additional security-related features may be tied to the DTD License.

Tiered Subscription Licensing Model

HPE Aruba provides three licensing tiers for EdgeConnect SD-WAN deployment: **Foundation**, **Advanced**, and **On-Premise**.

- **Foundation** provides a value-centric option including essential SD-WAN features and advanced NGFW features. It includes Cloud Orchestrator.
- **Advanced** provides maximum performance with advanced SD-WAN features and advanced NGFW features. It includes Cloud Orchestrator.
- **On-Premise** provides maximum performance for on-premise deployment with advanced SD-WAN features and advanced NGFW features.

Advanced is the recommended tier for most customers because it supports all the SD-WAN features and provides the most flexibility,

The table below outlines the differences in the tiers

	Foundation (AAS)	Advanced (AAS)	On-Prem
Bandwidth Tiers	3 tiers (100M, 1G, 10G)	20M/50M/100M/200M/500M	20M/50M/100M/200M/500M/1G/2G/Unlimited
BIOs	3 (RealTime, Critical & Default)	7	7
Network Segments/VRF	2 (auto-enabled, default & guest only)	64	64
Routing	BGP, OSPF, Subnet Sharing	BGP, OSPF, SS	BGP, OSPF, SS
Mesh Networking	No	Yes	Yes
Multi-region Topology	Max 4 regions, 4 hubs/region	Yes	Yes
AppExpress	No (Monitor only)	Yes (Monitor and Steer)	Yes (Monitor and Steer)
Orchestrator	Cloud Orchestrator Foundation	Cloud Orchestrator Advanced	On-premises

	Foundation (AAS)	Advanced (AAS)	On-Prem
Orchestrator	24h/7d/1mo (m/h/d)	72h/14d/3mo	Custom
Stats Retention			

Common features in all licensing tiers include:

- Unconstrained site count
- Path Conditioning
- NGFW
- Firewall protection profiles
- AVC/1st packet iQ
- Zero Trust Segmentation (Roles)
- DDOS detection/mitigation
- EC SD-WAN Fabric Orchestration
- Advanced Crypto
- EdgeHA
- ZTP / Templates.

Bandwidth License

Each site must be licensed for the appropriate bandwidth.

To calculate the throughput licenses required for a given gateway, determine the total WAN bandwidth the gateway will use. For example, for a 100 mbps Internet circuit and a 20 mbps MPLS circuit, the total is 120 mbps. For asymmetric circuits, such as a 50 mbps download and 5 mbps upload Internet circuit, use the largest number for the calculation. Since cellular back-up is used as path-of-last-resort only, it is not used for throughput license calculation.

Boost Licensing

Boost, though not currently used in this design, is sold in blocks of 100 mbps and allocated to the gateways in 1kbps increments. To determine the amount of boost licensing required, follow these steps:

- Determine the applications to accelerate.
- Determine the number of sites needed to accelerate the applications.
- Determine how much bandwidth the specific applications require to work optimally at each site.
- Purchase the total amount of boost to apply at the sites.

Dynamic Threat Defense

The advanced security license provides IDS capabilities in EdgeConnect. The example design does not include this feature. To determine the advanced security licensing required, use the guidance below.

Example Hub Design

The illustration below shows an example hub.

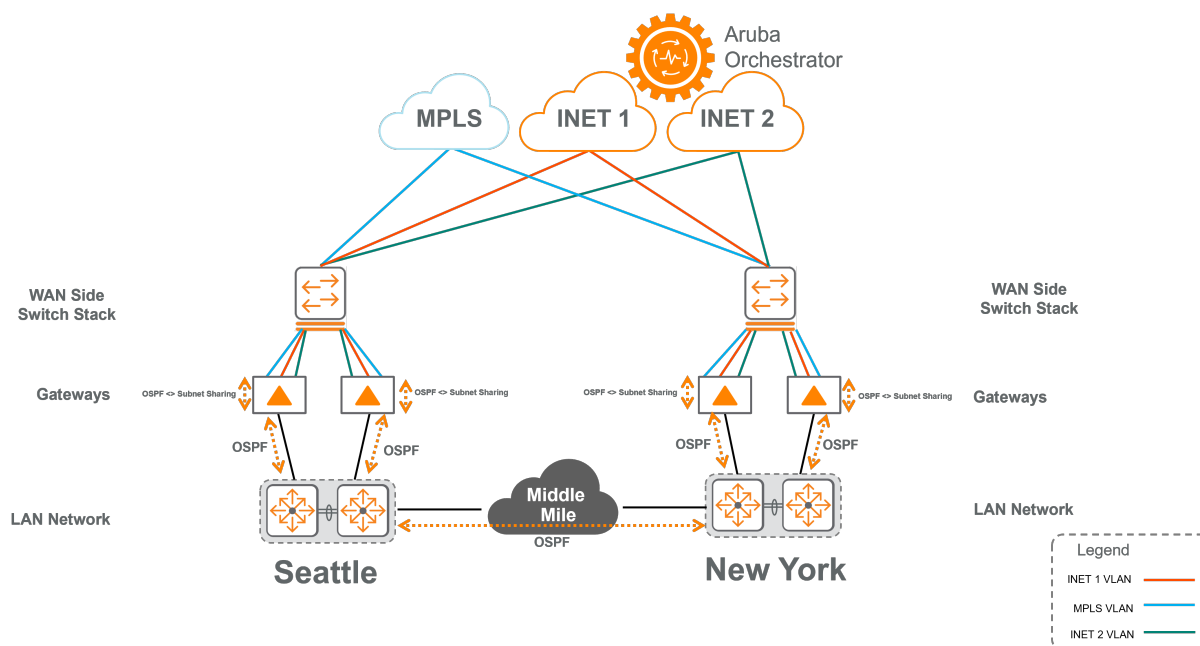


Figure 50: RA - Hub

Key elements of the design include:

- Gateways are placed inline.
- WAN connections plug into the upstream WAN side switch and are then connected to gateways to allow for traditional high availability. This requires additional IP space on the WAN transports.
- Gateways connect via L3 to the LAN, into a WAN aggregation block and peer OSPF.
- BGP adjacency is maintained with the MPLS provider to facilitate migration.
- A hub summary route and default route are advertised with Subnet Sharing.
- All branches use a template with peer priority to select a hub when the same route is presented from multiple hub gateways.
- High OSPF can be set when redistributing from Subnet Sharing (WAN) to OSPF (LAN).
- The design ensures that traffic between the data centers uses the Data Center Interconnect (DCI). The OSPF can be adjusted as needed.
- When redistributing from Subnet Sharing (WAN) to OSPF (LAN), assign a higher value to one gateway to ensure flow symmetry.

The following components are selected to meet the hub requirements. Hubs are purchased in five-year terms.

Hardware

Quantity	SKU	Consideration
4	EC-L-H (Recommended)	Two per hub. Provides high levels of hardware HA
	EC-XL-H (Alternative)	Use if high levels of Boost are required or 25G connectivity to the LAN / WAN
	EC-M-H (Alternative)	Use if less bandwidth is needed
1		
12	EC-SFP-SR (Recommended)	Customer currently has SR fiber plant
	EC-SFP-LR (Alternate)	

Licenses

If planning to replace MPLS circuits with higher-speed internet circuits in the future, consider buying licensing based on future needs or purchasing licenses with shorter terms.

Quantity	SKU	Consideration
4	HPE ANW EC Adv UL 5yr Sub SaaS	Customer bandwidth above 2 Gbps at the hubs
	HPE ANW EC Adv 2Gb 5yr Sub SaaS (Alternate)	Can be used when customer has lower bandwidth needs
	HPE ANW EC Adv 1Gb 5yr Sub SaaS (Alternate)	Can be used when customer has lower bandwidth needs
	HPE ANW EC Adv 500mb 5yr Sub SaaS (Alternate)	Can be used when customer has lower bandwidth needs

Example Branch Design

Small Site Design

The illustration below shows an example hub for a small site.

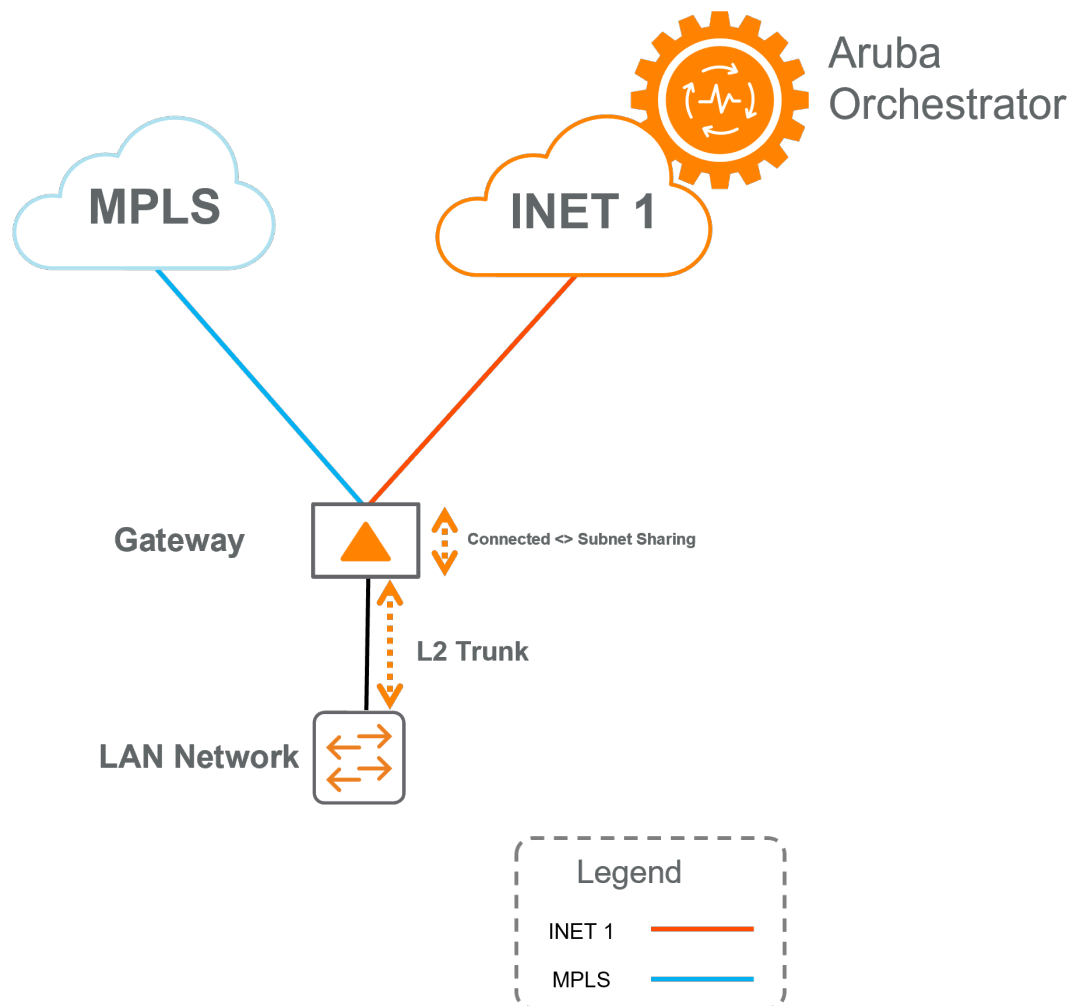


Figure 51: Small Branch

Key elements include:

- The gateway is placed inline.
- The gateway is connected to the switch via a single L2 connection.
- The gateway acts as default gateway for any subnet at the branch.

The following components were selected to meet small site requirements, purchased in five-year terms.

Hardware

Quantity	SKU	Consideration
1	EC-10104 (Recommended)	Cost efficient for low bandwidth / user count
	EC-10106 (Recommended)	Use for higher bandwidth applications or SFP+ is required

Licenses

If planning to replace MPLS circuits with higher speed internet circuits in the future, consider buying licensing based on future needs or purchasing licenses with shorter terms.

Quantity	SKU	Consideration
1	HPE ANW EC Adv 500mb 5yr Sub SaaS (Recommended)	Based on customers current needs, with room for growth
	HPE ANW EC Adv 200mb 5yr Sub SaaS (Alternate)	Can be used when customer has different bandwidth needs
	HPE ANW EC Adv 100mb 5yr Sub SaaS (Alternate)	Can be used when customer has different bandwidth needs
	HPE ANW EC Adv 50mb 5yr Sub SaaS(Alternate)	Can be used when customer has different bandwidth needs

Medium Site Design

The illustration below shows an example hub for a medium site.

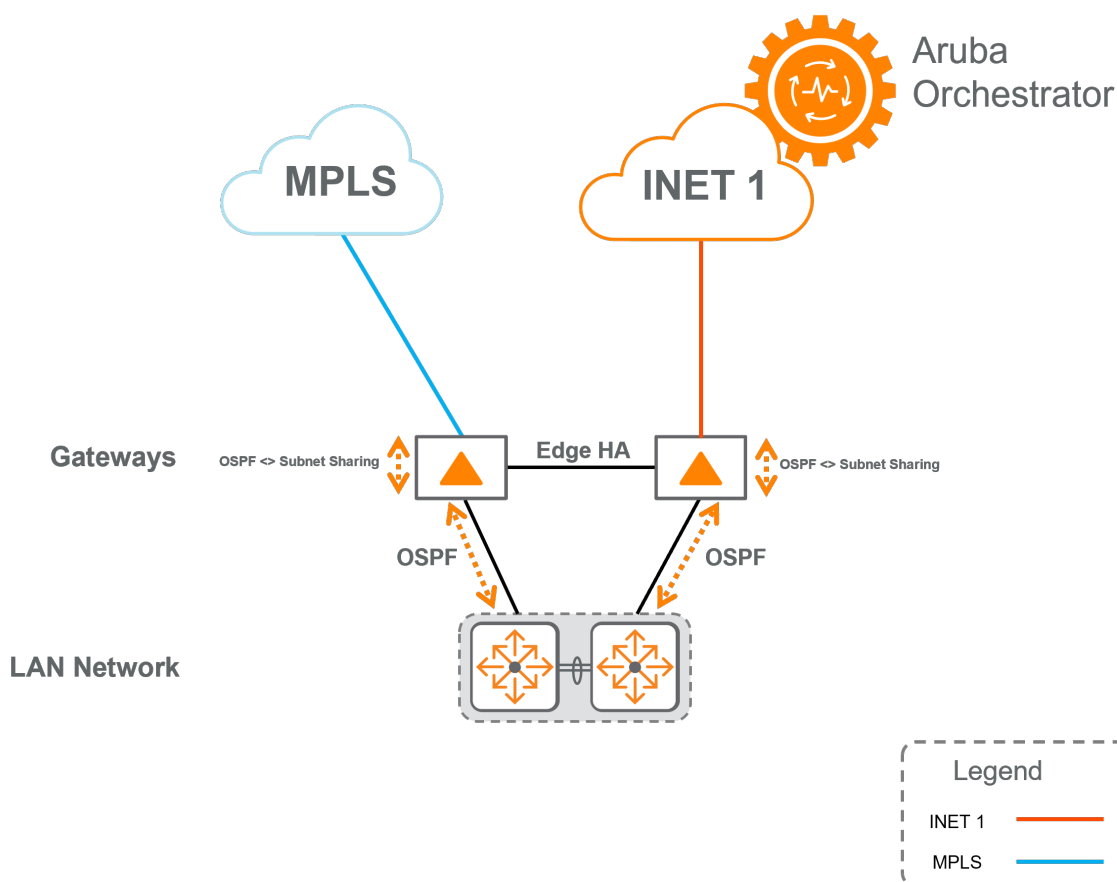


Figure 52: RA - Medium Branch

Key elements include:

- The gateway is placed inline.
- EdgeHA is used for high availability.
- The gateway is connected to a collapsed core using an L3 link and peer OSPF.
- When redistributing from Subnet Sharing (WAN) to OSPF (LAN), assign one gateway a better metric to ensure flow symmetry.
- When redistributing from OSPF to Subnet Sharing, assign one gateway a better metric to ensure flow symmetry.

The following components are selected to meet medium site requirements, purchased in five-year terms.

Hardware

Quantity	SKU	Consideration
2	EC-10106 (Recommended)	
	EC-S-P (Alternative)	Use for higher bandwidths and more hardware HA

Licenses

If planning to replace MPLS circuits with higher speed internet circuits in the future, consider buying licensing based on future needs or purchasing licenses with shorter terms.

Quantity	SKU	Consideration
1	HPE ANW EC Adv 1gbps 5yr Sub SaaS (Recommended)	Based on customers current needs, with room for growth
	HPE ANW EC Adv 500mb 5yr Sub SaaS (Recommended)	Can be used when customer has different bandwidth needs
	HPE ANW EC Adv 200mb 5yr Sub SaaS (Alternate)	Can be used when customer has different bandwidth needs
	HPE ANW EC Adv 100mb 5yr Sub SaaS (Alternate)	Can be used when customer has different bandwidth needs
1	HPE ANW EC Adv HA	HA license for the second gateway at the site. Be sure to match tier, bandwidth, and term.
1	HPE ANW EC DTD 5yr Sub SaaS	Optional DTD license
1	HPE ANW EC DTD HA 5yr Sub SaaS	Optional DTD HA license

Large Site Design

The illustration below shows an example hub for a large site.

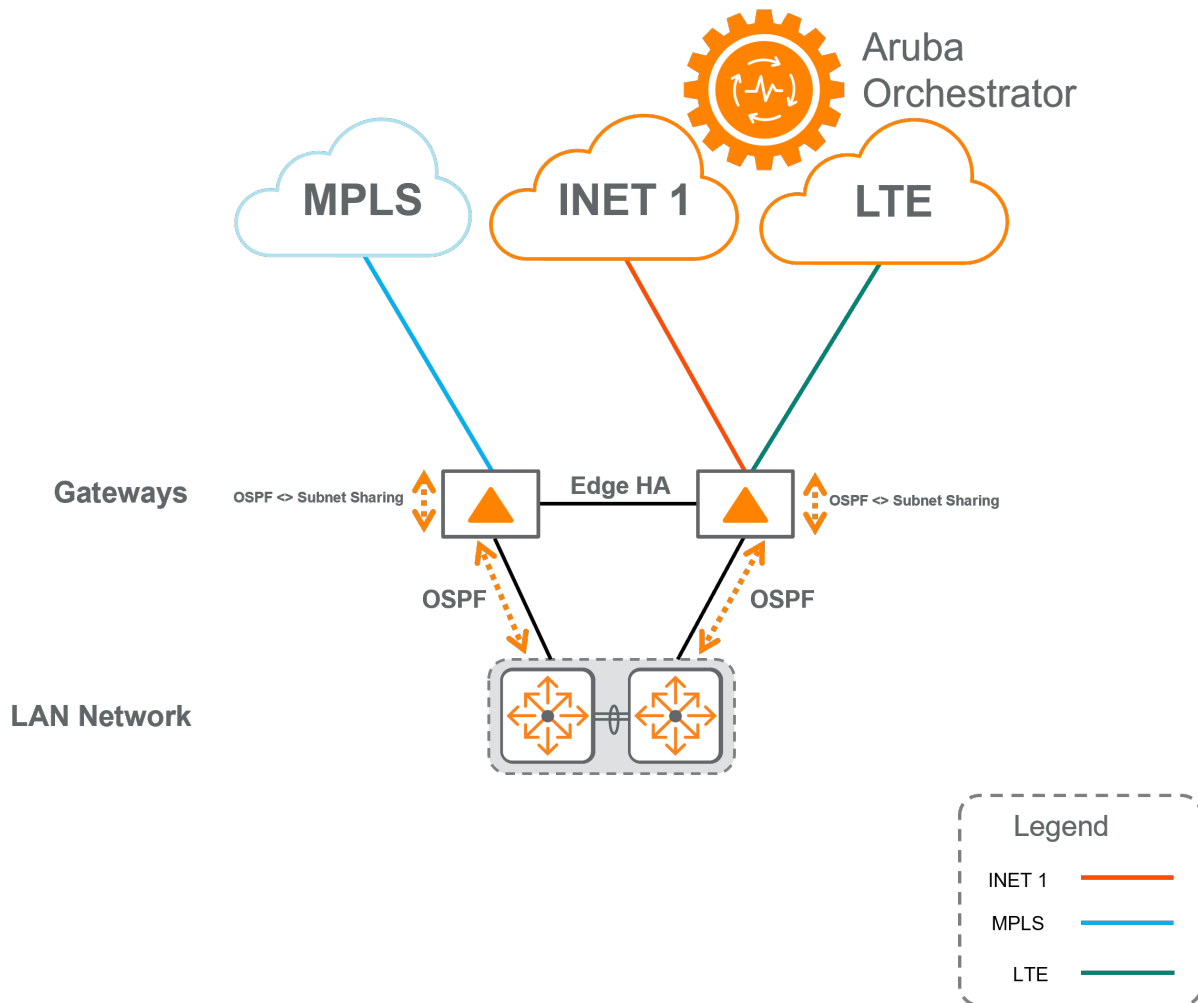


Figure 53: RA - Large Branch

Key features include:

- The gateway is placed inline.
- EdgeHA is used for high availability.
- The gateway is connected to a collapsed core via L3 link and peer OSPF.
- When redistributing from Subnet Sharing (WAN) to OSPF (LAN), assign one gateway a better metric to ensure flow symmetry.
- When redistributing from OSPF to Subnet Sharing, assign one gateway a better metric to ensure flow symmetry.
- LTE connection, which requires a third-party device to terminate the LTE with an ethernet handoff, is used as a path-of-last resort, forwarding traffic only if both circuits are down.

The following components are selected to meet large site requirements, purchased in five-year terms. The large site has the highest bandwidth, serving a large amount of users, with HA implemented.

Hardware

Quantity	SKU	Consideration
2	EC-S-P (Recommended)	Cost efficient for low bandwidth / user count
	EC-M-P (Alternative)	Use for higher bandwidths





Licenses

If planning to replace MPLS circuits with higher speed internet circuits in the future, consider buying licensing based on future needs or purchasing licenses with shorter terms.

Quantity	SKU	Consideration
1	HPE ANW EC Adv 2gbps 5yr Sub SaaS (Recommended)	Based on current needs, with room for growth
	HPE ANW EC Adv 1gbps 5yr Sub SaaS (Recommended)	Based on customers current needs, with room for growth
	HPE ANW EC Adv 500mb 5yr Sub SaaS (Alternative)	Can be used when customer has different bandwidth needs
1	HPE ANW EC Adv HA	HA License for second Gateway at site, ensure to match tier, bandwidth, and term
1	HPE ANW EC DTD 5yr Sub SaaS	Optional DTD license
1	HPE ANW EC DTD HA 5yr Sub SaaS	Optional DTD HA license

Example Overlay Design

The illustration below shows an example overlay.

Priority	Overlay	SD-WAN Traffic to Internal Subnets ↗					Breakout Traffic to Internet & Cloud Services ↗		
		Topology	Hubs +	Primary Interfaces	Backup Interfaces	QoS & Security +	Policy Order	Primary Interfaces	Backup Interfaces
1 = ×	RealTime Match Traffic Overlay ACL	Mesh 		INET2 INET1 MPLS1 MPLS2 High Availability Waterfall: Overall Quality	LTE If Pri & Sec Down		1 Break out 2 Backhaul	INET2 INET1 Waterfall: Auto	LTE
2 = ×	CriticalApps Match Traffic Overlay ACL	Hub & Spoke 		INET2 INET1 MPLS1 MPLS2 High Quality Waterfall: Overall Quality	LTE If Pri & Sec Down		1 Netskope 2 Backhaul	Waterfall: Auto	
3 = ×	BulkApps Match Traffic Overlay ACL	Hub & Spoke 		INET2 INET1 MPLS1 MPLS2 High Quality Waterfall: Overall Quality	LTE If Pri & Sec Down		1 Break out 2 Backhaul	INET2 INET1 Waterfall: Auto	LTE
4 = ×	DefaultOverlay Match Traffic Overlay ACL	Hub & Spoke 		INET2 INET1 MPLS1 MPLS2 High Quality Waterfall: Overall Quality	LTE If Pri & Sec Down		1 Break out 2 Backhaul	INET2 INET1 Waterfall: Auto	LTE

Features include:

- Cloud firewall integration is used to tunnel specific to-cloud filtering.
- EC-V is deployed in IaaS providers to facilitate better access to the services hosted there.
- Zone-Based Firewalling is configured on the gateways to provide basic filtering of Internet traffic when cloud firewalling is not used.

The table below illustrates the BIO design elements. All traffic matches start with the default overlay ACL and are modified as necessary.

BIO				
Name	Traffic Match	Topology	Link Bonding	Internet Egress
Real Time	Real-time communication applications	Mesh topology provides optimal traffic flows between branches	The High Availability link bonding policy should provide 1:1 FEC to ensure no loss of voice / video traffic	Direct-to-net is primary, with hub fallback
Critical Apps	Business-critical internal apps, such as the inventory system, and SaaS apps such as Salesforce	Hub and spoke	High Quality	Cloud firewall is primary, with a fallback to the hubs which host a backup security stack

BIO

Name	Traffic Match	Topology	Link Bonding	Internet Egress
BulkA	Large internal traffic flows, such as FTP and cloud-hosted file repositories	Hub and spoke	High Quality. Within the QOS policy ,this BIO is allotted only a certain percent of WAN bandwidth during times of congestion, ensuring that large flows do not saturate the WAN transports.	Direct-to-net is primary, with hub fallback, as these flows are large and must go direct to their destinations
Default	Match all other traffic, including guest	Hub and spoke	High Quality	Direct-to-net with a fallback of the data center, since guest traffic is critical to the business

NOTE:

If the Foundation licenses are used, only three BIOs are possible: real-time, critical, and default.

EdgeConnect SD-Branch Reference Design

This section of the guide demonstrates how to design an EdgeConnect SD-Branch deployment based on the [customer profile](#).

Component Selection

This section describes how to select the correct devices to use for any deployment. Use this guidance to select devices appropriate for the customer profiles presented below.

Selecting a Hub Device — There are four aspects to consider: maximum number of tunnels, the number of routes to be learned from the data center, interface count, and interface type.

Selecting a Branch Device — There are three aspects to consider when deciding the branch gateway to use: firewall sessions, interface count, and interface type.

See the [DataSheet](#) for SD-WAN gateway devices and scale numbers.

Licensing Options

Foundation - This license provides all features required for SD-Branch functionality in branch or headend deployments.

Foundation Base - This license provides all features included in a Foundation License, but can support only up to 75 client devices per branch site.

Foundation with Security - This license provides all features required for [SD-WAN](#) functionality in branch or headend deployments with additional security features.

Foundation Base with Security - This license provides all the features included in a Foundation with Security License, but can support only up to 75 client devices per branch.

Advanced - This license provides all the features included in a Foundation License, with additional features related to SaaS Express Net Conductor and AI Insights.

Advanced with Security - This license provides all the features of an Advanced License, with additional security features related to IPS and IDS, security dashboard, and anti-malware.*

Virtual Gateway (VGW) License - This license is available for AWS, Azure, and ESXi platforms and is licensed based on the bandwidth required. The license types available for VGW are VGW-500M, VGW-2G, and VGW-4G.

See the [Ordering Guide](#) for more detail.

Overlay Design

For each of the customer profiles below, the following general requirements and considerations apply:

- Improve the experience for users with IaaS (Infrastructure-as-a-Service) and SaaS (Software-as-a-Service) applications as the business migrates to the cloud.
- Protect certain sensitive corporate data going to a SaaS service by traversing IPS/DLP.
- Use Zoom and Microsoft Teams for real-time communications.
- Perform real-time inventory data queries on in-house SQL systems, hosted at data centers.
- Use bulk FTP file transfers to process transactions hosted at data centers throughout the environment.
- Use SaaS applications, such as Sales Force, to provide optimal Internet egress.

To address these requirements:

- Hub-and-spoke overlay is used.
- To improve SaaS performance, SaaS express is used to break out Sales Force, Zoom and Teams traffic locally.
- A high-priority DPS policy is used for inventory queries. Other applications use separate DPS policy.
- To ensure application security, stateful application-aware firewall is enabled along with IPS, Web Content filtering, and IP Reputation.

Hub Design

The customer profile includes the following requirements and considerations:

- Accommodate 100 branch sites with an expected 10% growth over five years.
- Improve the experience for users with IaaS (Infrastructure-as-a-Service) and SaaS (Software-as-a-Service) applications as the business migrates to the cloud.
- Decrease reliance on MPLS to reduce operational expense, with the goal to phase it out completely over time.

Design Summary

Model Selection	Max IP Sec tunnels	Considerations
9012 (Recommended)	512	Redundant pair of gateways for 8+ year growth. In a failure scenario, one box can handle all sites.
9106 (Alternative)	8k	Redundant pair of gateways for 10+ year growth. In a failure scenario, one box can handle all sites.
vGW-2G (Future)	4096	Future consideration for IaaS/SaaS migration

The following list summarizes the hub design elements:

- Gateways are placed inline.
- Both WAN transports (INET, MPLS) are connected to each gateway.
- Gateways connect via L3 to the LAN, into a WAN aggregation block and peer OSPF.
- DC routes are summarized when redistributing into the SD-WAN overlay.

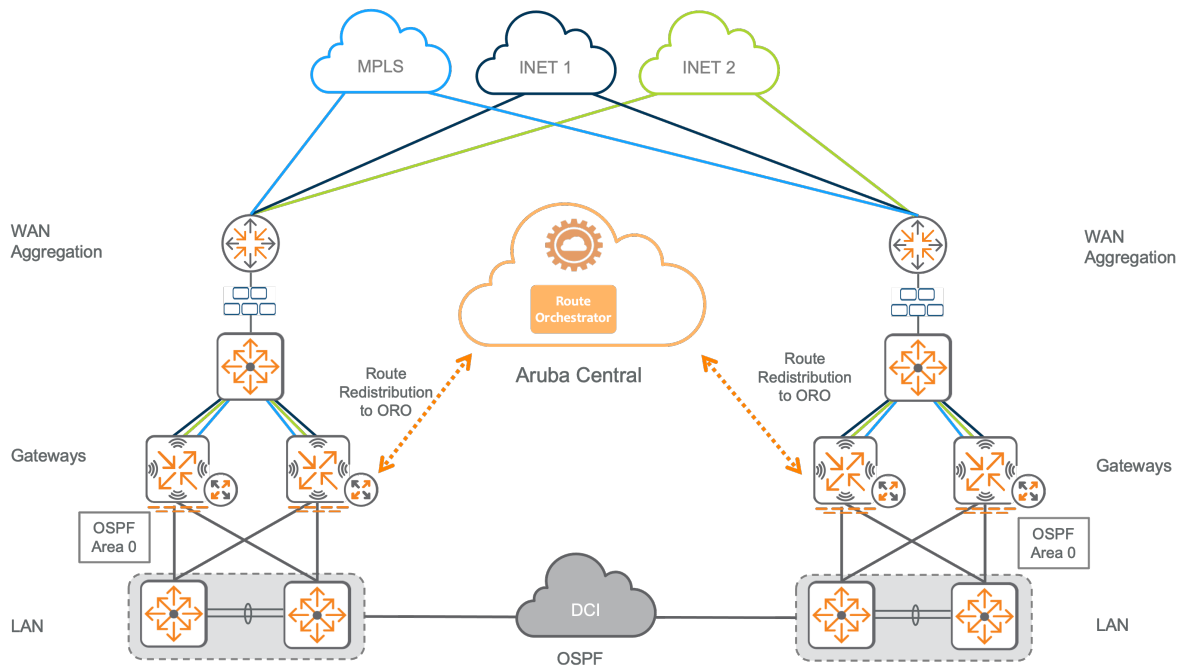


Figure 54: hub_desgin_ra

Branch Design

Based on the customer profile, there are three different branch site designs, requiring three different template groups for each site size. Medium and large sites are standardized at branch gateways; the small site is standardized on Microbranch.

Large Site

Based on the customer profile, large sites have the following requirements:

- The business has no tolerance for unscheduled downtime.
- Uptime is provided by the gateway HA and cellular backup.
- Certain sensitive corporate data going to a SaaS service by traversing IPS/DLP must be protected.
- The site has up to 200 users.
- The site uses an existing 40 mbps connection and plans to add 200/50 mbps commodity Internet circuits with a 5G LTE backup

To address these requirements:

- Dual gateways will be placed inline.
- MPLS will connect to one gateway, with INET connected to the second gateway.
- WAN Uplink sharing will be enabled.
- LTE connection will be used as a backup.
- Gateways connect via L3 to the LAN and peer OSPF to a collapsed core.
- Branch routes are summarized when redistributing into the SD-WAN overlay.
- Collapsed Core should be in a VSF stack.
- Collapsed Core to access switch connectivity should be LACP trunks.
- Tunneling is enabled for switching (UBT) and wireless.

Firewall		
Model Selection	Sessions	Considerations
9004 LTE (Recommended)	64k	Redundant pair of gateways, LTE built into gateways No SFP/SFP+ ports
6300 (Recommended)	—	Collapsed Core
6200 (Recommended)	—	Access switch
6100, 6300 (Alternative)	—	Access switch

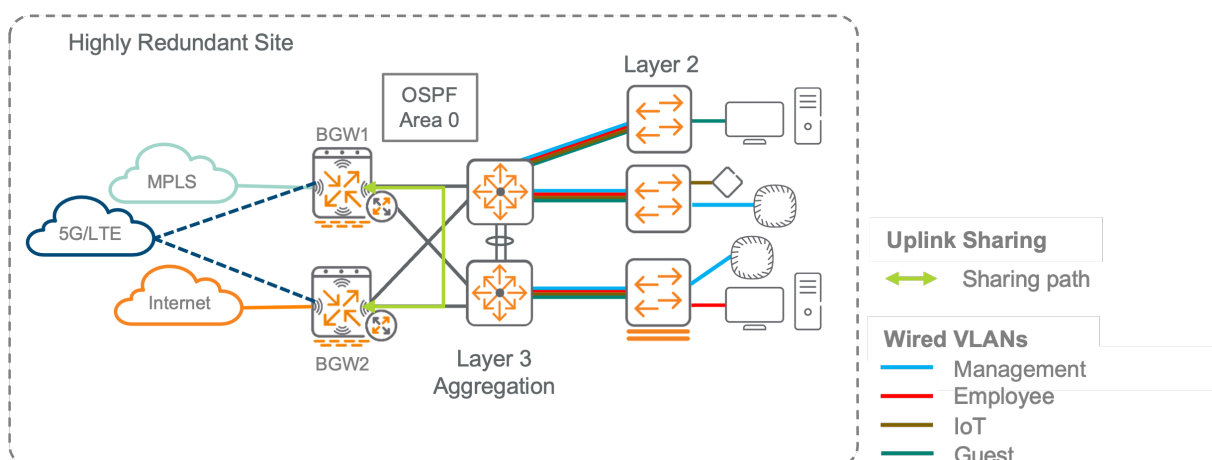


Figure 55: large_site_ra

Medium Site

Based on the customer profile, medium sites have the following requirements:

- The business has less tolerance for downtime.
- More uptime is provided by the gateway HA, but with no cellular backup.
- The site has up to 100 users.

- The site uses an existing 30 mbps MPLS connection and plans to add a 100/10 commodity internet circuit.

To address these requirements:

- Dual gateways will be placed inline.
- MPLS will be connected to one gateway, with INET connected to the second gateway.
- WAN uplink sharing will be enabled.
- Gateways connect via L3 to the LAN and peer OSPF to a collapsed core.
- Branch routes are summarized when redistributing into the SD-WAN overlay.
- Collapsed Core should be in a VSF stack.
- Collapsed Core to access switch connectivity should be LACP trunks.
- Tunneling will be enabled for switching (UBT) and wireless.

Model Selection	Firewall Sessions	Considerations
9004 (Recommended)	64K	Redundant pair of gatewaysNo SFP/SFP+ ports
6300 (Recommended)	—	Collapsed Core
6200 (Recommended)	—	Access switch
6100, 6300 (Alternative)	—	Access switch

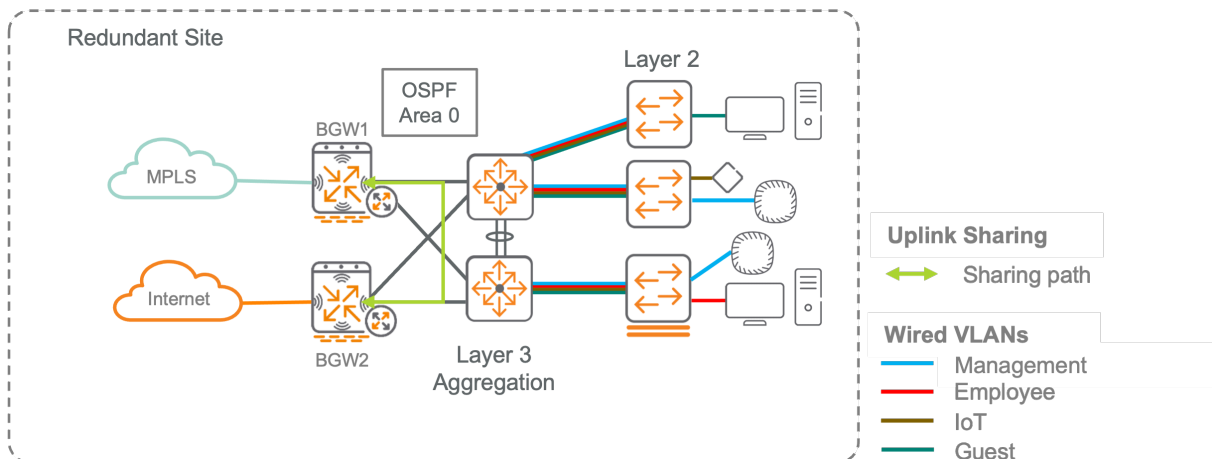


Figure 56: medium_site_ra

Small Site

Based on the customer profile, the small sites have the following requirements:

- The business can tolerate downtime.
- The site has up to 10 users.
- The site requires only a single gateway, with no device-level HA or cellular backup.

- The site uses existing 5 mbps MPLS connections and plans to add 50/10 mbps commodity Internet circuits.

To address these requirements:

- Single gateways will be placed inline.
- MPLS and INET will be connected to gateway.
- Gateway will act as Default gateway for all VLANs.
- The Guest network will use internet breakout.
- Branch routes will be summarized when redistributing into the SD-WAN overlay.

Model	Firewall Sessions	Considerations
9004 (Recommended)	64k	No SFP/SFP+ Ports
6100 (Recommended)	—	Extra ports for local devices.




Figure 57: small_site_1_ra

Alternative Small Site

To address these requirements:

- Single remote access point will be placed inline.
- INET will connect to AP.
- AP will act as that default gateway for all SSIDs.
- L3 will be routed for internal users.
- L3 will be NATed SSID for the Guest network.
- Branch routes will be redistributed into the SD-WAN overlay.

Model	Considerations
500H Series (Recommended)	Wi-Fi 6 ready
303H Series (Alternative)	Will not support next-generation Wi-Fi
6100 (Recommended)	Extra ports for local devices.

What's New in This Version

The following changes have been made since Aruba last published this guide:

- Addition of EdgeConnect SD-WAN design content.

© Copyright 2021 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein. Aruba Networks and the Aruba logo are registered trademarks of Aruba Networks, Inc. Third-party trademarks mentioned are the property of their respective owners. To view the end-user software agreement, go to: www.arubanetworks.com/assets/legal/EULA.pdf



www.arubanetworks.com

3333 Scott Blvd. Santa Clara, CA 95054
1.844.472.2782 | T: 1.408.227.4500 | FAX: 1.408.227.4550

See Confluence for Correct Doc Title