



**Implementing
Aruba
Campus
Access**

STUDENT GUIDE



Implementing Aruba Campus Access

23.11

Student Guide

January 2023

Implementing Aruba Campus Access

Copyright

© 2022 Aruba Networks, Inc. AirWave®, Aruba Networks®, Aruba Mobility Management System®, Bluescanner, For Wireless That Works®, Mobile Edge Architecture, People Move. Networks Must Follow., RFProtect, The All Wireless Workplace Is Now Open For Business, and The Mobile Edge Company® are trademarks of Aruba Networks, Inc. All rights reserved. All other trademarks are the property of their respective owners.

Open Source Code

Certain Aruba products include Open Source software code developed by third parties, including software code subject to the GNU General Public License ("GPL"), GNU Lesser General Public License ("LGPL"), or other Open Source Licenses. The Open Source code used can be found at this site:

http://www.arubanetworks.com/open_source

Legal Notice

The use of Aruba Networks, Inc. switching platforms and software, by all individuals or corporations, to terminate other vendors' VPN client devices constitutes complete acceptance of liability by that individual or corporation for this action and indemnifies, in full, Aruba Networks, Inc. from any and all legal actions that might be taken against it with respect to infringement of copyright on behalf of those vendors.

Warranty

This hardware product is protected by the standard Aruba warranty of one year parts/labor. For more information, refer to the ARUBACARE SERVICE AND SUPPORT TERMS AND CONDITIONS.

Altering this device (such as painting it) voids the warranty.

SKU: EDU-ICAC-ILT-v23.11

January 2023

Implementing Aruba Campus Access

STUDENT GUIDE TABLE OF CONTENTS

Module 0: Course Overview	1
Course Prerequisites	2
Introductions	3
Logistics	4
Sign Up for HPE Learner ID	5
Welcome to Aruba Central Training	6
Agenda	7
Resources	12
Module 1: Introduction to Aruba Solutions	13
Objectives	14
Aruba ESP Architecture	15
Topologies & Encapsulation	21
WLAN Review	36
Aruba Product Portfolio	45
Aruba Central Review	58
Automation and REST API	74
Knowledge Check	77
Lab Activity	85
Module 2: Building Wired Infrastructure	87
Objectives	88
VSX	89
Lab Activity 1	127
OSPF	129
Virtual Routing and Forwarding	166

Policy-Based Routing	175
Lab Activity 2	180
Aruba Central Template Groups	182
Knowledge Check	189
Lab Activity 3	196
Module 3: Wireless Infrastructure with Aruba Gateways	198
Objectives	199
AOS 10 AP & Central Groups	200
Deploying Aruba Gateways	208
Initial Gateway Setup	218
Introduction to Gateway Cluster	222
Knowledge Check	227
Lab Activity	231
Module 4: Tunneled WLAN Architecture	235
Objectives	236
Tunneled Mode WLAN	237
Clustering Continued	253
Client Monitoring	265
Knowledge Check	268
Lab Activity	274
Module 5: Secure Enterprise WLAN	277
Objectives	278
Introduction to Authentication	279
802.1X Authentication	291
Firewall Roles and Role Derivation	297
CoA Support	328
Knowledge Check	336
Lab Activity	343
Module 6: Guest and Captive Portal	346
Objectives	347
Introduction to Guest Access	348
Guest Authentication Process	356
Guest WLAN Configuration	363

MAC Caching	371
Knowledge Check	375
Lab Activity	380
Module 7: Wireless Authentication for IoT	382
Objectives	383
MPSK Introduction	384
MPSK Authentication Process	389
Local MPSK Configuration	396
Knowledge Check	403
Lab Activity	407
Module 8: Mixed Mode Architecture	409
Objectives	410
Aruba Mixed Mode WLAN	411
Mixed Mode Configuration	418
Knowledge Check	425
Lab Activity	428
Module 9: Gateway Cluster Deployments	430
Objectives	431
Cluster Deployment Options	432
Cluster Configuration	441
Cluster Monitoring	446
Secondary Gateway	451
Knowledge Check	454
Lab Activity	457
Module 10: Wired Port Access	459
Objectives	460
Wired Authentication	461
Configuring 802.1X	472
MAC Authentication (MAC-auth)	484
RADIUS Attributes & Roles	495
Lab Activity 1	518
Dynamic Segmentation	520
Wired Captive Portal	534

Additional Wired Authentication Features	555
Knowledge Check	564
Lab Activity 2	572
Module 11: VXLAN and GBP	574
Objectives	575
Distributed Overlay & VXLAN Introduction	576
VXLAN GBP	587
Knowledge Check	599
Lab Activity	605
Module 12: Security and Availability Features	607
Objectives	608
IDS/IPS	609
Authentication Survivability	622
Central Survivability	627
Management User Authentication	631
Knowledge Check	635
Lab Activity	642
Module 13: Traffic Optimization and QOS	644
Objectives	645
WLAN Optimization	646
AirGroup	656
Quality of Service	664
UCC	695
LLDP MED	701
Knowledge Check	708
Lab Activity	718
Module 14: Monitoring	720
Objectives	721
Central Alerts	722
Device Insights	730
User Experience Insight	737
Knowledge Check	757
Lab Activity	763

Next Steps.....765

Confidential - For Training Purposes Only

Rev 23.11



Course Overview

Implementing Aruba Campus Access

Hello and welcome to the course, Implementing Aruba Campus Access!

Confidential - For Training Purposes Only

Course Prerequisites: Ensure Your Success

IMPORTANT

For success in this course, you should know the content in
Aruba Campus Access Fundamentals

- Network components, models
- Access switch configuration
- STP, VRRP, VSF, VSX
- Aruba ESP, Central, ClearPass
- Onboard devices to Central
- Basic Wi-Fi concepts, WLANs
- Basic AAA, 802.1X, roles, rules
- Basic WLAN security
- Password recovery, factory reset
- Basic Central diagnostic tools

Your success is important to us! To do well in this course, you must know the content in the Aruba course Aruba Campus Access Fundamentals as well as having practical experience. This includes the specific topics you see here on this slide.

If this is not the right course for you, please contact the person you worked with to sign up for this course to get scheduled for the Get the Edge training. If you choose to leave before the end of day one of this course, you will not be charged for this course.

Introductions



When it is your turn, please introduce yourself. Please provide the following information:

Name

Location

Company

Job responsibilities

Networking experience

Class expectations

Logistics



Breaks



Lunch



Restrooms



This slide contains the logistics for the class: when breaks and lunch will typically occur and where the restrooms are located. If you are taking this class virtually, please remember to mute your microphone, when not interacting with the class, in order to minimize disruptions.

Confidential - For Training Purposes Only

First Things First: Sign Up For Your HPE Learner ID

GET STARTED: https://certification-learning.hpe.com/TR/get_started.html

- **Required** to register for certification exams on PearsonVUE
- **Required** to receive credit for Aruba courses in your training history
- View your certification and training history in HPE MyLearning
- Download your e-certificates and badges in HPE MyLearning



If you haven't already done so, Aruba strongly encourages you to register for your HPE Learner ID at the URL shown on here and included in the resources tab at the top of this course slide: https://certification-learning.hpe.com/TR/get_started.html.

You need a Learner ID to register for exams and receive credit for the courses in your training history. You can also access HPE MyLearning to view your certification and training history, and to download your e-certificates and badges. To access HPE MyLearning, visit: <https://www.mylearninghpe.com/hpcpbenefits/hpelogin.aspx>.

Welcome to Aruba Campus Access Training!

QUICK LINKS:

Aruba Certification & Training:
arubanetworks.com/certification

Aruba Training Schedule & Labs
https://inter.viewcentral.com/events/cust/cust_tracks.aspx?cid=aruba&pid=1

Airheads Community:
community.arubanetworks.com/


HPE Press Study Guides:
hpepress.hpe.com

Contact Us:
arubatraining@hpe.com



TAKE THE NEXT STEPS TO GET CERTIFIED

Start here!



*Study Guides on
HPE Press*



Thanks for choosing to learn about Aruba Campus Access Solutions! Aruba hopes you will take the next step to get certified on Aruba's technologies.

To learn more about our certification & training program visit Aruba online at:
<https://arubanetworks.com/certification>.

Aruba encourages you to join the Airheads community, a vibrant group of networking professionals (Aruba employees, partners and customers) where you can participate in discussions and watch free learning videos on the Airheads Learning Community.

To reinforce what you learn in this course visit HPE Press for study guides:
<https://hpepress.hpe.com/>.

If you have any questions or concerns we welcome you to contact the Aruba support team at arubatraining@hpe.com.

Agenda Day 1

1: Intro to Aruba Solutions

2: Building Wired Infrastructure

3: Wireless Infrastructure with
Aruba Gateways



This slide has the first set of modules you'll explore during the first day. Please note that this slide and the following slides provide a general guideline of what is covered each day and that no all classes will follow this schedule verbatim.

Agenda Day 2

4: Tunneled WLAN Architecture

5: Secure Enterprise WLAN

6: Guest and Captive Portal



This slide has the list of modules that are covered during the second day.

Agenda Day 3

7: Wireless Authentication for IOT

8: Mixed Mode Architecture

9: Gateway Cluster Deployments



This slide has the list of modules that are covered during the third day.

Agenda Day 4

10: Wired Port Access

11: VXLAN and GBP

12: Security and Availability
Features



This slide has the list of modules that are covered during the fourth day.

Agenda Day 5

13: Traffic Optimization and QOS

14: Monitoring



This slide has the list of modules that are covered during the fifth and final day.

Confidential - For Training Purposes Only

Resources



- Validated Reference Design (VRD) Guides
 - <https://community.arubanetworks.com/t5/Validated-Reference-Design/tkb-p/Aruba-VRDs>
- Aruba Support
 - <http://asp.arubanetworks.com/>
- Airheads
 - <http://community.arubanetworks.com/>
- Aruba Solution Exchange
 - <http://ase.arubanetworks.com>



You may find these resources quite valuable in your networking journey, especially if you intend to expand your learning and knowledge and participate in the active Aruba community.

Confidential - For Training Purposes Only

Rev. 23.11




Introduction to Aruba Solutions

Implementing Aruba Campus Access



Welcome to module 1! This module provides an introduction to Aruba solutions.

Confidential - For Training



Objectives

- Describe Aruba ESP architecture
- Leverage topology knowledge to design better networks
- Describe overlay technologies that can improve security, connectivity, and scalability
- Explain WLAN technologies
- Use your knowledge of the Aruba portfolio to choose correct solutions
- Describe Aruba Central and REST API automation

MOD 1-2



After completing this module, you will recognize fundamental concepts that are foundational for mastering network technology. You will be able to:


- Describe Aruba ESP architecture.
- Leverage topology knowledge to design better networks.
- Describe overlay technologies that can improve security, connectivity, and scalability.
- Explain WLAN technologies.
- Use your knowledge of the Aruba portfolio to choose correct solutions.
- Describe Aruba Central and REST API automation.

Aruba ESP Architecture

Leverage Aruba's Edge Services Platform (ESP) for an integrated, secure, intelligent network architecture. This section provides a quick overview of Aruba's ESP.

Confidential - HP Training

Aruba ESP Overview



Aruba ESP is a cloud-native architecture that automates, unified, and protects your systems. Inherent to this solution are four key components

Artificial Intelligence Operations (AIOps) helps you to identify network, security, and application performance issues before they affect users. It eliminates many manual troubleshooting tasks, and provides optimization tips as your deployment grows and changes.

A unified infrastructure eliminates management silos. Manage all wired, wireless, and SD-WAN components in one centralized, cloud-based place: Aruba Central. Users get a consistent experience across the campus, branch, data center, and remote locations.

Aruba Central is the hub of your Aruba ESP operations. You get cloud-based visibility, control, policy, management, and security.


Zero Trust Security gives you full visibility, control, and enforcement. Ensure the same controls applied to campus or branch networks also extend to remote workers. See what's connected, use identity and roles to enable "least access" to IT resources, and dynamically change access privileges based on real-time threat data.

This architecture supports your entire operations, right on down to the Intelligent Edge. This is where people, devices, and things connect to the digital world. Aruba solutions support and enhance your Intelligent Edge with features specific to Internet of Things (IoT) and mobility solutions.

Let's explore Aruba ESP in more depth.

Confidential - For Training Purposes Only

Aruba ESP Campus Network Architecture




MOD 1-5

That artificial intelligence we discussed on the previous slide is in the cloud: Aruba Central. You get a unified management plane for all platforms (wired, wireless, SD-WAN) and all locations. This intelligence is fed with telemetry data from your own unified infrastructure, and from Aruba's global awareness of real-time issues and machine learning.

Of course, Zero Trust Security does provide perimeter security, but that's only a start. Security is inherently embedded into all Aruba ESP components and solutions. Users get consistent, quality connectivity, regardless of their location. You get flexible and highly reliable designs, simplified operations, and accelerated service delivery.

Aruba ESP is the industry's first platform that is purpose-built for the new requirements of the Intelligent Edge. It's a best-in-class industry network solution for modern organizations. This is all neatly organized into three layers.

Aruba ESP Layers



With Aruba ESP you get onboarding, provisioning, orchestration, security, analytics, location tracking, and management. AI Insights reveal issues before they impact users. Intuitive workflow-centric navigation allows you to easily accomplish tasks, using views with multiple dimensions of correlated data. You create policies centrally, and easily implement them with features like dynamic segmentation. This is because the Aruba ESP architecture is built in distinct layers.

The Aruba ESP campus connectivity layer

The connectivity layer is implemented on Aruba CX Ethernet switches. You get low latency and high bandwidth on a fault-tolerant platform designed to carry campus traffic from the access layer to the core. Wireless connectivity is over industry leading APs. You can tunnel traffic to centralized controllers or for bridge locally at the AP.

Aruba ESP campus policy layer

The policy layer uses overlay technologies and traffic filtering to isolate user and application traffic. You can tunnel data traffic back to a gateway cluster for centralized enforcement or let the local switch fabric handle it - policy enforcement at every node in the network. You typically use ClearPass Policy Manager (CPPM) as a RADIUS authentication interface into an LDAP user directory service. Here you define user roles with associated enforcement policies. Device Insights determines endpoint security posture, from information gathered from the network.

Aruba ESP decouples the policies from IP addressing, for powerful, scalable policy management. Traffic tunneled to a gateway cluster is tagged at ingress with a user role. This determines how the gateway forwards traffic. Traffic in a distributed fabric is tagged using the VXLAN-Group Based Policy (GBP) feature to assign a role ID to every frame, ensuring consistent enforcement across LAN, WLAN, and WAN.

Note: In AOS 8, the term mobility controller (MC) was used to define controllers managing APs and users while gateway was used to define access for SD-WAN and SD-Branch access. Starting in AOS 10, these functions have been merged into one software set and the device performing these operations is referred to as a gateway.

Aruba ESP campus services layer

Aruba Central provides a cloud management and services delivery platform for the end-to-end Aruba ESP solution. Use this AI-driven platform for automation workflows. For example, Aruba Central's NetConductor eases the deployment of EVPN-VXLAN overlay fabrics and role-based access policy and MultiEdit makes it easy to manage the configurations of one or more switches simultaneously.

Most organizations purchase and deploy Aruba gateways, switches, and APs using Central as the management platform. But large organizations can deploy ESP services on-premise or in a private cloud. Organizations that operate the network as a managed service can deploy it through HPE GreenLake for Aruba.

Note about 5G

5G is the fifth-generation technology standard for mobile networks and includes the following benefits: larger channels to increase speed, lower latency to improve responsiveness. Aruba integrates cellular into the enterprise network. Aruba Air Pass™ enables Wi-Fi enabled devices with SIM credentials from major cellular network operators to automatically connect to enterprise networks. This improves the in-building cellular experience because Wi-Fi is automatically offloaded onto the enterprise network. Using Air Pass, users can send and receive Wi-Fi calls and text messages, and the Wi-Fi network can deliver high-speed data offload.

Topologies & Encapsulation



After that high-level architecture discussion, let's move on to topologies and encapsulation.


Confidential - For Training

Two-Tier Campus LAN

Overview: L2 access switches connect to L2/L3 core

Use case: Small buildings with few wiring closets/access switches

Use case: Large environments when fiber cables from each closet are homed to a single location



MOD 1-8



A two-tier wired architecture includes access switches or switch stacks connected to a dual-switch collapsed core. The access switches typically provide layer 2 services to connected endpoints and connect to core switches, which provide both layer 2 and layer 3 services. The two-tier design is well suited for small buildings with few wiring closets and access switches. It also works well in larger environments when the fiber cables from each wiring closet are homed in a single location. The figure illustrates this design.

Collapsed core layer

Use Aruba CX switches that support VSX redundancy to allow access switches and other devices the option of connecting over a redundant, VSX LAG (MC-LAG) layer 2 connection(s). A foundation for establishing a network overlay can be built within the two-tier topology by configuring OSPF peering between each access and core switch and adding a loopback interface to each access switch in the OSPF backbone area.

Access layer

Use Aruba CX switches that support VSF stacking for simplified growth in the network closet. In 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 from the access layer. The access layer enables Aruba ESP Colorless Ports by configuring port policies to allow 802.1x dynamic authentication and network configuration and enables layer 2 protection mechanisms such as Loop Protection, BPDU Filter, Root

Guard, and BPDU Protection. To simplify the network as much as possible, all routing is typically performed on the core devices.

Confidential - For Training Purposes Only



Moving to a Three-Tier Campus LAN

Network growth is

producing cross-campus traffic beyond capacity of single collapsed core
beyond a small number access aggregation points
beyond a few building aggregation points

Use case

Large networks, thousands of users
Physical wiring layout not suited for two-tier



Organizations move to a three-tier network design for the following primary reasons:

- Network growth producing cross-campus traffic beyond the capacity of a single collapsed core.
- Network growth beyond a small number access aggregation points.
- Network growth beyond a few building aggregation points.

The three-tier campus design is recommended for large networks with thousands of users or where the physical wiring layout is not suited for a two-tier design. Let's check it out.


Three-Tier Campus LAN

Aruba CX switches w/ports & speeds to service b/w requirements of aggregation layer

Most critical part! Routed traffic only: reduces complexity

Separate switches, dual ECMP to aggregation switches

Point-to-point links: propagate up/down changes immediately (with ECMP, millisecond convergence)



Layer 3 services for wired network hosts are moved from the core to VSX pairs of aggregation switches. A pair of core switches join the aggregation switches together using high-speed layer 3 links and equal-cost multipath (ECMP) links. Additional capacity between pairs of aggregation switches is added by increasing the number of links between the core and aggregation switches. The access switches typically remain layer 2 only. The slide illustrates this design. Let's explore the three layers.

Core layer

Use Aruba CX switches with sufficient ports of appropriate speeds to service the full bandwidth demands of the campus aggregation layer. The core layer of the LAN is the most critical part of the campus network and standalone cores. The core layer reduces network complexity by carrying only routed traffic. The standalone core uses separate switches, acting independently of one another, with dual ECMP connections to all aggregation switches. ECMP is an advanced routing strategy in which next-hop packet forwarding occurs over multiple paths with identical routing metric calculations.

When considering core topologies, it is important to use point-to-point links because link up/down changes propagate almost immediately to the underlying protocols, helping to speed convergence. Topologies with redundant ECMP links are the most deterministic and convergence is measured in milliseconds.

Aggregation layer

Use Aruba CX switches that support VSX redundancy with sufficient ports of appropriate speeds to service the full bandwidth demands of the campus access layer. In layer 2 access designs, use VSX MC-LAGs to provide efficient layer 2 connectivity to the devices.


In most designs, the aggregation layer of the LAN limits cable distance to access devices, isolates the network core from layer 2 traffic, and provides layer 3 services to access VLANs. In campus designs using layer 3 access, routing moves to the edge switches and the aggregation devices fulfill a simpler, transit-only function.

Consider running OSPF to provide loopback reachability to access layer devices even in networks relying on layer 2 access. Enabling loopback reachability to all devices in the campus LAN ensures the ability to leverage a growing range of options for network automation and orchestration.

Access layer

Use Aruba CX switches that support VSF stacking for simplified growth in the network closet. In 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. At the access layer, enable Aruba ESP Colorless Ports by configuring port policies to allow 802.1x dynamic authentication and network configuration. Also enable layer 2 protection mechanisms such as Loop Protection, BPDU Filter, Root Guard, and BPDU Protection. To reduce load on the network core, all routing is typically performed on the aggregation devices.


WLAN Bridge Mode Deployment



WLAN bridge mode provides an easy solution when tunneled traffic is not needed and advanced gateway features are not required. In this mode, wireless traffic is bridged directly from the AP into the wired infrastructure. The access switch ports for the APs are trunked to provide SSID-to-VLAN connectivity. The AP handles the packet encryption, user authentication, and policy enforcement functions. Aruba Central handles RF management, key management, live upgrades, monitoring, and troubleshooting.

Aruba Campus Access Associate training covers the details of bridge mode deployment in depth.


Tunnel Mode Deployment



You can add Aruba gateways to a greenfield design (completely new design) or an existing bridge mode deployment. The only requirement is to run the same software version on APs and gateways. A tunnel mode deployment offers robust security features and maximum operational flexibility. Deploy gateways individually, or clustered for increased redundancy and scale. Clusters are automatically created by adding gateways to the same group in Central (this is unlike ArubaOS 8, where you had to manually create controller clusters).

Tunnel mode increases visibility into applications, which helps prioritize business-critical data. This model also provides micro-segmentation, dynamic RADIUS proxy, and encryption over the LAN. Seamless roaming is supported across an entire layer 3 campus.

Overlay Networks



An overlay network is a virtual network that runs on top of a physical underlay network. Even if you have never heard about this terminology before, you have probably seen it. The slide shows a simple example: a GRE tunnel that runs on top of a physical underlay network.

Note: The slide shows just one overlay example: between an AP and a gateway. Other overlay tunnel scenarios might be from switch to switch or switch to gateway, as you soon learn in this course.

The underlay network's job is simple: its only job is to get packets from point A to point B. Of course, the underlay should be designed for stability, scalability, and predictability. When you implement a layer 3 underlay network, you can use routing protocol like OSPF and load balance traffic on redundant links.

Overlay networks provide a mechanism to deploy flexible topologies that meet the demands of ever-changing endpoints and applications. By fully decoupling user data traffic and associated policy from the physical network topology, overlays enable on-demand deployment of layer 2 (L2) or layer 3 (L3) services. Overlays also make it simpler to carry device or user role information across the network without requiring all devices in the path to understand or manage the roles.

Technologies such as GRE and IPsec have been used to create overlays in the campus and WAN space for many, many years. Virtual Extensible LAN (VXLAN) is now a new and standardized method to create distributed, scalable, and manageable overlay networks in the

campus.

Confidential - For Training Purposes Only


Aruba Centralized vs Distributed Overlays

User Based Tunnelling - GRE

Centralized Overlay

Simplified operations and advanced security features on gateway devices

AP or switch to gateway




EVPN-VXLAN

Distributed Overlay

Dynamic tunnels between switches

Consistent, continuous policy enforcement

Switch-to-switch tunnels




MOD 1- 14

Aruba provides the flexibility to choose between centralized or distributed overlays to address a range of traffic engineering and policy enforcement requirements. User-based tunneling (UBT) is a centralized overlay architecture that provides simplified operations and advanced security features on gateway devices. EVPN-VXLAN is a distributed overlay architecture that enables dynamic tunnels between switches anywhere on a campus, providing consistent and continuous policy enforcement across the network. Both overlay models support the colorless ports feature, which enables automated client onboarding and access control for ease of operations.

UBT: Centralized Overlay

- Tunnel user traffic to gateway cluster to enforce policy
- Selectively tunnels traffic based on a user or device role
- Client policies typically RADIUS server-assigned

Firewall, DPI,
app visibility, b/w control




MOD 1- 15

User-based tunneling (UBT) is a centralized overlay that allows you to tunnel specified user traffic to a gateway cluster to enforce policy using services such as firewalling, deep packet inspection (DPI), application visibility, and bandwidth control. UBT selectively tunnels traffic based on a user or device role. The policies associated with each client usually are assigned through a RADIUS server such as ClearPass Policy Manager (CPPM).

Aruba ESP centralized overlays use the Generic Routing Encapsulation (GRE) protocol. Let's examine GRE encapsulation.

GRE Encapsulation




MOD 1- 16

Aruba ESP centralized overlays are implemented using the Generic Routing Encapsulation (GRE) protocol. GRE is a standardized layer 3 IP protocol. GRE enables an access layer switch or AP to tunnel client traffic to a gateway or gateway cluster for policy enforcement.

Note: Because GRE adds additional overhead—GRE header, IP header, and an additional Ethernet header—It is highly recommended to enable jumbo frames across all paths between the Aruba APs, switches, and gateways.

EVPN-VXLAN: Distributed Overlay

- Tied to policy-based micro-segmentation, based on global roles, across the infrastructure
- Fully automated overlay that extends L2 connectivity
- **EVPN: Control Plane** - BGP-driven for overlays over IP networks
- **VXLAN: Data Plane** - tunneling protocol that supports 16 million L2 broadcast domains



Distributed overlays are an evolution of the traditional campus network design, built using EVPN-VXLAN on highly available underlays. These overlays are tied to a full policy-based micro-segmentation, based on global roles, across the entire network infrastructure. Role-based policies abstract policy from the underlying network and enable flexible and simplified policy definition and enforcement. This is provisioned by fully automating the overlay that extends layer 2 connectivity over an existing physical network and layer 3 underlay.


EVPN-VXLAN is an open standards suite to create more agile, secure, and scalable networks in campuses and data centers. It consists of:

Ethernet VPN (EVPN): a BGP-driven control plane for overlays that provides virtual connectivity between different layer 2/3 domains over an IP or MPLS network.

Virtual extensible LANs (VXLAN): a common network virtualization tunneling protocol that expands the number of layer 2 broadcast domains to 16 million from the 4,000 available using traditional VLANs.

Aruba ESP implements EVPN-VXLAN overlays on an IP underlay using redundant, layer 3 links for high-speed resiliency and maximum bandwidth utilization. Historically, EVPN-VXLANs were deployed in data centers to solve VLAN connectivity issues. However, many customers need the same type of scalability and flexibility in campus designs; therefore, the use of EVPN-VXLANs is becoming more commonplace today in campus networks.

VXLAN Encapsulation



Tip! Configure jumbo frames in the fabric! Accommodates additional encapsulation MOD 1- 18

VXLAN encapsulates L2 Ethernet frames in L3 UDP packets. These VXLAN tunnels provide both L2 and L3 virtualized network services to connected endpoints. A Virtual tunnel end point (VTEP) is a switch function that originates or terminates VXLAN tunnels. Like a traditional VLAN ID, a VXLAN network identifier (VNI) identifies an isolated layer 2 segment in a VXLAN overlay topology. The VNI associates tunneled traffic with the correct corresponding layer 2 VLAN or layer 3 route table so the receiving VTEP can forward the encapsulated frame appropriately. Symmetric Integrated routing and bridging (IRB) enables the overlay networks to support contiguous layer-2 forwarding and layer-3 routing across leaf nodes.

A VTEP encapsulates a frame in the following headers:

IP header: IP addresses in the header can be VTEPs or VXLAN multicast groups in the transport network. Intermediate devices between the source and destination forward VXLAN packets based on this outer IP header.

UDP header for VXLAN: The default VXLAN destination UDP port number is 4789.

Note: Configure jumbo frames on all underlay links in the fabric to ensure the accurate, non-fragmented transport of additional encapsulation.



WLAN Review

This section provides a review of WLANs. This topic is covered indepth in the Aruba Campus Access Fundamentals course.

Confidential - For Training

RF Bands and Wi-Fi Generations

2.4 GHz **5 GHz** **6 GHz**

1999	802.11b	
1999		802.11a
2003	802.11g	
2009	802.11n	802.11n
2013	802.11ac	802.11ac
2018	802.11ax	802.11ax
2020		802.11ax extension

Early IEEE 802.11 standards, like the original 802.11, 802.11b, a, and g standards, were great at the time, with data rates up to 54 Mbps, and support in both the 2.4 and 5GHz frequency bands. These standards are slow and inefficient by today's standards.

The 802.11n amendment added High Throughput (HT) speeds, theoretically up to 600Mbps, with 300 Mbps typical. It accomplished these speeds by increasing the channel width, using advanced 64-Quadrature Amplitude Modulation (QAM), and techniques like Multiple Input/Multiple Output (MIMO).

802.11ac only operates in the 5GHz frequency, with theoretical data rates up to 6.93 Gbps. It accomplishes this by using Very High Throughput (VHT) methods, with higher channel width and more complex modulation techniques. 802.11ac has sometimes been referred to as Wi-Fi 5.


802.11ax is the latest standard and operates on both 5 GHz and 2.4 GHz frequencies, with a maximum data rate of 4.8 Gbps. The big news here is a technique called Multi-User MIMO, enabling multiple transmitters at the same time. This early version of 802.11ax has also been referred to as Wi-Fi 6.

An 802.11ax extension was published in 2020, supporting up to 7.8Gbps using 1024-QAM with MCS 10/11. There are two big news items here: it operates in the 6 GHz band and supports a 1 gigabit channel width. This version of 802.11ax has been referred to as Wi-Fi 6E.

OK, that's the main IEEE standards for Wi-Fi. Let's look at the Wi-Fi Alliance next.

Confidential - For Training Purposes Only

WLAN Logical Configuration: BSS




A basic service set or BSS includes an AP radio and all associated clients in a coverage area. An AP with two or more radios has more than one BSS. When you define an SSID, it is assigned a different 48-bit MAC address on each of the AP radios that advertise it. The radio MAC address, known as the BSSID, identifies the BSS. All WLAN clients associated to the radio use this BSSID as the destination MAC address in their 802.11 frames.

All BSSIDs are based on the AP's 2.4, 5, and/or 6 GHz radio "base" MAC address(es). So the first SSID you define on AP1 might end with aa:aa:aa, as shown on the slide. The next SSID you define on that AP might end with aa:aa:ab, the next ends with aa:aa:ac, and so on.

An extended service set (ESS) includes all APs and clients associated to the same logical network name, commonly known as an SSID. The SSID name is case sensitive and identifies the WLAN to the client. An AP beacon sends both the BSSID and the SSID into the air for discovery of the WLAN.

WLAN Logical Configuration: ESS

One SSID broadcast on multiple radios



In this slide you see an ESS applied to multiple APs. This creates a blanket RF coverage through which users can seamlessly connect and roam. Although each AP radio advertises the same SSID, each radio (on every AP) has its own unique BSSID (wireless MAC address).

Types of 802.11 Frames

Management frames

Manage the wireless network

Control frames

Control access to the medium

Data frames

Carry client data



MOD 1- 23

There are three 802.11 frame types:

- Management
- Control
- Data frames

Management frames are used to manage the wireless cell or BSS. This includes probing, associating, roaming, and disconnecting clients from the BSS. Control frames are used to control access to the medium and are used for frame acknowledgement. Control frames only contain a header and trailer, no body. Data frames are used to carry the users' data.

Ethernet Frame vs 802.11 Frame

802.3 – two address fields

6	6	2	46-1500	4
Dst address	Src address	Type	Data	CRC

802.11 – four address fields


2	2	6	6	6	6	6	0-2312	4
Frame Ctl	Duration/A-ID	Address 1	Address 2	Address 3	Seq. Ctl	Address 4	Body	CRC



MOD 1- 24

The IEEE 802.11 frame structure is different from wired Ethernet 802.3 frames. As you can clearly see in the slide, the wireless frame is more complex, with four address fields as opposed to only two in Ethernet. In the upcoming slides, you will learn how these address fields are used in different scenarios.

802.11 Addressing



To DS	From DS	Meaning	Address 1	Address 2	Address 3	Address 4
0	0	Direct STA-STA; Mgmt and ctl	RA = DA	TA = SA	BSSID	N/A
0	1	From DS to STA	RA = DA	TA = BSSID	SA	N/A
1	0	From STA to DS	RA = BSSID	TA = SA	DA	N/A
1	1	From AP to AP	RA	TA	DA	SA

Address fields use is controlled by **To_DS** and **From_DS** bit

MOD 1-25



The way four address fields in the 802.11 frame work is controlled by two bits in the Frame Control (Frame Ctl) field: TO_DS and From_DS, where DS stands for distribution system.


802.11 Addressing Scenarios

Ad-Hoc




T	F	Addr.1	Addr.2	Addr.3
0	0	00:01	00:05	AA:BB
0	0	00:05	00:01	AA:BB

Infrastructure



T	F	Addr.1	Addr.2	Addr.3
0	1	00:01	00:02	00:04
1	0	00:02	00:01	00:04

Repeater



T	F	Addr.1	Addr.2	Addr.3	Addr.4
1	1	00:02	00:03	00:01	00:04
1	1	00:01	00:02	00:01	00:04

MOD 1- 26

The To_DS (To distribution system) and From_DS (From Distribution System) fields indicate address field usage. A DS is typically a device like an access point (AP).

Ad-hoc mode is a direct STA (station) to STA communication, like two PCs interacting with each other directly. There is no distribution system (AP) to go to or from, so it makes sense that both bits would be low. Address 1 is used as the RF radio receiver address, which is also the end destination address for this frame. Address 2 is the RF radio transmitter address which is also the source address for this frame. Like all WLAN communications an SSID is used for Ad-Hoc, the WLAN is technically called an independent SSID or IBSS. This is placed in the address 3 field.

When the To/From bits are set to 01, the frame is transmitted from the wired network (the DS; i.e., the AP) toward a WLAN client. Address 1 is the receiving client's MAC address, address 2 is the AP transmitter's BSSID MAC address, and the SA is MAC address of whatever wired Ethernet device originated this frame. This is most likely the WLAN client's default gateway (DG).

When the To/From bits are set to 1/0, the WLAN client (STA) is likely sending a frame back to its DG, via the AP. So the AP's BSSID is the receiver address (RA), the STA is the transmitter address (TA), and the DA is the client's DG on the wired Ethernet system or DS.

When the To/From bits are set to 1/1, this is used in a repeater scenario, when frames must go from DG-to-AP-to-AP-to-STA or from STA-to-AP-to-AP-to DG.

Aruba Product Portfolio

Let's now look at Aruba's product portfolio and some important Aruba features and technologies.

Confidential - For Training

Aruba CX Switching

Next-gen switching designed for the network operator

AOS-CX

Simplicity from access to DC




Aruba Central

Unified mgmt, intelligent configs



Aruba NAE

Visibility with analytics built-in




Aruba VSX live Upgrades

High Availability



Dynamic Segmentation

IoT and User Security



Gen7 ASIC

Data center class performance



No extra license required to enable ASIC functions

MOD 1-28

The following are advantages of AOS-CX switches:


- A single operating system across your entire environment spanning from edge to data center.
- Aruba Central can be used as the management tool. Central provides a unified infrastructure single pane of glass Central with multi-edit mode for advanced configs and to orchestrate common configuration changes across multiple CX switches.
- Aruba Network Analytics Engine (NAE) is a tool to help you with troubleshooting. NAE provides network operators with distributed, real-time insights into network health for a faster, more proactive way of monitoring, detecting, and troubleshooting IT- and user-impacting issues.
- Aruba Virtual Switching Extension (VSX): technology bundles multiple physical switches into a single logical switch. This ensures high availability and support for live upgrades to eliminate maintenance windows and reduce downtime.
- Dynamic segmentation is a way to unify network access and network security regardless of whether the user is using wired or Wi-Fi connection.
- The top-shelf, Gen7 ASIC is a microprocessor specialized in processing network traffic at faster than wired speeds. It offers programmable, data center class performance with no

bottlenecks and with no additional switch software licensing or subscriptions to leverage it.

Note: In a lot of Aruba documentation, you might see terms like ArubaOS CX, Aruba CX or AOS-CX in reference to the Aruba switches running the CX software/operating system. To keep it simple, this course will typically use the term AOS-CX when referring to the Aruba switches running the CX software. Please note that the term AOS-S or just AOS, refers to the older ArubaOS switches based originally on the ProCurve hardware platform. This course only covers the use of AOS-CX switches in campus deployments.

Confidential - For Training Purposes Only

Aruba CX Product Portfolio



The CX portfolio includes modular access, aggregation, and core switches. At the access and aggregation layers you get PoE, multi-gig, secure segmentation and stacking. At the campus core and data center, you get high availability, bandwidth, and speed, with 40/100 GbE scalability.

Here is a brief summary of the AOS-CX switches:

Aruba CX 6000 and 6100

- Overview: ideal for smaller networks of around 100-500 devices
- Positioning: branch offices, mid-market enterprise, and SMBs looking for a reliable, cost-sensitive entry level access. The CX 6100 has 10G uplinks.

Aruba CX 6200

- Overview: These fixed-port access switches support 8-member stacking with built-in 1/10 GbE uplinks and PoE.
- Positioning: Ideal for branch offices, mid-market enterprise, and SMBs looking for a reliable, simple to deploy, cost-sensitive access solution.

Aruba CX 6300

- Overview: A top-of-rack (ToR) model with back to front airflow designed for cost-effective 1 GbE ToR and OOBM data center deployments.
- Positioning: Ideal for use in business-critical networks with redundant power and fans, with IT initiatives driving to Wi-Fi 6, plans for IoT, and looking for automation to help with

rollouts. Supports bandwidth intensive applications, plus investment protection with future-ready Smart Rate multi-gig and high-power PoE.

Aruba CX 6400

- Overview: Versatile 5- and 10-slot high availability (HA) chassis with up to 28 Tbps capacity, 1/10/25/40/50/100 GbE line rate interfaces, and high-density Smart Rate multi-gig and high power 60W PoE.
- Positioning: Ideal for use from mid-large enterprise campus access to core and some data center use cases that require performance, scale, and high availability.

Aruba CX8300

- Overview: Compact 1U (1 unit) switches with 6 Tbps of capacity with high speed 1/10/25 GbE, 1/10/40 GbE and 40/100 GbE connectivity and redundant hot swappable power and fans and options for front-to-back or back-to-front cooling for flexible hot-cold aisle deployments.
- Positioning: Ideal for enterprise aggregation, core ToR and end-of-rack (EoR) deployments. Aruba Fabric Composer speeds infrastructure provisioning and reduces risk by utilizing pre-engineered and tested Aruba and HPE integrations.

Aruba CX 9300

- Overview: The Aruba CX 9300 switch series is a next-generation 25.6 Tbps, 1U fixed configuration switch supporting 32-ports of 100/200/400GbE. The switch is an ideal solution for flexible, cost-effective, high-density 400GbE connectivity server, storage, and intra-fabric connectivity.
- Positioning: Designed for spine, core and/or aggregation, ToR or leaf, or EoR in the data center.

Aruba CX 8400

- Overview: A carrier-class, 8-slot chassis providing up to 19.2 Tbps of capacity with high density 10/40/100 GbE connectivity with fully redundant fabric, management, power and fans.
- Positioning: Ideal for large enterprise campus core and data center networks where High availability is a must and the largest scale (table sizes, buffers) is required.

Aruba CX 10000

- Overview: Provides advanced services to the data center such as distributed stateful firewall, encryption services, and stateful network address translation (NAT), and Zero Trust Networking.

- Positioning: On-premises enterprise data centers and private cloud.

CX 4100i: a layer 2 switch with ruggedized design and is a special use-case for challenging environments. Two fanless models can survive extended temperatures, ranging from -40 to +70 degree Celsius.

Confidential - For Training Purposes Only

AOS-CX Switch Virtualization: VSX & VSF

VSX on 10000, 9300, 8400, 8300, or 6400


Agg/Core layer- VSX

Manual: must config on both switches.

Separate control, mgmt, data planes mitigates “shared fate”

Reliability is vital – Agg/core connects to nearly everything


Easy to manage – sync mgmt. planes, MAC/ARP, and more



Access layer - VSF

Plug-and-play: a switch automatically joins with another switch.

Ease of deployment is important – large number of access switches



Virtual Switching eXtension (VSX) and Virtual Switching Framework (VSF) are similar virtual or logical switching Aruba technologies: to each other, and even to earlier virtualization techniques like Virtual Router Redundancy Protocol (VRRP). With these virtualization techniques multiple devices are perceived as a single device. Switch virtualization allows you to the use of the eliminate Spanning Tree Protocol (STP) and have active/active traffic forwarding so all uplinks are being used at the same time. This improves resiliency while reduces configuration and management overhead. But there are distinct differences between VSX and VSF, as summarized in the slide.

You use VSX at the campus network aggregation and core layers. You use VSF primarily at the campus access layer. The differences in features and function are intentional, since priorities differ between access and core.


Configuring VSF and backplane stacking is a plug-and-play affair: a switch, at default settings, automatically unites with another switch. This ease of deployment is especially important at the access layer, since you might have dozens, or hundreds of access switches.

You must manually configure VSX on both switches. VSX maintains separate control plane and management planes, and much of the data plane. This mitigates the shared fate nature of common planes, for improved uptimes: you get high availability and redundancy. This is especially important when aggregating all those access layer switches. And it remains easy to manage, because you can synchronize the management planes, and optionally, other features, MAC and ARP addresses, and protocol configurations, and more.

Dynamic Segmentation: Problem Solved

Challenge Unsecured devices

Some devices lack security mechanisms like 802.1X/EAP
Security cameras, IoT sensors, medical gear, card scanners
Low endpoint security increases potential attack vectors - risk



Solution Dynamic Segmentation

Improves your security stance while easing administration
ClearPass authenticates devices, tunnels traffic to an GW
Centralized, policy-based firewall and other security features
Eliminates attack vectors, mitigates low endpoint security

MOD 1-31



All modern desktops, laptops, tablets, and smart phones support robust security mechanisms, like 802.1X/EAP for authentication. But many devices lack these security mechanisms. This includes specialty devices, and some devices that rely on Power-over-Ethernet (PoE), such as security cameras, payment card readers, Internet of Things (IoT), and medical devices, to name a few. This lack of endpoint security can pose a serious risk to your infrastructure.


An Aruba switch with dynamic segmentation improves your security stance while easing administration. Dynamic segmentation has the switch authenticate these devices using ClearPass, and tunnels the traffic to an Aruba gateway, which centralizes formerly disparate access control lists (ACLs) and firewall rules into a single, intuitive policy system. You can continue to use these edge devices without compromising your network security.

Aruba Access Points

Wi-Fi 6 campus APs



Wi-Fi 6E campus APs



Wi-Fi 6 outdoor APs



Wi-Fi 6 ruggedized APs



Wi-Fi 6 remote AP



MOD 1- 32

Aruba APs provide edge connectivity that's intelligent, safe, and fast. With AI-powered mobility and enterprise-class design, Aruba APs deliver superior services and segmentation without the complexity that other vendors add. And now with Wi-Fi 6 and 6E, Aruba delivers critical security and reliability to mobile users, IoT, and latency-sensitive apps—even in crowded areas. Enterprises w/5G strategies also benefit from Aruba Wi-Fi 6E, serving as an on-ramp to 5G (in place of small cell/DAS solution). Different AP models represent different levels of performance and capacity as well as different operational environment: indoor, outdoor, and remote/hospitality.

Aruba Gateways



9004 and 9012 Series
Small/midsized branch
4 Gbps firewall throughput



7010/7024
Midsized/large branch w/16-24 ports
and PoE
4 Gbps firewall throughput



9200
Medium/large enterprise
40 Gbps firewall throughput



7200
Large enterprise
100 Gbps firewall throughput

When to use gateways:

- ✓ Network design simplification
- ✓ SD-WAN
- ✓ Security: tunneled WLANs, segmentation/Guest WLANs, network encryption
- ✓ Greater mobility: L2/L3 roaming across domains
- ✓ Dynamic Radius Proxy
- ✓ Scale: 5,000 clients/site or 500 APs/site*

* Validated best practice numbers/ not a hard technical limit

MOD 1- 33

The main reason you need Aruba gateways in a large-scale deployment is to simplify the network design and to make the solution easier to manage. The Aruba gateway is responsible for wireless traffic termination: all wireless traffic is being tunneled to the gateway and then passes through its firewall, and from there it is forwarded to its final destination. This eliminates an excessive span of VLANs. Aruba offers both gateways to support the needs of larger enterprises who require additional scalability, 24x7 operations, and greater security.

Note: The use of gateways are optional in a campus deployment (you can have a campus solution with only APs). However, gateways provide a much more scalable solution than using an AP-only solution.

You should use gateways when one or more of the following are true or desired:

Simplified network design

SD-WAN capabilities

Security: tunneled WLANs, segmentation/Guest WLANs, network encryption

Greater mobility: L2/L3 roaming across domains

Dynamic Radius proxy

Scalability: 5,000+ clients per site and/or 500+ APs per site

Confidential - For Training Purposes Only

AOS 8 to AOS 10 Comparison

AOS 8/InstantOS 8	AOS 10
IAP, CAP, RAP	AP, Microbranch AP
Mobility Conductor	Aruba Central
Mobility Controller/VC	Gateway
SD-Branch Gateways	SD-Branch & Mobility
Local or Cloud Mgmt	Cloud Management



MOD 1- 34

Aruba AOS 10 software simplifies AP deployment methods to only two AP modes: AP and microbranch AP. AOS 10 does not include an IAP cluster or IAP virtual controller concept (Instead, this functionality is inherent in Central for AOS APs). Additionally, AOS 10 does not include a mobility conductor (formerly known as mobility master) since all those features are now defined in Aruba Central. In other words, you get a simpler network architecture with AOS 10 software!

Additionally, in AOS 10, Aruba has renamed controllers to gateways, and merged the firmware code for SD-Branch and mobility implementations into one single firmware platform.



Aruba Central Review

In this section, you will review key aspects of Aruba Central. Please note that this was covered in the corresponding Aruba Campus Access Fundamentals course.

Confidential - Internal Training

Aruba Central

High security


Group config, less workload



SD-Branch, APs, gateways

Views, reporting

AI analytics, ML



Cloud-based management platform for Aruba networks

MOD 1- 36



Aruba Central supports group configuration, to provision and manage similar devices with less work. Continuous monitoring and AI-based analytics by Central provide real-time visibility and insight. AI machine learning (ML) leverages a growing pool of network data and deep experience. Aruba Central is very secure, with encrypted HTTPS connection and certificate-based and cloud authentication policy features.

You get comprehensive network views, device status and health, and application usage. You can monitor, identify and address issues with a data-driven dashboard, alerts, reports, and troubleshooting workflows. Aruba Central also runs part of the control plane for some network devices and solutions, using the Aruba SD-Branch solution and Aruba APs and gateways running AOS 10 firmware.

Flexible and Cost-Effective Cloud Networking

HARDWARE

Limited lifetime warranty

APs

IAPs
AOS 10



Switches

Aruba OS
Aruba CX



Gateways

SD-Branch
AOS 10



CLOUD MANAGEMENT

License per managed device
Transferrable between devices

1/3/5/7/10-year subscriptions
Includes tech support




MOD 1- 37



Aruba Central supports management, monitoring, and configuration for the Aruba IAPs (AOS 8), switches (older AOS and newer AOS-CX), and SD-Branch gateways, and gateways (AOS 10). And since it's a cloud-based management system, it only requires a simple 1, 3, 5, 7, or 10-year subscription license (per device) which includes technical support. Additional services/features are available.

Aruba Central 2.5.5 Licensing for Campus



Licenses based on 1,3,5,7, and 10-year subscriptions

MOD-1-38

Aruba Central supports foundation and advanced license keys. These keys are for APs, switches, SD branch gateways, and gateways running AOS 10. Foundation keys enable you to onboard, monitor, and manage Aruba devices. Advanced keys add functionality and features. Every device Aruba Central manages requires one license (subscription). Licenses are based on 1,3,5,7, and 10-year subscriptions. The licensing model provides the license types shown in the slide.


For gateways, there is also a Virtual Gateway (VGW) license, 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.

Security gateway licenses add security features related to IPS and IDS, security dashboard, and anti-malware. Base Licenses have the same features as the respective licenses but are limited to 75 clients.

Note: There are different licenses for different device families; i.e, switch Foundation licenses for the 6200 switch series and foundation licenses for 6400 switch series, and advanced gateway licenses for the 7000 gateway family, and advanced license for the 7200 gateway family.

To check the latest information on Aruba Central License keys, navigate to the Aruba Central online help (Help support icon > Documentation center) and select Aruba Central Licenses under Key Features.

Device to Aruba Central Communication



MOD 1- 39

Central uses a secure HTTPS connection and provides a strong mutual authentication mechanism based on certificates for all communication with Aruba devices. These certificates ensure the highest level of protection.

What are Central Groups and Sites

Groups

- Uses: Manage, configure, monitor, filter, report
- A device can belong to only one group at a time

Sites

- Uses: Monitor, filter, report
- Should follow physical location
- A device can belong to one site at a time



MOD 1- 40

There are three main mechanisms of controlling management and viewing functions: groups, sites, and labels. These were covered in the Aruba Campus Access Fundamentals course and HPE Press book. Here is a quick review of groups and sites.

Groups

In Aruba Central, groups are a primary management and configuration element: a container for device management, monitoring, and maintenance. Groups enable you to combine devices with identical configuration requirements and manage these devices efficiently by using either a UI-based configuration workflow or CLI-based configuration template or editor. Device groups enable you to:

Provision multiple devices in a single group. For example, a group can consist of multiple AP. These APs can share common configuration settings. You can also apply a common security policy for the devices deployed in a specific geographical location.

Provision different types of devices in a group. For example, a group can consist of APs, gateways, and switches.

Create a configuration base and add devices as necessary. When you assign a new device to a group, it inherits the configuration that is currently applied to the group.

Create a clone of an existing group. If you want to build a new group based on an existing group, you can create a clone of the group and customize it as per your network


requirements.

Sites

A site refers to a physical location where a set of devices are installed: a campus, branch, building, or venue. In Aruba Central, you use sites as a primary navigation element. Therefore, if your devices are deployed in a campus, you could create a site called CampusA. You can also tag the devices in CampusA using labels. If the campus consists of multiple buildings, you could label the devices deployed in the campus as Building1 or Lobby. If the devices in a specific location or an area within a specific location must have similar configuration, you can group them together.

Confidential - For Training Purposes Only

Aruba Central group



MOD 1- 41

A group's main purpose is to simplify the configuration. All devices in the same group will receive the same configuration. However, you can also use groups to monitor. When you select a group to monitor, you'll only see the devices in that group.

APs advertise SSIDs based on their group membership and configuration. APs will inherit the group configuration where they belong. A device can only belong to one group at a time. If you want all APs to advertise the same SSIDs, then you only need one group.

The slide shows an AP group named AP-Building-1. In this example, these APs must support three SSIDs: employee, voice, and guest. Thus, all APs are placed in the same group. Certain employee homes have a microbranch setup, so the scenario includes a Home-Office group. This group supports an SSID-named employee. This SSID configuration secure tunnels corporate traffic back to the office. The internet SSID configuration bridges traffic directly to the user's local LAN, and on to the internet. This example also includes a group called Switch-Building-1 for the Aruba switches' configuration in the building.

Group Types

UI group

Configure via a graphical user interface (GUI)

Template group

Configure with CLI-based templates



MOD 1- 42

There are two group types: UI and template groups. The type controls how you will manage devices placed in a group.

You configure UI group device members using a web interface (GUI), which allows you to customize and manage configuration parameters using UI workflows. For example, you can configure APs in a UI group using the menu options under Organization > Groups.

You manage template group device configuration with CLI-based configuration templates. You can combine devices with similar configuration requirements in a single template group and apply a common configuration template to these devices.

The default group is a system-defined container that consists of devices running factory default configurations. When a new device with factory default configuration connects to Aruba Central, it is automatically added to the default group. You can use the default group as a placeholder for new devices, with factory default configurations, allowing the network admin to find and provision those devices easily.


Why is my device marked as unprovisioned? If a device has a customized configuration and connects to Aruba Central, Aruba Central marks this device as unprovisioned. If you want to preserve the device configuration, create a new group and assign this device to the newly created group. If you want to overwrite the configuration, move the unprovisioned device to an existing group.

You will learn to manage groups, sites, and labels during the lab activities.

Confidential - For Training Purposes Only

Group Persona: Overview

- Represents the device's role in a network deployment
- Defines the type of devices and features to be available and managed by the group



- ✓ Customize configuration workflows
- ✓ Automate parts of configurations
- ✓ See default config and relevant settings
- ✓ Customize monitoring screens
- ✓ Troubleshoot workflows




MOD 1- 43

A device persona represents the device's role in a network deployment. The group persona defines the type of devices and features that are to be available and managed by the group. Creating personas for devices helps you to customize configuration workflows, automate parts of configurations, and see the default configuration and relevant settings for the device. It also helps you to customize the monitoring screens and troubleshoot workflows appropriate for the device.

Groups: Template and UI-based Configuration

Define a **preferred configuration method** during group creation



MOD 1- 44

Aruba Central allows you to manage configurations for different device types, like Aruba APs, gateways, and switches in your inventory. You can configure these devices using either UI workflows or configuration templates. You define a preferred configuration method when creating the group.

Aruba Central allows you to create a single group with different configuration methods defined for each device type. For example, you can create a group with the name Group1 and within this group, you can enable a template-based configuration method for switches and a UI-based configuration method for APs and gateways. Aruba Central identifies both groups under a single name (Group1). If a device type in the group is marked for the template-based configuration method, the group name is prefixed with TG, (TG Group1). You can use Group1 as the group ID for workflows such as user management, monitoring, reports, and audit trail.

After you assign devices to a group and access the configuration containers, Aruba Central automatically displays relevant configuration options based on the configuration method you defined for the device group. The Central Groups page allows you to create, edit, or delete a group, view the list of groups provisioned in Aruba Central, and assign devices to groups.

Managing Groups and Personas: Task Overview

Create group

Add a name, choose device type: APs, gateways, switches
UI-based config is default, or config via templates
Choose architecture: AOS 10 or AOS 8: Cannot be changed

Define group persona

AP: Campus/branch or MicroBranch
Gateways: Mobility or branch
Switches: AOS-CX only, AOS-S only, or both

Add device to group

Often from the *Unprovisioned* group to one of your defined groups
After move, device receives new group config



MOD 1- 45

There are three steps to managing groups and personas:

- Create and name the group
- Define the group persona
- Add a device(s) to the group

Create the group

Group creation starts with a group name and the device types to be members: access points, gateways, and switches. Then choose UI-based configuration (the default) or template-based configuration, and choose which device types in the group shall be managed using templates.

You can mix and match UI and templates in a single group. For example, a UI configuration method for access points and gateways and templates for switches. This adds flexibility for you to choose the type of management interface to use per device type within a group. Keep in mind that you cannot change the type of configuration method between UI and template once the group has been created.

Define the group

Then you define the group. Choose an architecture: ArubaOS 10 or Aruba OS8. You choose the group persona by choosing the architecture for each device category. These selections

affect the configuration options at the group and device levels. Also, you cannot change the group persona after you create the group. For example, if you selected ArubaOS 8 at the time of the creation, the only way to manage ArubaOS 10 devices is by creating a new group with the AOS 10 architecture. Note that you cannot mix the two OS firmware types within the same group!

Add devices to the group

After you define groups, you can of course move devices into those groups as members: perhaps devices from the unprovisioned group or default group to one of your defined groups. Be careful, though. After a device is moved to a new group, it reboots and inherits the new group's configuration. This means that the device will erase its current configuration and download a new one!

Confidential - For Training Purposes Only

Sites


Physical location where a set of devices are installed




Site – Sao Paulo
Site – San Jose
Site – Paris

Site Name	Address	Device Co...
All Devices		6
Unassigned		0
Paris	4 Rue Paul Lafargue	0
San Jose	6280 America Center Dr	0
Sao Paulo	alameda rio negro	6
Tokyo	2 Chome-2-1 Ojima	0


AP - Sao Paulo



AP+ switches - Paris



AP+ Gateway + switches - San Jose



MOD 1- 46

A site refers to a physical location where a set of devices are installed; for example, campus, branch, building, or venue. Sites are more than just a tool to filter views, place devices in a map and generate reports. You can use sites for several Aruba Central services such as live upgrades, VisualRF, AirMatch, install manager, network health dashboard, User Experience Insight (UXI), and etcetera. As a best practice, you should assign every device to a corresponding Aruba Central site. Devices can only belong to one site. Sites are also used for Clustering and Live Updates functionalities.

When creating a new site, you need the site name, street address, city, country, and region. The more information you add, the more accurately the site shows on the map. Most information is mandatory, but the zip/postal code is not.

Note: VisualRF currently does not support hierarchical arrangement for a location. If a campus has two or more buildings, it is recommended to create an Aruba Central site for each building.


Automation and REST API



You can really improve efficiency and consistency with automation and the REST API. This section introduces the REST API and the Network Analytics Engine (NAE).

Confidential - For Training

Network Analytics Engine (NAE) with AOS-CX




MOD 1- 48

The NAE engine is integrated with AOS-CX system configuration and time series databases. This enables you to examine historical trends and predict future problems due to scale, security, and performance bottlenecks. With that information, you can create software modules that automatically detect issues, take appropriate actions, and reduce the time spent on manual tasks.

NAE uses python-based agents, developed by Aruba and the Aruba community, to automate network monitor and troubleshoot functions. Agents monitor the state, settings, and statistics for system components in the current state database (CSDB)—perhaps the number of bytes transmitted on an interface or the switch's operating temperature. The agent then writes ongoing results to the time series database (TSDB) for ongoing analysis.

You can interact with the NAE through the AOS-CX switch's Web UI analytics dashboard. You can also automate the NAE through the REST API using tools such as Postman, Ansible, or another of your choice. You can create or obtain Python scripts for new agents and load the agents on the switch using the Web UI or the REST API. You can also deploy a pre-defined set of NAE scripts to AOS-CX switches using the Aruba Central Configuration UI, and custom NAE scripts can be pushed via group template configurations.

Aruba devHub

A screenshot of the Aruba devHub website. The header features the 'aruba' logo in orange and 'devHub' in dark blue. Navigation links include 'Get Started', 'APIs & Guides', 'Code Exchange', and 'Community'. The main visual is a cluster of hexagons; a central orange hexagon is labeled 'API', surrounded by several blue hexagons containing icons for various network functions like gear settings, a laptop, and a shield. To the right, the heading 'Master the game with the complete API reference' is followed by a paragraph: 'Take your knowledge to the next level. Build automated workflows and applications using Aruba products documentation and API guides.' Below this are five links with orange arrows: 'ArubaOS CX APIs & Guides', 'ArubaOS 8 APIs & Guides', 'Aruba Central APIs & Guides', 'ClearPass Policy Manager APIs & Guides', and 'User Experience Insight APIs & Guides'. The bottom right corner shows 'MOD 1-49'.

Aruba devHub is a place where you can learn how to integrate with Aruba products. Here you can find tutorials, documentation, and curated learning content to start a project:

<https://devhub.arubanetworks.com/home>.

You can take your knowledge to the next level and build automated workflows and applications using Aruba products documentation and API guides. You can also speed up developments with code samples and browse all available SDKs, packages, and code snippets posted by Aruba engineers, partners, and customers.



Knowledge Check

Self-check on key learning points



It's time for a knowledge check to review what you have learned.

Confidential - For Training Purposes Only

Question #1

You are leveraging Aruba's Edge Services Platform (ESP) in a new deployment. Which feature would you use to help identify issues before they affect users?

- a. Unified Infrastructure
- b. AIOps
- c. Zero Trust Security
- d. Cloud-native architecture

a

Knowledge Check

MOD 1- 51




b

Confidential - For Training Purposes Only

Question #2

Match the design layer to an appropriate description.



Knowledge Check



MOD 1- 52



Two-tier design	Small buildings with few wiring closets
	Large environment with fiber cables from each closet homed to a single location
Three-tier design	large network with thousands of users
	A more complex physical wiring layout

Question #3

Your colleague is confused about why you used a WLAN bridge mode deployment in some situations and tunnel mode deployments in other cases. Which statements below will improve your colleague's understanding?

- a. Both modes require you to run the same software version on APs and gateways
- b. Bridge mode deployments are great for simpler solutions that do not require advanced gateway features
- c. Tunnel mode deployments offer robust security features and greater operational flexibility
- d. Tunnel mode deployments increases visibility into applications

a

Knowledge Check

MOD 1- 53




b,c,d

Confidential - For Training Purposes Only

Question #4

Match the overlay method (centralized or distributed) to an appropriate description.



Knowledge Check ✓



MOD 1- 54



Centralized overlay User-based tunneling

Uses Generic Routing (GRE) encapsulation

Tunnels created from APs or switches to gateways

Distributed overlay Uses EVPN-VXLAN

Tunnels created between switches

Consistent, continuous policy enforcement

Question #5

Various endpoints may use different MAC addresses when communicating with the same AP radio, depending on which logical WLAN those clients are on.

- True
- False

a

Knowledge Check

MOD 1- 55



true

Confidential - For Training Purposes Only

Question #6

Which of the following statements accurately describe using Aruba Central groups and sites?

- a. Both groups and sites can help you to do monitor, filter, and report functions
- b. Both groups and sites can help you to manage and configure devices
- c. A device can only belong to one group and one site at a time
- d. You leverage both template-based and UI-based configurations on devices in the same group

a

Knowledge Check

MOD 1- 56



a, c, d

Note: Answer d is correct, since a single group can contain both switches and APs for example. Switches can be configured with templates, APs with UI.

Question #7

You have been directed to deploy an Instant AP (IAP) solution in a medium-sized office space. Which Aruba solutions would you leverage for this scenario?

- a. AOS 10 APs
- b. AOS 10 or AOS 8 APs
- c. AOS 8 APs
- d. Aruba gateways

a

Knowledge Check

MOD 1- 57



c

Confidential - For Training Purposes Only



Lab Activity

Let's review the lab activity that is associated with this module.

Confidential - For Training

Lab Tasks

Lab 01.01: Testing Lab Connectivity

Task 1.1

- Access the remote lab environment

Task 1.2

- Test connectivity to each device
- Access various lab devices

Module 2 – Building Wired Infrastructure



MOD 1- 59


The slide provides a brief review of tasks performed in this lab. Please see your Lab guide for details of this lab.

Building Wired Infrastructure

Implementing Aruba Campus Access

Welcome back everyone! This is Module 2—Building Wired Infrastructure. This module primarily focuses on VSX (it's operations, configuration and management, and split-brain issues), multi-area OSPF operation and configuration, virtual routing and forwarding (VRFs), policy-based routing, and the use of Aruba Central templates to manage AOS-CX switch groups.

Note that some of these topics were introduced in the Aruba Campus Access Fundamentals class, like VSX, and Central templates, but are covered in more depth in this course.



Objectives

- Deploy Aruba VSX to improve network resiliency and performance
- Deploy more advanced OSPF solutions to improve L3 performance and connectivity
- Use PBR for better load sharing, economy, QoS, and provider costing
- Create Aruba Central template groups to ease and automate switch configuration management

MOD 2- 2





This module offers more advanced techniques to really make your network sing and dance to YOUR tune. After completing this module, you will be able to:

- Deploy Aruba VSX to improve network resiliency and performance.
- Deploy more advanced OSPF solutions to improve L3 performance and connectivity.
- Use PBR for better load sharing, economy, QoS, and provider costing.
- Create Aruba Central template groups to ease and automate switch configuration management.

VSX

You can improve performance and resiliency by implementing VSX on your campus core and aggregation switches (6400 series and higher models).

Switch Virtualization Technologies





MOD 2- 4

The left example shown in the slide implements the Virtual Router Redundancy Protocol (VRRP) and the Spanning Tree Protocol (STP). This is an active/standby solution. This means that only one VRRP member actively processes traffic, and thus produces a sub-optimal solution. Of course, you can enable Multiple STP (MSTP) and configure multiple VRRP groups. Thus, one switch is a VIP master and the other is a VIP backup for a given VLAN. But this only load-shares between traffic from different VLANs (per VLAN load-sharing). What if you want switches to forward traffic and use all uplinks, even if the traffic comes from the same VLAN?

The solution is to use a switching virtualization technology and/or technologies that allows true active/active traffic processing, like Aruba's Virtual Switching Framework (VSF) and/or Virtual Switching eXtension (VSX). Depending on the technology, you can combine two or more switches to create one virtual switch. Now you can use link aggregation instead of STP/VRRP between the access and core/aggregation layers! This means that all uplinks can be used for load sharing—even for traffic from the same VLAN. For IP routing, both switches share the same IP address, and both can do L3 forwarding. You get true active/active L2 and L3 traffic processing!

AOS-CX Switch Virtualization: VSX vs VSF




MOD 2- 5

Both VSX and VSF allow active/active traffic forwarding. A significant difference is that VSX delivers high availability during software upgrades, with near zero downtime and continuous packet forwarding. This is due to the fact that a pair of VSX switches can run different version of the AOS-CX firmware, where this is not possible with VSF.

VSX is Aruba's core and aggregation layer technology while VSF is an access layer switch stacking feature. Let's dig a bit deeper.

Switch Virtualization Solution Comparison




MOD 2- 6

VSX virtualizes the control plane of two switches to function as one device at layer 2 and as independent devices at layer 3. From a data path perspective, each device does an independent forwarding lookup to decide how to handle the traffic. Some of the forwarding databases, such as the MAC forwarding database and neighbor tables, are synchronized between the two devices using a proprietary VSX control plane. Some of the forwarding databases are built independently by each switch. Configuration synchronization is one aspect of this VSX solution where the primary switch configuration is synced to the secondary switch using an opt-in process. With few exceptions, you must specify what you want synchronized between the VSX switch pair.

VSF allows multiple AOS-CX switches of the same model to act as a single virtual device. VSF helps to provide fast failover, scalability, manageability, and high availability. A VSF fabric has one active management plane and one active control plane, which the commander switch provides. The VSF commander proxies the control plane to the standby member. Both members of the fabric participate in the forwarding plane. Interface modules from both members are combined as if over one large switch connected by the VSF link.

VSF was covered in the Aruba Campus Access Fundamentals course. For now, let's focus on VSX.

VSX Benefits



Control plane

- Dual control plane: Better resiliency
- Unified management: Synced config, easy troubleshooting
- Independent software upgrades: Near zero downtime
- In-chassis (8400/6400) and device level redundancy

L2 distributed LAGs

- No spanning-tree requirement
- Loop-free L2 multi-pathing (active-active)
- Simple configuration
- Rapid failover

MOD 2-7


The slide shows an example of an Aruba-recommended topology, with VSX configured at the aggregation layer.

You can deploy VSX at the core if the topology requires it, but you may prefer to keep redundant core switches entirely independent. A dual control plane provides the best resiliency and software upgrades with near zero downtime. Meanwhile VSX permits unified management for features that you choose to synchronize.

The VSX pair uses link aggregations (LAGs) with links on both VSX switches to connect to access layer switches. This type of LAG is called a multi-chassis (MC-LAG), distributed, or VSX LAG. These VSX LAGs operate at layer 2 because the access layer switches typically are not configured to route, and the aggregation layer VSX pair acts as the default gateway for the access VLANs.

The layer 2 LAGs eliminate the need for a spanning tree protocol to run on the redundant links. They provide a loop-free path that permits the use of all the links in the aggregation (active-active design). Configuring a LAG-based topology is simple; and if a link fails, failover occurs very quickly.

VSX Benefits: Layer 3



L3 links

- Various options:
 - Routed only ports (ROPs)
 - L2 ports associated with dedicated SVIs
 - VSX LAGs associated with dedicated SVIs
- Unified data path: active-active L2/L3 forwarding
- L3 ECMP + L2 VSX LAG (highly fault tolerant)

Active-active gateway


- Active-active first hop gateway (VIP)
- No VRRP/HSRP
- Simple configuration
- No gateway protocol overhead
- DHCP relay redundancy

MOD 2- 8

The VSX pair has layer 3 links to the core. Each VSX switch can connect independently to each core switch on a routed-only port (ROP) or on a layer 2 port associated with a layer 3 VLAN interface dedicated to that port, also known as a Switch Virtual Interface (SVI). Alternatively, both VSX switches could connect to each core switch together on a VSX LAG. This LAG would operate at layer 2 and associate with one or more SVIs. Whichever option you select, VSX can support active-active layer 3 routing through Equal Cost Multipath (ECMP) and an efficient path for unicast traffic. With the final option, the combination of layer 3 ECMP and the VSX LAG creates a highly fault-tolerant system.


Remember that the VSX pair acts as the default gateway for the access layer VLANs. To do so, the VSX pair uses the active-active gateway feature. This feature allows each switch in the pair to act as an active default gateway for the VLAN using a shared virtual IP address (VIP) and virtual MAC address. It eliminates the need for Virtual Router Redundancy Protocol (VRRP) or Hot Standby Router Protocol (HSRP); and unlike VRRP, the active-active gateway feature allows both VSX switches to process traffic for the VLAN(s). Simple to configure, the active gateway feature relies on VSX operations. Thus, it adds no protocol overhead. It also supports redundancy for DHCP relay (IP helper) functions.

Data Plane Virtualization



Data plane virtualization supports VSX LAGs, so other LAG-connected devices perceive the VSX switch pair as a single device. In this example, those are both downstream access layer switches and upstream core switches.

Management Plane Synchronization



Database driven architecture

- Active-active components know peer state
- CLI/REST/WebUI exposes both control planes

Config and troubleshooting simplicity

- Continuous common config sync (opt-in)
- Show commands aggregate/contrast info from both switches
- Joint view of the VSX system

Hitless upgrade orchestration

- Proactively move traffic from the upgrading switch

Active-active analytics

- NAE agents cross monitor each others database
- Detection of discrepancies that remain for too long
- Validation that the overall solution is healthy

MOD 2- 10

VSX management plane synchronization is covered in the following sections.

Database-driven architecture

The VSX dual management plane supports synchronization, made possible partially through AOS-CX's innovative database-centric design. VSX peers can synchronize portions of their database so that each peer can actively forward traffic while also knowing the state of its peer.

VSX simplifies management functions: the CLI, REST interface, and web UI exposes both control planes to admins and business applications. In addition to simplifying management, this design enables analytics across the redundant pair.

Config and troubleshooting simplicity

You can choose which specific components to synchronize, such as an interface or a specific VLAN or an entire protocol. Configuration for that component is then continuously synchronized for easier configuration and fewer misconfigurations. show commands aggregate information from several configuration pieces across both switches and compare them, quickly picking out mismatched settings. You can add the vsx-peer option to many show commands, to view information about both the local switch and/or the peer: you get a simple joint view of the VSX system.

Hitless upgrade orchestration


VSX synchronization also helps to orchestrate hitless live upgrades. When one member needs to undergo an upgrade, the other member knows, and the pair can proactively move traffic away from the member that needs to be upgraded toward the other member.

Active-active analytics

While each switch in the VSX pair maintains its own NAE agent, which stores sensor values locally, NAE is aware of VSX. The agent on each member can connect to the other member's database and cross monitor it to detect discrepancies. For example, each NAE agent monitors the number of objects in the primary and secondary members' databases. If synchronization is working correctly, the number of objects should be identical. An alert can trigger if, for example, there are 20 percent more objects in the primary database than the secondary database for at least five minutes. Active-active analytics helps to validate that the overall solution is healthy.

Confidential - For Training Purposes Only

VSX Components Overview



MOD 2- 11


VSX supports multi-chassis link aggregation (MC-LAG): two or more links from two switches form a LAG that acts as a single logical interface. While the IEEE standard 802.3ad is limited to aggregating links on a single switch or device, VSX uses a proprietary technology to overcome this limitation to support link aggregation across multiple switches in the same VSX stack.

An Inter-Switch Link (ISL) connects the two switches, and a keep-alive mechanism is used for resiliency.

Also, VSX provides node-level network redundancy if when one switch fails. You must configure the downstream device as a standard 802.3ad LAG interface: the access switches perceive the VSX pair as a single device with a single system MAC address. Downstream devices can be anything that supports 802.3ad Link Aggregation Control Protocol (LACP), including switches, APs, firewalls, gateways, and other products. In VSX, one device acts as primary and the other device as secondary, but both devices forward traffic, for an active-active gateway solution.

And you will soon learn about the linkup delay feature, which enhances initial forwarding readiness.

VSX ISL



- Must **directly connect** the VSX peers
- Single link or **LAG** ← Strongly recommended: up to 8 physical links
- Use 10, **40 or 100 GbE** ← Preferred
- Span long distances: transceiver dependent
- ISL bandwidth should be equal or higher than speed of the uplink
- Apply typical QoS/ACL policies



MOD 2- 12

A VSX Inter-Switch Link (ISL) is a layer 2 interface between two VSX peer switches. You must configure each VSX switch with an ISL that directly connects to its peer VSX switch. The ISL is used at all three switch processing planes:

- Data plane: forward production traffic
- Control plane: exchange VSX protocol messages
- Management plane: synchronization


The ISL can consist of a single physical link, but Aruba strongly recommends link aggregation with two or more physical links. You can have up to eight physical links for resiliency, and all must be the same speed and type. You can use 10 Gbps, but are strongly encouraged to use 40 or 100 Gbps to ensure adequate bandwidth for the data plane traffic. The ISL uses front plane ports on the member switches and can use any media (but in the ISL link, all media types must match). If you select fiber media, the ISL can span long distances, depending on the fiber and transceiver type.

The ISL interface is a member of all VLANs by default. You can change the ISL membership through the CLI, but you must ensure that any VLAN that is carried on a VSX LAG is also carried on the ISL. In the data path, traffic is forwarded natively with no additional encapsulation, unlike with VSF. Also note that you can control the ISL with QoS and ACL

policies much as you do other interfaces on the switch. Of course, you must be careful that the ACL does not block required traffic. The ISL also preserves DSCP markings and you can optionally configure a VSX LAG to re-mark traffic with a specific DSCP.

Confidential - For Training Purposes Only

VSX ISL Resiliency




MOD 2- 13



Recall that the VSX switches use LAG links to connect to core and access switches. If a local VSX switch's link fails, it transmits the traffic across the ISL so its peer can forward it. The packet keeps the DSCP value that it would have had if transmitted over the LAG locally. However, the ISL interfaces on the other peer must be configured to trust the DSCP settings on ingress frames (this is not the default on AOS-CX switches, but you will learn how to configure QoS in Module 13).

This all works with the help of the Inter-Switch Link Protocol (ISLP).

VSX ISL Protocol



MOD 2- 14


The Inter-Switch Link Protocol (ISLP) runs over the ISL link. When you first configure the VSX pair, ISLP synchronizes the switches. Then it periodically synchronizes LACP states, MSTP states, MAC tables, ARP tables, and configuration information (the latter must typically be opted-in by you, the administrator). VSX also supports interoperability with Multiple Spanning Tree Protocol (MSTP), but not with RPVST+.

ISLP runs at layer 2. In other words, it is encapsulated in Ethernet frames. Since VSX switches are directly connected, there is no need for IP connectivity between the switches, by default. ISLP also helps peers to detect each other's status. They periodically exchange hello packet at one second intervals by default. The dead interval is 20 seconds by default. It determines when a switch decides that the peer is no longer available and initiates split detection. While the dead interval might seem long, understand that each VSX peer maintains its own control plane. If one peer goes down, all links in VSX LAGs go down. Traffic instantly fails over to the links on the other peer, which continues to actively forward without interruption. In most cases, peers determine that they must initiate split detection when the ISL link is down. The ISL link is considered down when all physical member links are down for the hold time, which is 0 seconds by default.

ISLP uses version control and provides backward compatibility regarding VSX synchronization capabilities. Processes used for VSX include `vsxd` for VSX LAGs, `vsx-synchd` for synchronization, and `vsx-mgmt` to synchronize REST, NAE, and configuration settings across the pair.

Keepalive: Protection from ISL Failure

- Identify and mitigate a split-brain situation →
 - VSX peer still up?
 - ISL link is down?
- If split-brain, force VSX LAG links down on secondary switch




You should configure each VSX switch with a keepalive connection to its peer. If the ISL goes down, each VSX peer uses keepalive communications to identify whether the peer is still up, but unavailable over the ISL—called split-brain condition. If the keepalive communications do indicate that the peer is still up, the primary VSX switch keeps its VSX LAG links up, and the secondary VSX switch forces its VSX LAG links down. The primary switch was selected when VSX established based on a user-configured role setting. If you configure the switches with the same role the lowest MAC address acts as a tie-breaker.

Peers exchange keepalives over a routed network; the path can consist of a direct layer 3 link or an indirect link through the upstream layer 3 network (the keepalive link should use a completely different interface than what it used for the ISL connection). You can set the source for keepalive packets to a loopback interface for stability. In either case, though, make sure that keepalives do not pass through the ISL.

Keepalive packets are UDP based. By default, they use port 7678, but the UDP port is configurable. By default, keepalive hello packets are sent every one second, but you can configure the interval to a value from 1 to 5 seconds. The keepalive dead-interval is 3 seconds, but has a configurable range of 2 to 20 seconds. If a device does not receive a keepalive packet from its peer within the dead interval, it treats the peer device as out-of-service and does not implement the split-brain protection mechanisms.

Keepalive: Connection Options



MOD 2- 16



You have two options for establishing the keepalive link. The most crucial rule is that the keepalive communications must not pass through the ISL.

- Use a directly connected, dedicated ROP (routed-only port) on each peer or a directly connected route LAG. Aruba recommends this type of configuration.
- Use the indirect routed path option. Keepalives pass through the upstream Layer 3 network. Aruba does not recommend passing the keepalives through the downstream network. The indirect option is cost-optimized because it requires fewer links.

You can choose an ROP, routed LAG, SVI, or loopback interface for the keepalive source. A loopback interface provides the most stability because the peers can always reach each other's loopback interfaces if even one Layer 3 path is available between each other. If you tie the keepalives to a particular physical interface or SVI, that interface might go down. Remember to advertise the source loopback interface in the dynamic routing protocol.

Note: For additional hardening, Aruba recommends assigning keepalive source interface to its own VRF, which is completely separate from your production and management VRFs. You'll learn more about VRFs later in this module.

Example Configuration for the Primary Switch

1. ISL interface

```
interface lag 128
no shutdown
no routing
vlan trunk native 1
vlan trunk allow x,y,x
lacp mode active
interface 1/4/28
no shutdown
lag 128
interface 1/4/32
no shutdown
lag 128
```

2. Set up the keepalive interface

```
interface 1/1/5
ip address 192.168.100.1/24
```

3. Set up VSX

```
vsx
role primary
system-mac 02:01:00:00:aa:bb
inter-switch-link lag 128
```

4. Creating a multi-chassis interface

```
interface lag 1 multi-chassis
no shutdown
no routing
vlan trunk native 1
vlan trunk allowed 11
lacp mode active
interface 1/1/1
no shutdown
lag 1
```



MOD 2- 17

This slide introduces the VSX configuration. First, set up the ISL link: a normal layer 2 LAG. When a native VLAN is defined, as shown here, the switch automatically executes the `vlan trunk allowed all` command to ensure that the default VLAN is allowed on the trunk. In this example, LAG 128 is being used as the ISL.

The same list of VLANs that are trunked on the access and/or core VSX LAGs must also be configured on the primary and secondary VSX switches in the global configuration. The list of VLANs can be synced to the secondary switch if the `vsx-sync` command is used in the VLAN context. Also verify that the VLAN set is also permitted on the ISL on the primary and secondary VSX switches. To configure VLAN trunking on the ISL, enter the `vlan trunk allowed [<VLAN-LIST> all]` command. If a native VLAN is defined, the switch automatically runs the `vlan trunk allowed all` command to ensure that the default VLAN is allowed on the trunk. To allow only specific VLANs on the trunk, enter the `vlan trunk allowed <VLAN-LIST>` command, for example: `vlan trunk allowed 2,3,4`. Then activate the LAG on the physical interface(s) between the two switches. Remember that to follow best practices, you should minimally use 2 links in the LAG.

Next, define the layer 3 interface that will be used for sending keepalives between the two switches. This can be a direct link between the switches. Each switch will need unique IP addressing.

After defining the keepalive interface, you need to set up VSX. Enter the VSX context and

define the role: primary or secondary. Next, specify the LAG that will be used for the ISL link. In the above example, this is LAG 128. That is all that is necessary to initially set up VSX. For detailed information on configuring, tuning, and troubleshooting VSX, see the Virtual Switching Extension (VSX) Guide. In the example, a system MAC address was also defined.


To set up a multi-chassis LAG between this VSX pair and another device, define a new LAG, but specify the multi-chassis parameter: this parameter tells the VSX pair that there is at least one link on each switch connecting to a neighboring device, like a switch or gateway, for example.

Confidential - For Training Purposes Only

Periodic Synchronization

All entries learned by both peers

- Even if all traffic flows through one switch, peer learns by syncing
- Sync has no functional impact to normal data-path learning



VSX split and rejoin scenario - periodic sync resumes after the bulk sync. **Exceptions:**

- You clear ARP entries or shut down/re-enable an SVI on 1 VSX switch, entries out of sync
- **Best practice: Perform every operation on both VSX switches**

clear arp • interface vlan, then shut/no shut a VLAN

MOD 2- 18

Every second, each VSX switch synchronizes its learned MAC addresses with its VSX peer using ISLP. Each VSX switch also buffers and sends its new neighbor ARP/ND entries to the VSX peer every time its buffer is full (64 entries) or when it hits a timeout of 5 seconds. In a VSX scenario where all traffic from the core to the access layer flows through one VSX switch only, the other VSX switch still learns the necessary ARP/ND entries from its VSX peer. VSX has no functional impact on the VSX switches' normal datapath based learning.

If a VSX split occurs and the switches then rejoin together, the switches perform a bulk sync; after that, periodic syncs resume.

Inter-VRF route leaking (IVRL) leaks neighbor entries from a source VRF to a destination VRF. This process creates two neighbor entries, one learned through the datapath in the source VRF and one that was induced in the destination VRF. The VSX switches do not sync the induced entries either during the initial ARP sync or the periodic syncs. Instead, each switch learns the ARP entries in the source VRF (either through the normal datapath process or through the ARP sync). The local IVRL process then induces the entries into the destination VRF on each VSX switch individually.

Keep one caveat about ARP syncing in mind: if an administrator clears ARP entries or enables and disables a VLAN interface on one VSX switch, but not on the other, ARP entries will go out of sync between the switches. The administrator must perform such operations on both VSX switches so that both devices will re-learn entries and synchronize correctly.

Therefore, if you enter one or more of the following commands on one VSX switch, but not on the other VSX switch, the ARP entries on both switches will become unsynchronized.


clear arp

interface vlan

- shutdown for a VLAN
- no shutdown for a VLAN

Confidential - For Training Purposes Only

Initial Sync and Linkup Delay



MOD 2- 19

After a VSX switch reboots, it has no entries for ARP, MAC, and routes. If downstream VSX LAG ports activate before all this information is relearned, the traffic is dropped. To avoid a traffic drop, VSX LAGs on the rebooted switch stay down until the restoration of LACP, MAC, ARP databases, and MSTP states, if MSTP is used.

The learning process for the VSX LAGs has two phases:


- Initial sync phase: The LACP states, MAC address table, ARP table, and potentially MSTP states download from the forwarding switch to the freshly rebooted switch.
- Link-up delay phase: The downloaded entries are installed into the ASIC. Router adjacencies with core nodes and learned upstream routes are also established. You can configure the link-up delay phase with the linkup-delay-timer command. The default value is 180 seconds. Set the link-up delay timer to the maximum value of 600 seconds for a network with many MAC addresses, a large ARP table, or a large routing table.

When both VSX switches reboot, the link-up delay timer is not used because both switches must relearn the LACP states, MAC address table, and ARP table. To get upstream router adjacencies established during the link-up delay, you must exclude the upstream LAGs (LAG101, say) from the scope of the link-up delay. Run the linkup-deal-timer exclude lag-list <LAG-LIST> command for identifying the LAGs for exclusion. Until linkup delay timer expires, all SVIs associated with VSX LAGs are kept in a pseudo-shut state.

Let's get more detail on this two-phase learning process.

Confidential - For Training Purposes Only

Initial Sync and Linkup Delay Phases



Initial sync phase

- Rebooted node learns all LACP+MAC+ARP+MSTP entries from its VSX peer via ISLP
- Initial sync timer is not configurable

Linkup delay phase

- Time required to:
 - Install downloaded entries to ASICs
 - Establish router adjacencies with core
 - Learn upstream routes
- Linkup delay timer
 - Default: 180s
 - Set up to 600s, based on network size, ARP/routing table size


MOD 2- 20

The slide summarizes the two phases of the linkup delay.

The initial sync phase is the download phase, in which the rebooted node learns all the LACP, MAC, and ARP database entries from its VSX peer through ISLP. The initial sync timer is not configurable. It is the required time to download database information from the peer.

The linkup delay phase should last for as long as is required to install the downloaded entries to the ASIC, to establish router adjacencies with core devices, and to learn upstream routes. During this phase, links are forced down. The linkup delay timer's default value is 180 seconds. Depending on the network size, ARP table size, and routing table size, the timer may need to be set to a higher value, where the maximum timer is 600 seconds.

Configuration Synchronization



Most configs same across VSX pair

- Exceptions: IP addresses, router IDs, hostnames, etc.
- Global VSX configuration sync available

These features support synchronization

Individual features

- VLANs
- ACL, QOS, PBR
- Object-group class
- Policies
- Rate-limits
- Active-gateway

Group features

- AAA, SNMP
- COPP_policy
- DNS, SSH, Time, sflow
- vxslag_interfaces
- qos_global
- static_routes


MOD 2- 21

Data plane sync and config-sync are two different things. Data plane sync was covered earlier. Let's dive into config-sync (the vsx-sync parameter). You can use the vsx-sync parameter from the CLI or NetEdit to ensure that the configurations are synced between the VSX peers.

The two VSX switches should have nearly the same configuration with a few exceptions. These exceptions include host IP addresses on layer 3 interfaces, router IDs, hostnames, and a few other settings. Configuration synchronization makes it simpler to ensure identical configs. By default, VSX configuration synchronization is enabled along with VSX. However, enabling synchronization globally has no effect until you select individual configuration items on which to enable synchronization. The vsx-sync attribute enables synchronization on an item. You can disable configuration synchronization globally, in which case you can no longer apply the vsx-sync attribute to any configuration items.

The slide shows individual and group features supported by vsx-sync.

Configuration Synchronization Behavior



- Primary switch config auto-pushed to secondary
- If secondary down, config pushed when it comes up
- Config sync fails show up in a CLI status command
 - You must fix problem manually
 - No fail-safe that auto-disables VSX if sync fails


MOD 2- 22

You must set the `vsx-sync` attribute on the desired configuration items from the VSX primary switch. The primary switch will automatically push this configuration to the secondary switch. If the secondary switch is not available when you configure the primary, the configuration is pushed to the secondary switch after the secondary switch boots up. If a certain config sync operation fails, the CLI status commands will indicate the failure. You must manually fix the inconsistency. There is no fail-safe mechanism that disables VSX if configuration synchronization fails.

You can configure `vsx-sync` on a secondary switch if `config-sync` is disabled or if the ISL is down.

Important: The `vsx-sync` attribute uses this opt-in model due to customers' and partners requirements and needs. In other words, with few exceptions, no commands or configurations are synchronized on the two switches unless you explicitly tell the VSX primary to do synchronization for a configuration component.

Configuration Overlap and Protection



Config element is set on Secondary without being synced/pushed from Primary

No problem - valid scenario

Same configuration item on Primary is now vsx-synced

Primary overrides previously set config element on Secondary

vsx-sync removed from config element on Primary

Config item stays on Secondary - without vsx-sync keyword

MOD 2- 23

It is perfectly valid to configure a setting manually on the secondary switch without that setting being synced from the primary switch. However, if vsx-sync is now enabled for this feature on the primary switch, the setting on the primary switch overrides the previous setting on the secondary switch. For example, you could add VLAN 10 on the secondary switch with description "employees." You then add VLAN 10 with no description on the primary switch and enable vsx-sync. The secondary switch now has VLAN 10, but with no description.

Firmware Upgrade

During firmware upgrade, software versions are temporarily mismatched



vsx-syncd process:

- Stops any synchronization attempt
- On-going check to determine if it can sync
- Detects when peers are on same version
- Restarts on-going sync



Graceful upgrade for routing features

Prior to upgrade, withdraws from route advertisements, traffic shifts to alternate path

MOD 2- 24

During a firmware upgrade, software versions will temporarily mismatch on the VSX peers. Therefore, the vsx-syncd process, which is still running, stops any synchronization attempts. It then performs an ongoing check to determine if it can sync. Once both VSX peers are running the same firmware, this ongoing check sees that it can restore the synchronization. The vsx-syncd process then continues the on-going sync. The database connection between the peers works if the firmware versions are the same.


As of AOS-CX 10.2, VSX also supports graceful upgrade for routing features. When a member is about to be updated, it takes action such as withdrawing itself from route advertisements, so traffic is shifted away from it during the upgrade process.

VSX LAG

Combines two or more links across two switches in a LAG, acts as a single logical interface

- Appears as one device to a LAG partner
- Ports must be same speed and type
- Up to 4 physical links per peer for 8 links total
- Requires LACP L2 LAG only

- Sync LAG states between VSX peers via ISLP
- Always prefer local VSX LAG links to switch traffic



MOD 2- 25

A VSX LAG allows you to connect a switch or another networking device to both physical peers in a VSX fabric. In the slide, switch Access-1 connects to both Agg-1 and Agg-2 of the VSX pair and perceives a standard LAG link: a single logical interface to a single logical switch. The same holds true for Access-2 and the core switches.


The two VSX switches appear as a single peer ID to the upstream or downstream switches. Ports must have the same speed. The LAG can include up to four physical links per peer for eight links maximum per VSX LAG. As of the publication of this course, VSX LAGs are layer 2 only, and VSX LAGs must be LACP-based. A non-LACP or static VSX LAG is not supported.

The VSX peer switches synchronize their LAG states using ISLP. You can configure a lag-hash scheme, which dictates how the switches load-share traffic across the links in the LAG. (Note: The 8400 provides L4-based hashing, but the 8320 does not.) VSX LAGs are locally optimized, meaning if a switch with local links is in a LAG then it constrains the LAG hash-scheme to look at only local links, and it will only forward traffic over the ISL for transmission on the peer links in the LAG if local links of a VSX LAG are down.

The number of VSX LAGs supported is dependent on the switch model: 255 on 8400, 53 on 8320, 55 on 8325.

Note: Originally, Aruba used the term multi-chassis LAG (MC-LAG) to describe a LAG between an external device and a VSX switch pair. Today, Aruba uses the term VSX LAG₁₄

Active Gateway Overview



Active-active Layer 3 default gateway

- Eliminates single point of failure
- Shared VIP and VMAC on SVI
- Permits active-active forwarding
- Eliminates need for VRRP
- Switch interface MAC is source for routed traffic

Active gateway is *highly* recommended for first-hop redundancy protocol



MOD 2- 26

Active gateway is a first-hop redundancy protocol that eliminates a single point of failure for default gateway services on an access network. The active gateway feature improves the reliability and performance of the host network by enabling a virtual router to act as the default gateway for that network.

To configure this feature, you simply configure a shared virtual IP address (VIP) and a shared virtual MAC address (VMAC) on an SVI on both VSX switches. The VIP/VMAC serves as the default gateway for each access layer VLAN. Which peer in the VSX pair receives a particular piece of traffic from an access VLAN depends on which link in the LAG the downstream access switch's LAG has selected. Whichever peer receives the traffic will then forward it onwards.

When you use the active gateway feature on a VLAN, the VLAN has no need for VRRP; the features are mutually exclusive. As with VRRP, routed traffic from the VSX peer is sourced from the switch interface MAC, not the VMAC address. Every three minutes (non-configurable), each active gateway broadcasts a hello packet from the VMAC to avoid VMAC aging on the access switches. The hello from both peers transmits on the same VSX LAG, so no issues occur with the same MAC address appearing on different ports. The hello is a non-IP Ethernet packet using HP EtherType, 080009 Hewlett-Packard, and a destination reserved multicast address.


The VSX pair supports up to 16 different VMACs, which are not shareable between IPv4 and

IPv6. For example, the pair could use eight VMACs for IPv4 simultaneously with eight VMACs for IPv6.

You should understand some key differences between VRRP and the active gateway feature. VRRP requires multiple configuration options, including the VRID on the VLAN, the roles on each VRID, the virtual IP address on each VRID, the advertisement timer configurations, and so on. Active gateway only requires a single line of configuration per VLAN interface: the configuration of the virtual IP/MAC addresses. VRRP and active gateway also differ in terms of how they forward traffic. VRRP uses an active-standby data plane with only the VRRP master routing traffic. With active gateway, both VSX switches forward traffic.

Confidential - For Training Purposes Only

Active Gateway: South-North Unicast Traffic



Normal operation, no traffic over ISL due

- VSX LAG local link optimized usage
- Active gateway




MOD 2- 27

If you configure VSX LAGs with links on each peer, the normal operation unicast traffic sent between the access layer and the core (south-north) uses all access-to-aggregation links and does not flow over the ISL. From a Layer 2 perspective, VSX LAGs span two switches and operate in active-active mode. Traffic between the access layer and aggregation layer switches can be forwarded to any of the active links. There are no loops and no need for spanning tree protocol or blocked ports.

From a datapath perspective, local link optimized usage ensures that when a VSX switch receives traffic, it always uses its local LAG links to forward the traffic to its destination. The VSX switch only uses the ISL link if the local LAG links are down. An active gateway ensures that the first VSX switch that receives traffic from the access layer can route the traffic across to the layer 3 domain.

VSX LAG and Upstream Unicast Routing



Both VSX peers run independent control planes

- Separate OSPF/BGP processes
- Unique router IDs

VSX peers in the data path

- Function as a single router
- Active-active forwarding

Design options

- L2 or L3 links with upstream core nodes?
- Static, OSPF, BGP?
- Single VRF / multiple VRF?
- Sizing / limitations?
- Best practice for high availability

MOD 2- 28


You will now examine how the VSX pair operates on the connections to upstream core switches, which are generally routed connections. Remember that each VSX peer runs an independent control plane with its own OSPF or BGP processes. Therefore, each VSX peer appears as a different router with its own router ID to the core switches. However, as far as the data path is concerned, the VSX peers function as a single router and support active-active forwarding. This is possible because a core switch uses Equal Cost Multipath (ECMP) routing to learn routes to the subnets behind the VSX pair through both VSX peers and because the VSX pair supports a feature called active forwarding, about which you will soon learn. Also remember that VSX LAGs can only operate at layer 2.

Keeping these principles in mind, you can consider the options that you have for designing the upstream VSX LAGs and unicast routing. You could choose to connect with routed ports or layer 2 ports associated with layer 3 VLAN interfaces. Or you can use layer 2 VSX LAGs, which are associated with layer 3 VLAN interfaces. For routing, you could choose static routing, OSPF, or BGP. You could use a single VRF or multiple VRFs, depending on the requirements of the scenario. Your choices will be informed by sizing limitations and by best practices for high availability.

Upstream Connectivity: ROP

ROP, single VRF

- Each of the 4 links has its own /30 or /31 subnet
- OSPF network-type point-to-point




MOD 2- 29

There are several options for links between the aggregation VSX pair and the core: one involves a single VRF implementation while the other options involve multiple VRFs.

The first option uses route-only ports (ROP). It could also use routed LAGs, but those routed LAGs would only have links on the local switch: they cannot be multi-chassis. Each VSX switch has its own routed link or LAG to each core switch, and each of these four links is its own /30 or /31 subnet. In OSPF, this would be a point-to-point network. This simple VLAN-free configuration is best suited for a network that runs on a single VRF domain. If the network had multiple VRFs, each VRF would need its own ROP, which is less feasible.

The other options are beyond the scope of this course.

VSX and Upstream Routing: Supported Scenarios




MOD 2- 30

The following chart summarizes the supported scenarios for the upstream connections between a VSX pair and the core. When you have a single VRF, ROPs provide the simplest option. However, layer 2 ports associated with dedicated SVIs, as well as layer 2 VSX LAGs associated with dedicated SVIs, work as well.

All of these options work with static routing, OSPF, or BGP.

Split-Brain Scenario



ISL is DOWN!

- Both switches still active, but cannot exchange info
- Switches become out of sync – split brain

Keepalive prevents split brain

- OOB mechanism detects whether link or switch is down
- If link down, disable Agg2-access switch links
- Traffic forced to use Agg-1 only

MOD 2- 31

The diagram in the slide shows a situation in which the ISL between the two VSX switches is down. Both switches are still active, but the switches cannot exchange information, which causes the switches to become out of sync. This situation is called a split brain, and it occurs if the keepalive function is not enabled. The keepalive function prevents split brain issues by providing an out-of-band mechanism to detect if the link between the VSX switch is down or if the VSX switches are not working. The keepalive functionality brings down the link between Agg2 and Access Switches. The traffic is forced to go from access switches through Agg1 only.

Split-Brain Scenario

- VSX LAGs on secondary are disabled to prevent "split-brain"
- When ISL is restored, there is no reboot of secondary node
- Secondary VSX LAGs are brought up after linkup-delay timer


	Keepalive-Established	Keepalive-Init
ISL "In-Sync"	Forwarding OK Protection OK	Forwarding OK Protection NO
	STATE • ISL is down • Peer is UP	STATE • Peer is DOWN
ISL "Out-of-Sync"	ISSUE • VSX tables will get out of sync	ISSUE • Forwarding at 50% capacity
	ACTION • On secondary: stop forwarding over VSX-LAGs	ACTION • Keep forwarding

MOD 2- 32

The table in the slide shows what happens with various ISL and keepalive statuses on the VSX pair, as described below.

- If the ISL is up and the peers are in sync: VSX pair forwards traffic normally
- If the keepalive connection is up, system is protected from split-brain condition
- If the keepalive connection is down, but ISL is fine: VSX pair can forwards traffic normally, but system is at risk. (Protection No in the table)
- If the ISL goes down and peers are out of sync, but keepalive is up: each switch knows that their peer is still up. They know that they are at risk of becoming out of sync, so the secondary switch shuts down its VSX LAGs.
- If the ISL goes down and keepalive is down: The switch assumes that its peer is down and that the system is functional but forwarding at 50 percent capacity. It keeps forwarding even if it is the secondary peer. It is important that you form the keepalive connection correctly so that this assumption is correct.

ISL Failure: SVI Impact



ISL failure, but keepalive established

- SVIs associated to any VSX-LAGs are shut down -> VSX-LAG interfaces are brought down
- VLAN 20 still UP with external connectivity
- Server 3 port up, but VLAN SVI 10 is down due to ISL failure
- Single-homed systems connected to SVI supported on MC-LAG should connect to primary VSX member

Server 3 not reachable as SVI 10 is **DOWN**
Server 3 should connect to Agg-1

Always plan for redundant connections

MOD 2-33

An ISL failure has no effect on the up/down state of VLANs and associated SVIs that are not part of any VSX LAG but part of at least one orphan port (beside the ISL LAG, which is not a VSX LAG).

During ISL failure (before initial sync)

If the secondary VSX node has a port that is a member of a VSX LAG, the associated SVI of the VLAN—transported by said VSX LAG—has the VSX secondary node shut down, whether or not there are orphan ports carrying that given VLAN.

During link-up delay timer (after initial sync)

As long as the secondary VSX node has a port that is a member of a VSX LAG, the associated SVI of the VLAN—transported by said VSX LAG—is kept shut down on the VSX secondary node, whether there is an orphan port(s) carrying that given VLAN or not. The associated SVI of the VLAN transported by a VSX LAG restores ON/UP on the VSX secondary node, only if the following conditions are met:

You have used the command `link-up delay timer exclude lag` to exclude the said VSX LAG.

No VSX LAG not part of the exclusion set allows the given VLAN.

In the topology in the slide above, after an ISL failure, ports and SVI associated to a VSX¹²³

LAG are automatically shut down, ports and SVIs not related to and VSX LAG are still operating. In this case, server 2 and VLAN 20 are still up and running and have external connectivity. The server's 3 port remains UP, but VLAN 10 SVI is shutdown due to ISL failure. This also means that VLAN 10 subnet will no longer be advertised in a routing protocol by Agg-2. As a best practice, always connect single-homed system to primary VSX member.

Confidential - For Training Purposes Only



Lab Activity

Let's review the lab activity that is associated with this part of the module.

Confidential - For Training

Lab Tasks

Lab 02.01 - Campus Wired Aggregation

Task 1

- Review initial configuration

Task 2

- VSX Basic Configuration

Task 3

- Configure VSX LAG

Task 4

- Configure VSX L3 SVI with Active Gateway

Task 5

- VSX Link Up delay

Task 6

- VSX Split-brain detection

Lab 02.02



MOD 2- 35

The slide provides a brief review of tasks performed in this lab. Please see your Lab guide for details of this lab.

OSPF



OSPF is the most commonly-used internal routing protocol. To be the hero you know you can be, you need strong OSPF skills. Please note that a single area OSPF configuration was covered in the Aruba Campus Access Fundamentals course. This section will provide a quick review of single area designs, but will primarily focus on multi-area designs and external networks.

OSPF Review

- Dynamic routing protocol for complex, redundant topologies
- IGP - used within an organization or AS
- Link state routing protocol →
 - Advertise links
 - Calculate lowest cost route based on topology
 - Communicate updates and converge quickly
 - Uses Dijkstra (SPF) algorithm
- Hierarchical, area-based design



MOD 2- 37


Many companies have multiple routers and routing switches that send traffic across the campus. Dynamic routing protocols enable routers to discover and update routes automatically. As an Interior Gateway Protocol (IGP), OSPF enables route info exchange among routers in the same autonomous system (AS): a system under a single organization's control, like your company.

OSPF is a link state protocol, where routers advertise information about their connected layer 3 interfaces, or links, along with the cost (or metric) associated with that link. These link state advertisements (LSAs) propagate across participating routers, which use the advertisements to build a network topology—a map of the network—all known links, the routers connected to them, and associated costs. Using the topology, routers calculate their own lowest cost route to each destination subnet.

Link state protocols like OSPF tend to choose better routing paths and to converge faster than distance vector protocols like Routing Information Protocol (RIP). OSPF uses the Dijkstra's algorithm, also called the shortest-path first (SPF) algorithm to find the least cost, loop-free paths. For scalability, OSPF supports a two-layer hierarchical design: a backbone area and other areas connected to the backbone.

OSPF Overview

- 1 Create routing process
- 2 Define RID: Loopback IP recommended
- 3 Create an area
- 4 Enable OSPF on L3 links, place in area



MOD 2- 38

Let's review the basic steps for setting up OSPF. Of course, configuring OSPF to meet your organization's needs has many more nuances, many of which are covered in this module. However, these basic steps can get your AOS-CX switches routing traffic across the network:

Create an OSPF routing process – an arbitrary number that has local significance only. Although they need not match between routers, it's best to do so for consistency. Most folks simply set it to 1.

Optionally define a router ID (highly recommended). Unlike the routing process, this must be unique for each router in your OSPF system. For stability, use a loopback interface. If you use a physical interface as the router ID, and that link fails, the router ID changes – which is sub-optimal.

Create an area.

- OSPF uses areas for scalability—to segment the network into smaller pieces and to enable route aggregation. Your initial configuration will use one area, area 1. You will then learn about multiple area solutions.

Enable OSPF on layer 3 interfaces. For AOS-CX switches, these are either routed or VLAN interfaces. Enabling OSPF on an interface has two effects:

- The routing switch sends OSPF messages on the interface and can form neighbor relationships with other routers on the same interface. Neighbors that fully exchange their

topology information are called fully adjacent.

- The router advertises this interface in OSPF. For example, if the interface is associated with IP address 10.255.0.1/24, the router advertises that it has a link to 10.255.0.0/24.

VLANs that need to be advertised in OSPF include:

Those that connect routers (like VLAN 255 on Access-1 and Core-1)


Those that support endpoints (like VLAN 10 on Core-1)

If you are only enabling OSPF on the VLAN because you want to advertise its subnet, and you do not want the router to form adjacencies with other routers on the VLAN, configure the VLAN as a passive OSPF interface.

When you enable OSPF on a VLAN, you also specify the VLAN's OSPF area. The area ID must match on all OSPF routers connected on the same VLAN or subnet.

Confidential - For Training Purposes Only

OSPF Network Types



MOD 2- 39

OSPF routers run the protocol on layer 3 interfaces, or networks. This includes VLAN interfaces, physical interfaces set to routed mode, link aggregations set to routed mode, and loopback interfaces. AOS-CX switches support two types of networks: broadcast and point-to-point. Ethernet is a broadcast protocol, and a VLAN is a common example of a broadcast network.

You can connect OSPF routers to the VLAN in some sort of redundant tree topology, directly to each other, or through Layer 2 switches. The figure in the slide shows an example of a layer 3 connection with a circle, to which all routers connect. This “circle” is likely a layer 2 switch (the layer 2 connections are not shown in the diagram). In a broadcast network, a designated router (DR) is elected to establish adjacencies with other routers on the link (subnet). DRs also receive and flood link state updates (LSUs). A backup DR (BDR) is also elected to take over if the DR fails. OSPF routers multicast messages to the DR and BDR at 224.0.0.6. The purpose of DR/BDR election is efficiency. Non-DR routers need not communicate with several other routers on the link – only the DR/BDR.

When you directly connect two OSPF routing switches, it is generally preferable to set up a point-to-point network. This network is typically associated with a /30 or /31 subnet. On AOS-CX switches, GRE tunnel interfaces are point-to-point by default, and loopback interfaces are a special loopback type. It could also include VLAN interfaces, but only the two routers belong to the associated VLAN. But most interfaces are broadcast networks by default, including physical route-only ports and VLAN (SVI) interfaces. However, you can set route only ports

and link aggregations to point-to-point, assuming that the routed link connects directly to a routed link on another OSPF device. You can even set VLAN interfaces to point-to-point if they follow the rules. (Remember, though, that a VSX LAG between a VSX pair running OSPF and another OSPF router should be a broadcast network because the network has three routers: each VSX switch and the connected router.) Since a point-to-point network has only two routers, it does not require the DR and BDR roles.

OSPF also distinguishes between transit and stub networks. A transit network is any network on which a router discovers at least one neighbor; it could be a broadcast or point-to-point network. A stub network is any network on which the router does not detect another OSPF device. The reason that the router runs OSPF on the stub network is to advertise the network to other routers in the AS. Often a stub network is a VLAN interface for which the OSPF router is the default gateway.

Note: Two routers (even a VSX core) on same subnet, with passive interfaces, also results in stub network.

You can define the passive option for interfaces that will not connect to other OSPF peers, and/or to ensure that the interface will not form an OSPF adjacency should any OSPF neighbors be connected to that interface. This can slightly reduce OSPF overhead on the link, and improves security by ensuring that rogue routers will not connect to your system - especially if used along with OSPF passwords.

Note: OSPF point-to-point networks converge faster than OSPF broadcast networks and therefore are more preferred when applicable.

OSPF LSA Introduction

Type	Name	Advertised by	Scope of advert	Included information
1	Router	All routers	Area	Links (ID, type, and metric)
2	Network	DRs for broadcast (or NBMA) networks	Area	DR IP, network mask IDs of routers on the network
3	Summary	ABRs	Area	Network in another area (network IP, mask, and metric)
4	ASBR	ABRs	Area	ASBR router ID and metric
5	AS External	ASBRs	All normal areas in the AS	External network, mask, metric, metric type
7	NSSA External	ASBRs	Area	External network, mask, metric, metric type
9	Opaque Link Local	All capable OSPF routers	Link	Extensions (graceful restart)
10	Opaque area-local	All capable OSPF routers	Area	Extensions (traffic engineering)
11	Opaque AS	Capable ASBRs	Non-stub areas	Extensions




MOD 2- 40

The table in the slide summarizes key information for various LSA types. This table does not show Type 6 LSAs because they advertise multicast groups for Multicast OSPF (MOSPF), a protocol that routers rarely use and that AOS-CX switches do not support.

Type 1 LSAs are router advertisements: the router advertises information about itself. Type 2 LSAs are Network advertisement: routers advertise their connected subnets. You will learn about LSA types 3 to 7 in the rest of this module.

Other LSA types include extended or opaque LSAs, which enable features such as OSPF graceful restart and traffic engineering. Routers only advertise opaque LSAs to neighbors that indicate that they are opaque capable during adjacency establishment. You will learn about graceful restart later in this module. Other opaque LSAs are beyond the scope of this course. Not included in the table is the Type 8 LSA, which has special functions in OSPFv3, used for IPv6 routing. OSPFv3 is also beyond the scope of this course.

LSA Types and Usage




MOD 2- 41

The figure in the slide illustrates where you'll find the common LSA types and who advertises them.

You see Type 1 and 2 LSAs inside an area. ABRs inject Type 3 summary LSAs into an area. ASBRs send type 5 external LSAs, and sometimes Type 7 LSAs, while Type 4 LSAs are information about those ASBRs.

Single OSPF Area: Limited Scalability



- Every router must maintain LSAs from every other router in the area in its LSDB
- LSDBs become large. SPF algorithm requires ever-more processing overhead
- No benefit: Segment 1 routers only have 2 paths to reach Segment 2 – via Core-1 and Core-2
- Stability: Flapping interface in Segment 1 affects every router in the entire area



MOD 2- 42

The figure in the slide shows a single area OSPF network with two segments. The segments might be a campus LAN and a data center, a campus LAN and a large branch office, buildings in a large three tier-campus topology, or any other segments within a larger network. Segment 1 and Segment 2 each represent many routers, which ultimately connect to two core routing switches. These core routers connect through a network backbone, which might also include more routers and routing switches. The figure only shows two segments for simplicity, but a large enterprise network might have several such segments.

Every OSPF router in an area must maintain LSAs from every other router in the area in its LSDB. The top example shows how a single area implementation forces each of these Type 1 and 2 LSAs to propagate across the system. Plus, the figure only mentions LSAs associated with segment 2. Each OSPF router must also maintain LSAs for routers and networks in segment 1 and the backbone.

As a domain grows, the LSDB becomes excessively large and difficult to process. Routers must dedicate valuable processing resources to running the SPF algorithm. Often this resource consumption adds little value. As shown in the figure, networks often feature aggregation such that one segment connects to another through a few core devices. With only two exit points from segment 1 towards segment 2, the routing switches in segment 1 have no reason to maintain so much granular information about networks in segment 2.


In addition, suppose that two routing switches in segment 2 provide a connection to the

subnet A. One of those switches has a flapping interface. It loses its connection to subnet A and sends an LSU withdrawing that subnet from its list of links. Every routing switch in the network backbone and segment 1 receives this LSU, adjusts the LSDB, and re-runs the Dijkstra algorithm on the topology. However, these routers still reach subnet A through the same core routing switches. When the flapping interface comes back up, the switch sends another LSU, which propagates throughout the entire network, causing every OSPF router to recalculate shortest paths, again without any significant change.

What is the solution?

Confidential - For Training Purposes Only

Multi-Area OSPF: Large Scalability



- Aggregate inter-area routes
- Much smaller LSDBs
- Reduce LSA/LSU flooding
- Minimize SPF recalculations
- Increase stability and control



MOD 2- 43

OSPF areas segment the domain into more stable and manageable pieces. Each OSPF routing device maintains the LSDB and runs the shortest path algorithm only for its own areas. OSPF ABRs forward routing information between areas. Type 1 and Type 2 LSAs do not flow across area boundaries. ABRs summarize this huge number of Type 1 and Type 2 LSAs into a few Type 3 summary LSAs.

This all serves to dramatically decrease LSDB size and improve stability: routers do not have to re-run the SPF algorithm as often, and when they do, it runs on a smaller LSDB. Understand that ABRs automatically summaries LSAs (the advertisements about routes). They do not automatically summarize the routing tables themselves. If you had 100,000 routes in your routing table with a single area design, you will still have 100,000 entries after using areas.


However, now that you have ABRs, you can configure route summarization, which can drastically reduce your LSDB and routing table sizes. Routers in Area 1 may only have the few dozen or hundred routes in that area, plus a default route (0.0.0.0) to reach all other subnets via the ABRs. These benefits still apply even for powerful routers that could support very large areas.

This is especially true as relates to network stability. With ABRs and route summarization, you have effectively segmented your network into smaller failure domains. A flapping route in Area 1 has absolutely no effect on Area 0 or Area 2. This can be a big boost in network stability.

and ease of troubleshooting.

Confidential - For Training Purposes Only

ABRs



- Interface in area 0 and at least one other area
- Separate LSDB per area
- Advertisises inter-area routes as Type 3 summary LSAs
 - Generates Type 3 LSAs for local areas and adds to area 0 LSDB
 - Copies Type 3 LSAs from area 0 LSDB into local area LSDBs



MOD 2- 44

You learned how each OSPF area defines an isolated region beyond which Type 1 and Type 2 LSAs and shortest path topologies do not extend. ABRs receive these Type 1 and 2 LSAs from one area, summarize them into a Type 3 summary LSA, and inject them into another area.

Area 0 has a special role in multi-area networks: the backbone. All areas must connect directly to the Area 0 backbone, and all inter-area traffic must traverse the backbone. In the figure in the slide, Area 1 reaches Area 2 via the backbone, and vice versa, but always through an ABR. Each ABR must have at least one interface in area 0, and at least one interface in some other area. So, ABRs have interfaces in at least two areas, one of which is area 0. OSPF design requires all areas to connect to this backbone. This is nearly always a physical link, but there is a logical “virtual link” that is sometimes used.

Each ABR maintains one LSDB per connected area. Think about it. All those Type 1 and 2 LSAs are synchronized among all routers in Area 1: all Area 1 routers have an identical LSDB. This is true for all areas. These Type 1/2 routes are summarized into Type 3 LSAs, and injected into the other areas. So, Area 2 does not have any topology information for Area 1 (connectivity detail) but area 2 knows about all Area 1 networks. It simply knows that it can reach these subnets via any connected ABR.

In this example, Core-1 is an ABR. It synchronizes Type 1 and Type 2 LSAs with other OSPF routing devices in area 1. The figure shows just two other routing devices, Access-1 and


Access-2, for simplicity. In the real world, the area could have dozens or hundreds of routers. Access-1 and Access-2 reside entirely within area 1 and are called internal routers. That connected ABR, Access-1, and Access-2 have identical area 1 LSDBs.

The ABR follows the same process in area 0, exchanging Type 1 and Type 2 LSAs and synchronizing its area 0 LSDB with routers such as Router-1 and Core-2.

In each area, the ABR also creates Type 3, or summary LSAs, which advertise routes that it knows in other areas. The term “summary” is kind of misleading since ABRs do not summarize routing information automatically—you must configure this (you’ll see this later in this section). The sections below elaborate on how the ABR conducts this advertisement.

Confidential - For Training Purposes Only

ABRs: Type 3 LSAs



- ABR Core-1 creates a Type 3 LSA per Area 1 subnet, places them in Area 0 LSDB
- Advertises to Area 0 internal routers
- ABR Core-2 receives these Area 1 Type 3 LSAs, injects them into Area 2
- Type 3 LSAs include a metric (cost), for best-path selection between multiple ABRs



MOD 2- 45

The Core-1 ABR creates a Type 3, or summary LSA, for each network in area 1. It places these Type 3 LSAs in its LSDB for area 0. The ABR can then advertise its routes to area 1 networks to other routers in area 0.

For example, Core-1 learned of an inter-area route to 10.1.20.0/24 from Access-1's Type 1 LSA, contained in its area 1 LSDB. From that LSDB, Core-1 creates an intra-area route to 10.1.20.0/24. Core-1 then generates a Type 3 LSA for 10.1.20.0/24 and places it in its area 0 LSDB. Other area 0 routers, such as Core-2, receive Core-1's Type 3 LSA and add it to their area 0 LSDB. They can then calculate an inter-area route to 10.1.20.0/24. Core 2 also injects these Type 3 LSAs from Area 1 into Area 0.

In addition to the network address, the Type 3 LSA includes a metric (cost), which enables routers to select the best inter-area route when multiple ABRs advertise a Type 3 LSA with the same network ID.

Of course, the same thing happens in the other direction: Core-1 learns creates Type 3 LSAs for the networks in Area 0 and injects them into Area 1.

And Core-2 has generated Type 3 LSAs for each network in area 2 and advertised them in area 0. Core-1 calculates an inter-area routes to each area 2 network, generates Type 3 LSAs for them, and injects them into area 0. They arrive at Core-1, which injects them into Area 1.

If Core-1 connected to another area, such as area 3, it would place Type 3 LSAs for each internal area 3 route in the area 0 LSDB and the area 1 LSDB. In other words, for each area, an ABR creates Type 3 LSAs with routes to other areas to which the ABR connects. Remember, though, that the inter-area routes always traverse area 0 even if the same ABR supports multiple non-backbone areas. That is, you cannot configure an OSPF router with area 1 and area 3 and expect the router to advertise area 3 networks in area 1 and area 1 networks in area 3. The router must have interfaces in area 1, area 3, and area 0.

The figure in the slide shows how internal routers Access-1 and Access-2 receive the Type 3 LSAs for area 0 and area 2. Note that, although area 1 routers learn routes to area 2, the Type 3 LSAs themselves do not cross area boundaries. Instead, Core-2, connected to area 2 and area 0, generates Type 3 LSAs for area 2 in area 0. It is the advertising router for these LSAs. Core-1 receives these LSAs in area 0 and generates new Type 3 LSAs for advertising these routes into area 1, and it is the advertising router for these new LSAs.


Confidential - For Training Purposes Only

Need for Route and LSA Aggregation Between Areas

Challenge

Routers do not require granular knowledge of other areas

- Causes unnecessary disruptions when small changes occur
- Clutters the LSDB and routing table



MOD 2- 46

Simply dividing an AS into multiple areas begins to decrease the size of OSPF routers' LSDBs and to reduce the complexity of SPF calculations. However, to achieve the true benefits of multiple areas, you should further reduce the number of LSAs advertised between areas.

In this example network, both area 1 and area 2 have only two ABR exit points toward the backbone and other areas. Internal area 1 routers and backbone routers might benefit from some level of detail in the information that they receive about area 2. For example, each ABR's Type 3 LSAs reflect the metrics of reaching the network in question, enabling routers in other areas to select the best ABR for reaching each network in area 2.

However, with only two entry points into area 2, routers (and, in particular, internal area 1 routers) are probably receiving too much granularity about each subnet. This information clutters the routers' LSDB and routing tables.

In addition, small changes in the area 2 topology could flood LSUs through the backbone and area 1, causing routers to re-run the SPF algorithm. For example, network 10.2.40.0/24 in area 2 becomes unavailable. The two area 2 ABRs withdraw their Type 3 LSA for that specific network in area 0, and then the two area 1 ABRs do the same in area 1. However, the area 1 routers have no other path to this network. It might be best to simply not inform routers beyond area 2 of the small change. If they continue to forward traffic for this network, the area 2 ABRs simply drop it.

You should aggregate these Type 3 LSAs to hide the unnecessary information and to smooth over small changes, thereby gaining the true benefits of a multiple area OSPF configuration.

Note: Although this example focuses on aggregating advertisements for area 2 networks, you would also aggregate summaries of area 1 or area 0 networks for similar reasons.


Confidential - For Training Purposes Only

Problem Solved: ABR Route Aggregation

Solution

ABR area summary

- Generate a single Type 3 LSA for the specified range
- Suppress any Type 3 LSAs within that range



MOD 2- 47

You can configure an ABR to aggregate its summary routes. Rather than generate a Type 3 LSA for every network in an area, it generates a few Type 3 LSAs, each of which includes many networks in its range. In other words, the Type 3 LSA advertises a network with a shorter prefix length such as a /20 or a /16 subnet that includes many /24 subnets.

You configure area summaries on the ABR that connects to the area to be aggregated. In the figure, you configure area summaries for area 2 on the two ABRs connected to Area 2. These ABRs then advertise the summaries to other backbone routers and ABRs in area 0. The other ABRs can then copy them into their own areas. In this example, the aggregated summary 10.2.0.0/16 shows up as a route on ABR A and on internal routing switches in area 1.

Similarly, you could configure the ABRs connected to Area 1 to summarize Area 1 routes and inject them into Area 0. All ABRs connect to Area 0, so you can configure them all to summarize area 0 routes and inject them into other area.

You see the drastic reduction in routes. Before summarization, Area 1 had 255 entries for Area 2. Now it has 2 entries. The figure in the slide shows the syntax to configure this aggregation. When you enter the area x range command, the ABR automatically suppresses the generation of individual Type 3 LSAs for subnets within the specified network. However, it continues to advertise individual Type 3 LSAs for networks that do not fall into the range. Thus, it is best to assign contiguous IP addresses within an area—it ensures easy route aggregation.


The ABRs for area 1 now receive one Type 3 LSA with the aggregated summary from each of the two ABRs that connect to area 2. Each area 1 ABR selects the best LSA for the summary range and creates an inter-area route to 10.2.0.0/16 through the best next-hop. Each area 1 ABR then generates an area 1 Type 3 LSA to advertise this inter-area route to routers in area 1. The internal area 1 routers now have two LSAs for all area 2 networks, and they learn routes to only the one aggregated network (up to two routes if the ABRs advertise equal metrics and the internal routers support Equal-Cost Multi-Path [ECMP]).

Note: After you configure area summaries on an area's ABR, you can see the summaries in the LSDB and routing tables for routers not connected to that area.

You typically want to create as few area summaries as possible, but this is not always possible. Areas might have grown organically and without advance planning, creating non-contiguous ranges or ranges that do not fit neatly within a CIDR subnet. For example, area 2 might include networks in the 10.2.0.0-10.2.255.255 range as well as 10.3.0.0-10.3.23.255. Area 3 includes networks in the 10.3.0.0-10.3.19.255 range and the 10.3.30.0-10.3.39.255 range. You cannot simply use 10.2.0.0/16 for the area 2 summary and 10.3.0.0/16 for the area 3 summary, due to the non-contiguous and overlapping ranges. Instead, you must use several area summaries for area 2 and several summaries for area 3.

Important: Examples throughout this module use the /16 prefix length for summary ranges to illustrate the concept in a simple way. When you create a range for your area, choose an appropriate range. For example, your area might include 16 /24 subnets, 10.1.0.0/24 to 10.1.15.0/24. This is equal to the 10.1.0.0/20 network, and you should specify that in the area <ID> range command.

Need for Redistributing External Routes into OSPF



MOD 2- 48


As a link state protocol, OSPF routers advertise network interfaces that run OSPF. But sometimes an OSPF router has other routes in its IP routing table, which must be advertised to OSPF neighbors. Perhaps one or more routers connect to an ISP with static default or BGP routes. You must inject these routes into your OSPF domain. Or suppose you are migrating your network from RIP to OSPF. During the upgrade, some segments run RIP. You must inject these routes into the already-migrated OSPF AS, and the old RIP segments must know about the OSPF networks. For security and efficiency reasons, you might choose not to enable OSPF on a stub network. Therefore, you redistribute connected routes into OSPF, to be advertised to the rest of your AS.

Use route redistribution to connect OSPF to any foreign routing protocol.

External Routes and ASBRs

```
router ospf 1
 redistribute static route-map My_Optional_Map
 default-metric 100000
```

ASBRs redistribute routes, advertised as Type 5 LSAs



ABRs inject into area as Type 4 LSAs



MOD 2- 49

The router that redistributes external routes is referred to as an autonomous system border router (ASBR) because it communicates routing information between the OSPF AS and other routing processes. A redistributed route becomes an external OSPF route, sometimes referred to as an AS external route. All routes calculated from Type 1, Type 2, and Type 3 LSAs, including both intra-area and inter-area routes, are internal OSPF routes.

ASBRs advertise external routes in Type 5 LSAs. The LSA includes the network address for the redistributed route, the router ID of the advertising ASBR, and the metric, which you assign when you redistribute the route into OSPF. OSPF routers propagate the Type 5 LSA in LSUs throughout the entire AS, making Type 5 LSAs different from Type 1, Type 2, and Type 3 LSAs. (Remember: although Type 3 LSAs advertise inter-area routes, each area has its own set of Type 3 LSAs.)

To calculate an external route from the Type 5 LSA, a router refers to the LSA's advertising ASBR. The router has calculated a shortest path to each ASBR. The forwarding interface and next hop for that path become the forwarding interface and next hop for all external routes advertised by that ASBR. In other words, routers do not add the external networks to the directed graph for the shortest path algorithm. Instead, they add them to a list of networks for which the advertising ASBR is the gateway. Therefore, changes to the particular external networks advertised by an ASBR do not affect the area topology.

The figure in the slide shows configuration syntax for ASBRs to redistribute static routes into


OSPF. You could also specify connected (directly connected links not running OSPF) or BGP. The route map syntax is optional. By default, all routes of the specified type are redistributed unless you employ route maps. When you redistribute routes into OSPF, they are assigned a seed metric to them: a beginning cost value. If you specify a default metric with the default-metric command, that would be applied to redistributed routes; if you want to have different metrics for different external routes, then use a route map to accomplish this task.

Note: Intra-area routes are always preferred over inter-area. Inter-area routes are always preferred over external routes. These rules are applied before the cost evaluation.

Confidential - For Training Purposes Only

External Routes and Route Maps

```
ip prefix-list PL_Net10 seq 10 permit 172.16.0.0/16
route-map RM_Net10 permit seq 10
  match ip address prefix-list PL_Net10
router ospf 1
  router-id 1.1.4.1
  area 0
    redistribute static route-map RM_Net10
```



MOD 2- 50

By default, when you redistribute a certain type of route into OSPF, the AOS-CX switch redistributes all routes of that type. However, you can filter the routes that are redistributed at the ASBR by applying a route map to the redistribute command, as shown in the figure in the slide.


The route map refers to an IP prefix list that specifies the networks to be redistributed into OSPF. In this example, the ASBR has several static routes, but you only want to redistribute 10.10.0.0/16 into OSPF. Optionally you can use various set parameters with match criteria, to modify how routes are advertised. You will learn more about this later in this module.

Verify External Routes

```
ip prefix-list PL_Net10 seq 10 permit 172.16.0.0/16
route-map RM_Net10 permit seq 10
match ip address prefix-list PL_Net10
router ospf 1
router-id 1.1.4.1
area 0
redistribute static route-map RM_Net10
```

```
ASBR# show ip route
...ipv4 routes for forwarding
'[x/y]'denotes [distance/metric]
10.8.0.0/16, vrf default
via 10.1.100.1, [1/0], static
10.9.0.0/16, vrf default
via 10.1.100.1, [1/0], static
172.16.0.0/16, vrf default
via 10.1.100.1, [1/0], static
```

```
ASBR# show ip ospf lsdb external
OSPF Router, ID (1.1.4.1) (Process ID 1 VRF default)
=====
AS External Link State Advertisements
-----
LSID          ADV Router  Age Seq#       Checksum
-----
172.16.0.0    1.1.4.1    305 0x80000002  0x0000adef
```




MOD 2- 51

The figure in the slide shows the ASBR routing table. Notice that the last static route matches the prefix list, as highlighted in green in the figure.

The figure also shows the LSDB entry for the redistributed route. Note that on the ASBR, the local static route, because its administrative distance is 1 is placed in its local routing table; however, other adjacent OSPF routers will receive the LSDB redistributed route and place that in their routing table if it happens to have the best distance and metric.

Need for Eliminating Excessive External Routes



Examine each non-backbone area: Does it really need all external routes?

If no ASBRs, all external traffic must traverse the ABRs

If only 1 ABR, then only 1 path out of the Area – just use a default route



Even if 2 ABRs, default routes from ABRs may suffice

MOD 2- 52


Look at the scenario in the slide, where 200 routes from ISP 1 and 150 routes from ISP 2 are injected into Area 0 as Type 5 LSAs. The ABRs then advertise these into Area 1, increasing route table size by 350 routes, along with associated memory and CPU usage. The internal routers in some OSPF areas rarely require all external routes.

Remember the primary benefit of individual Type 4 and Type 5 LSAs. Type 5 LSAs help routers choose the best ASBR for reaching an external network. If that ASBR is in another area, Type 4 LSAs further help the routers choose the best ABR to reach that ASBR, and perhaps to select the best ASBR. You should examine each non-backbone area and determine whether routers in the area truly need to choose between paths for particular external networks.

First, consider whether the area includes any ASBRs. If not, then all external traffic must traverse the ABRs. If there's only one ABR, there's only one path for external traffic. A default route through the ABR provides the same benefit as 350 individual external routes.

If the area has multiple ABRs, internal routers might benefit from choices for external paths. But ABRs often offer roughly equivalent paths to external networks, so internal routers can simply receive a default route from both ABRs.

Special Area Types



Often an area only has one or two ABR exit points: internal routers will choose the same next hop or hops for all external routes. This is when you might choose to filter the Type 5 LSAs at the area boundary. You do so by configuring the area as a special type: stub or totally stubby. These area types reduce the type 3, 4 and/or 5 LSAs by only advertising a default route.


Stub Areas Eliminate External Routes

Internal routers in stub area:

- No route redistribution

ABRs for stub area:

- Suppress Type 4/5 LSAs
- Inject a default route, IA routes



MOD 2- 54


To eliminate external routes and advertisements from an area, define the area as a stub area. ABRs for stub areas do not forward Type 4 or Type 5 LSAs into those areas, and internal routers in those areas do not generate or accept them. You must define the stub setting on every OSPF router in the area: ABRs and internal routers. Use the CLI syntax shown in the middle of the slide. Routers with mis-matched area types will not become neighbors –routing is broken.

When you configure an ABR area as a stub, it automatically generates a default route and injects it into the stub area as a Type 3 LSA. As usual, the ABR assigns a metric to the default route, and you can manipulate metrics to set up one ABR as a preferred exit for external traffic.

Notice that Type 3 LSAs for inter-area (IA) routes are still injected into the area, just like a normal area. Internal stub area routers receive summary LSAs for other areas through ABRs; the summary LSAs include individual or aggregated routes as configured on the ABRs. Therefore, internal routers in the stub area can continue to select the best path for inter-area traffic, based on the cost of the summary routes advertised by the ABRs.

In larger networks, you might also want to prevent these IA (inter-area) routes from being injected into the area.

Need for Eliminating IA and External Routes




MOD 2- 55

The figure in the slide shows a scenario where area 1 is a stub. All traffic destined to it terminates there. The area has only one ABR exit for all non-local traffic, including traffic destined to the backbone, to other OSPF areas, and to the ISP. In such an environment, the internal area 1 routers do not gain a benefit from multiple inter-area routes.

You can configure the OSPF system to suppress all inter-area routes in this area in addition to all external routes.

Totally Stubby Areas Eliminate IA and External Routes



MOD 2- 56

You define a totally stub area by configuring the ABR or ABRs connected to that area to suppress summaries. Use the OSPF configuration shown in the figure on all ABRs that connect to the area. Configure internal, non-ABR routers in the area with the regular area 1 stub command.

The no-summary option on the ABR(s) transforms a stub area to a totally stub area: it suppresses all external route advertisements as normal for a stub area. It also prevents the ABR from generating non-aggregated Inter-Area (IA) summary routes for this area. For example, the ABR in the figure would not generate area 1 summary LSAs for known routes in area 0 and area 2. Instead, the internal routers in area 1 use the default route to reach the other areas.

The number of routes eliminated may seem fairly trivial in this example, but in more complex networks, totally stubby areas can greatly reduce the size of LSDBs and route tables.

Setting Interface Costs

```
Switch # show ip ospf include Bandwidth
Reference Bandwidth: 100000 Mbps
```

```
Switch# show ip ospf interface
```

```
IP address 10.0.13.1/24,
State Backup-dr, Status up,
Link Speed: 1000 Mbps
Cost Configured NA, Calculated 100
```

$$100000 / 1000 = 100$$

OSPF perceives
same cost for
these interfaces

Bandwidth	Cost
1 Tbps	1
100 Gbps	1
20 Gbps	5
10 Gbps	10
4 Gbps	25
2 Gbps	50
1 Gbps	100

Default reference bandwidth

Bandwidth	Cost
1 Tbps	1
100 Gbps	10
20 Gbps	50
10 Gbps	100
4 Gbps	250
2 Gbps	500
1 Gbps	1000

```
router ospf 1
reference-bandwidth 1000000
```

Increase reference bandwidth

MOD 2- 57

OSPF path selection is based on link cost, which is based on the reference bandwidth, which defaults to 100000. This is shown in the figure, with the output of the show ip ospf command. AOS-CX devices divide this reference bandwidth by interface link speed in Mbps to determine cost. So, the default reference of 100000 Mbps / 1000 Mbps (1 Gbps) = 100. This is proven in the output of show ip ospf interface.

One problem with this default is shown in the figure in the slide. The default reference of 100000 Mbps / 100000 Mbps (100 Gbps) = 1. OSPF cannot use fractions, and so all interfaces faster than 100 Gbps will also have a cost of 1. Thus, OSPF will perceive a 100 Gbps path and a 1 Tbps path as being equal cost. If you have no links faster than 100 Gbps, this is not an issue.

If you do have links faster than 100 Gbps, then you should change the reference bandwidth to a higher value. In the figure, the reference bandwidth was changed to 1,000,000. Thus, 1 Tbps links have a cost of 1, while 100 Gbps links have a cost of 10, and so on, as shown in the figure. OSPF can now distinguish between these fast links.

Another issue is that by default, AOS-CX VLAN interfaces have a link speed of 1000 Mbps (1 Gbps), and so have a cost of 100. Because every VLAN interface has the same cost, route selection becomes a matter of hop count, rather than true lowest cost path, if you are using layer 3 VLAN interfaces. If you modify the reference bandwidth to 1,000,000 as shown, then all VLAN interfaces would have a cost of 1000.

Changing the reference bandwidth does not mitigate the issue of VLAN interfaces having the same cost. If this is an issue, you can change the individual cost values on a per-interface basis:

```
Switch(config)# interface vlan41
```


```
Switch(config-if)# ip ospf cost <1-65535>
```

As you can see in the generic example above, you can set each individual interface to an OSPF cost value of your choice (any value between 1 and 65535 inclusive).

Confidential - For Training Purposes Only

OSPF and BFD

Objective Fast detection of lost neighbors on broadcast networks



Without BFD

Router B fails, 4-second dead timer expires
Then A converges on new best path via C

Without BFD

Router B fails, Router A quickly knows
Then A converges on new best path via C



MOD 2- 58

There are several ways you can help OSPF respond quickly and gracefully in the event of topology changes such as failed link. Bidirectional Forwarding Detection (BFD) tests the connectivity between two IP addresses in a BFD session. BFD reports to OSPF when connectivity is lost. Routers use that information to take appropriate actions, depending on functions tied to BFD.

When you enable BFD on an OSPF layer 3 interface, BFD assumes responsibility for testing neighbor connectivity. If the BFD session fails, the router removes the BFD neighbor relationship, and so can remove all LSAs through that neighbor and update routing paths. The interface still sends out OSPF hellos, but BFD will generally detect any problems first.

In the figure in the slide, three routers connect on subnet 10.100.0.0/24. Routers B and C also connect to network 10.100.1.0/24. Router A uses Router B as the next hop for 10.100.1.0/24 because Router B advertised the lowest cost.

If BFD is not in use and Router B fails or its link on 10.100.0.0/24 fails, Router A continues sending traffic to 10.100.1.0/24 through 10.100.0.2 until its neighbor relationship times out. This could be 40 seconds if the dead timer is at its default value, or as fast as 4 seconds if you lower the dead timer to the recommended safe minimum. Because 10.100.0.2 is no longer up, this traffic does not reach its destination.

With BFD, Router A quickly detects that 10.100.0.2 is no longer available, where typically, this

is less than 1 second. BFD removes the neighbor relationship, causing Router A to lose adjacency with Router B. Router A can then immediately start using Router C as the next hop to 10.100.1.0/24.


Note: If an OSPF neighbor does not support BFD, the switch continues to try to establish a BFD session with that neighbor. It lists the session as down, but, because the session was never up, the down state does not cause the switch to tear down the neighbor relationship.

What makes BFD so much faster?

Confidential - For Training Purposes Only

BFD echo mode


Objective Runs over BFD session to provide faster failure detection



MOD 2- 59

When you configure BFD, echo mode is enabled by default. In echo mode, routers periodically send BFD echo packets at millisecond intervals. The peer router returns the received BFD echo packets back without processing them. If the sender does not receive BFD echo packet from the peer within the specified interval, the session is considered down. Because peers respond to echo packets without processing them, detection is fast.

BFD Configuration



1 `bfd`

Enable globally

`router ospf 1`
`bfd all-interfaces`

Enable on all
OSPF interfaces

2

`interface 1/1/1`
`ip ospf bfd`

Or...

Enable on some
OSPF interfaces



MOD 2- 60

First enable BFD globally, from global config mode, as shown in the figure. When you enable BFD, it automatically enables in echo mode. There is a `bfd echo disable` command, which disables echo mode, but this is not recommended. Use the `no` form of the command to re-enabled.

Next you can either enable OSPF BFD on all interfaces in the OSPF context. Or you can enter the desired Layer 3 interface context, and enable BFD on that interface specifically, as shown in the figure in the slide.

Because AOS-CX switches use active mode, the interface automatically attempts to establish a BFD session with each OSPF neighbor on that interface. It transmits a unicast BFD control packets on UDP port 3784 to each neighbor. As you see in the figure, a BFD session is point-to-point. In this example, the 10.100.0.0/24 network has three OSPF routers, each of which implements OSPF BFD. Each router attempts to establish a session with each of its two neighbors for six total BFD sessions on the network.

To establish a BFD session, the neighbors simply determine that both support BFD and negotiate timers.


Using OSPF with VSX

Transit OSPF interfaces

- L2 VSX LAG associated with L3 VLAN(s)
- Active forwarding on VLAN

Interfaces for which VSX pair is DG

- Active gateway
- Possibly passive OSPF interface



MOD 2- 61

Each VSX member runs OSPF separately. In the figure in the slide, aggregation layer switches Agg-1 and Agg-2 are a VSX pair. They have a VSX LAG to Core-1, and each switch has an IP address on the Layer 3 VLAN interface assigned to the LAG. If you were using multiple VRFs, the LAG would carry multiple VLANs; but in this example, the LAG has just one VLAN.

Core-1 learns ECMP routes to VLAN 10 and 20 subnets through both Agg-1 and Agg-2. Because the ECMP next hop selection algorithm will not necessarily match the link aggregation selection algorithm, the core switch might send traffic through next hop 10.0.0.1 but on the link to Agg-2. The active forwarding, which is set per-VLAN, enables Agg-2 in this situation to route the traffic itself even though the traffic was destined to 10.0.0.1. You should typically enable active forwarding on every OSPF interface that is a transit network.

The VSX pair is probably the default gateway for subnets such as those associated with VLANs 10 and 20. You should typically set up the active gateway feature on those VLANs. Recall that this feature enables the VSX switches to act as active-active default gateways for a subnet. You might run OSPF on these subnets, but in passive mode: VSX peers need not become OSPF neighbors on all subnets. Alternatively, you could redistribute connected routes into OSPF.



Virtual Routing and Forwarding

Virtual routing and forwarding (VRF) is a great way to improve routing flexibility, especially in situations with overlapping IP address space. This chapter provides an introduction to VRFs.

Confidential For Training

VRF Overview

One physical device



Multiple logical routers

- Separate tables and processes
- Overlapping addresses without conflict

Use cases

- Provider that supports multiple customers
- Integrate two merging companies
- Logical separation of business units
- Separate business functions for security



MOD 2- 63

Virtual Routing and Forwarding (VRF) allows a single physical device to host multiple logical routers, with complete L3 isolation. Because the device maintains separate route processes and tables for each VRF, they can have overlapping address spaces. Network functionality is improved because network paths can be segmented without requiring multiple routers.


Perhaps the scenario shown in the figure in the slide represents a single service provider, hosting separate environments for three customers. Customers can use any address space they choose, without fear of conflicting with other customers. Another use case is when you are merging with another organization. You can physically connect the two organizations without having to worry about addressing conflicts. As you normalize the address space between the two entities, you can eliminate the VRF configuration as appropriate. Another common use case is to maintain logical separation of business units. For example, a banking company might maintain separation between their retail, lending/mortgages, investments, and internal applications.

A common use case is to separate business functions. Place network management functions in one VRF, with the rest of the network in another VRF. Management network routes are unknown to the rest of the network, providing a logical separation between the two—you get enhanced security. Or in a data center place public resources in one VRF, with private or internal services in a separate VRF. Even though the both sets of services are in the same data center, they are logically separated since their routes are in different routing tables, thus enhancing security.

Let's analyze some scenarios in more detail.

Confidential - For Training Purposes Only

VRF Management Use Case



MOD 2- 64

In this example, network management traffic is logically separated from the rest of the campus traffic. One set of switches support two separate VRF instances, each with their own, separate set of routing tables. All campus users and services connect to interfaces that have been placed in the Campus VRF. This logically separates the campus management and data plane functions. Since nobody in the campus VRF can “see” the management VRF routes, they can’t reach them. You get solid security between the two domains, with no access control lists (ACLs) required.


Of course, the management VRF does not have campus routes in its routing table and so devices in the management VRF cannot reach resources in the campus VRF.

VRF Service Provider Use Case

VRF: Client1

Routing table

10.1.1.0/24
10.1.2.0/24
10.1.3.0/24
10.1.4.0/24
...
10.1.254.0/24



VRF: Client2

Routing table

10.1.1.0/24
10.1.2.0/24
10.1.3.0/24
10.1.4.0/24
...
10.1.254.0/24

Separate route tables



No access to each other's networks



MOD 2- 65

In this example, a service provider uses AOS-CX switches and VRF to logically separate two companies, Client1 and Client2. Notice that even though both companies have an overlapping route, the VRF implementation logically separates the two companies so that they cannot directly interact with each other.

By default, the two companies require an intermediate firewall to interact with each other, using address translation to deal with the overlapping address space. The two companies would have their own public address space and thus would translate their internal address space to a unique public address space in order to interact with other resources in the data center or on the Internet.

Note: Different interfaces must be used for the two VRFs: a VRF cannot share an interface.

Configure VRF

```
hostname Agg-1
vrf client1

interface loopback 10
vrf attach client1
ip address 10.1.0.2/32

interface vlan101
vrf attach client1
ip address 10.1.1.2/24

interface vlan141
vrf attach client1
ip address 10.1.2.2/24
```

Agg-1

```
Loop0 10.2.0.2
Loop10 10.1.0.2
v1:10.2.1.2/24
v101:10.1.1.2/24
V41:10.2.3.2/24
V141:10.1.2.2/24
```



```
hostname Agg-1
vrf default (implied)

interface loopback 0
ip address 10.1.0.2/32

interface vlan 1
ip address 10.1.1.2/24

interface vlan 41
ip address 10.1.3.2/24
```



MOD 2- 66

To create a VRF, use the global config command `vrf <vrf-name>`. The VRF name can be up to 32 alphanumeric characters. You are allowed a total of 65 VRFs on your routing switch. By default, two exist on the switch:

default: used if you do not specify a VRF with a configuration command

mgmt: used for out-of-band management (OBBM)

This means that you can create an additional 63 VRFs. And remember, an interface can only belong to one VRF.

The figure in the slide shows an example of one physical infrastructure that supports two separate VRFs. The blue configuration to the left is for VRF “client1”. This configuration is reflected in the network diagram) just focus on the blue addresses and ignore the grey addresses). Notice that each interface is configured with `vrf attach client1`. This is how you tell the router which interfaces belong to which VRF. If you enter the command `no vrf attach client1`, the interface is returned to the default VRF, and all interface configuration is removed. This is true anytime you move an interface to a different VRF.

The configuration to the right is for VRF default, also reflected in the network diagram. Just focus on the grey addresses and ignore the blue addresses. The command `vrf default` is shown but need not be configured – it is the default.

Note: This configuration matches the Service Provider use case you just reviewed. You would probably use VRF “client1” and VRF “client2”. The default VRF was used for client2 to ensure you understand that all interfaces are in the default VRF unless otherwise specified.

You have now split one physical device into two logical routers, but no routing protocol is configured. Let’s add OSPF routing to this deployment.

Confidential - For Training Purposes Only

Configure VRF with Routing

```
hostname Agg-1
vrf client1

router ospf 10 vrf client1
area 0.0.0.0

interface loopback 10
vrf attach client1
ip address 10.1.0.2/32
ip ospf 10 area 0

interface vlan101
vrf attach client1
ip address 10.1.1.2/24
ip ospf 10 area 0

interface vlan141
vrf attach client1
ip address 10.1.2.2/24
ip ospf 10 area 0
```

Agg-1

Loop0 10.2.0.2

Loop10 10.1.0.2



v1:10.2.1.2/24

v101:10.1.1.2/24

V41:10.2.3.2/24

V141:10.1.2.2/24

```
hostname Agg-1
vrf default (implied)

router ospf 1
area 0.0.0.0

interface loopback 0
ip address 10.1.0.2/32
ip ospf 1 area 0.0.0.0

interface vlan 1
ip address 10.1.1.2/24
ip ospf 1 area 0.0.0.0

interface vlan 41
ip address 10.1.3.2/24
ip ospf 1 area 0.0.0.0
```



MOD 2- 67

Many routing switch commands allow you to specify the VRF name to which the command applies. The figure shows the previous example, but with OSPF configuration added to the blue client1 configuration on the right. The syntax to configure OSPF is as typical, except for the addition of the VRF name at the end.

Many configuration and show commands allow you to reference a VRF name and thus compartmentalize your routing functions. If you do not use an optional VRF name, then the command applies to the default VRF.

Consider a static routing example, with the command `ip route 10.254.0.0/16 10.5.5.5`. This tells the router to use next-hop router 10.5.5.5 to reach destination network 10.254.0.0/16. Since you did not specify a VRF, this route is placed in the default VRF route table.


To place this route in the client1 VRF, use `ip route 10.254.0.0/16 10.5.5.5 vrf client1`

Validate VRF


```
Agg-1# show vrf
VRF Configuration:
-----
VRF Name   : default
Interfaces -----
loopback0      up
vlan1          up
vlan41         up

VRF Name   : client1
Interfaces -----
loopback10     up
vlan101        up
vlan141        up
```

VRF default



VRF client1



MOD 2- 68

The figure in the slide shows how to validate which VRFs exist, and which interfaces are associated with each VRF. To look at each VRF separately, you could use `show vrf default` and `show vrf client1`.


Policy-Based Routing

Use policy-based routing (PBR) to really take control over network routing decisions. Traditionally, you would configure a dynamic routing protocol, like OSPF as your IGP and configure some static routes for external connectivity. But there might be cases where you don't want to use the route decisions that OSPF or static routing has defined, but override these in certain situations. This is where PBR comes into play. This section provides a brief introduction to PBR.

PBR Overview

Normal routing

- Destination IP compared to route table entries
- Choose longest match destination in route table




Routing table:

Dest Network	Next Hop
10.1.1.0/24	10.2.2.2
10.1.0.0/16	10.3.3.3

Policy-based routing

- Compare various packet attributes to a classifier
- Take actions specified in a policy



Routing table:

Dest. IP	Next Hop
10.1.1.0/24	10.2.2.2
10.1.0.0/16	10.3.3.3

MOD 2- 70

The figure in the slide compares normal routing to Policy Based Routing (PBR). Normal routing is purely destination-based. As each router receives an inbound packet, it compares the packet's destination IP address to the route table, then chooses the entry with the longest match as the best path. In the example, the packet destined to 10.1.1.1 matches both entries in the route table, but only 16 bits match in the top entry, while 24 bits match in the second entry. So, the router routes this packet via next hop 10.2.2.2. Next, a packet arrives destined to 10.1.2.2. This destination only matches the top entry, and so 10.3.3.3 is the next hop.

PBR lets you manipulate the path of a packet based on various packet attributes. Packets must be destined to a valid subnet on your network. As each router receives an inbound packet, it compares appropriate header and tag information to a classifier that you configure. The router then takes action on this matching traffic based on a policy. Matching traffic with the same destination can be routed over different paths, to improve traffic handling, performance, and/or security for various traffic types.


In the example, both packets are for the same destination subnet. However, you have created a classifier that matches on the range of UDP ports used for Voice over IP (VoIP) traffic, perhaps UDP ports 5004-5065. This traffic is routed over its own path via destination 10.2.2.2. Note that this action overrides what normal routing would do by default – the ingress router only has next-hop 10.3.3.3 for 10.1.0.0/16. However, 10.2.2.2 is a valid path, and it is directly connected to the ingress router, as required by PBR. In other words, recursive lookups are not supported.

All other, non-VoIP traffic does not match the classifier, and so is routed as normal, via 10.3.3.3. PBR only affects the local router's routing decision. You might have to implement the PBR policy on multiple routing devices in order to achieve the desired outcome if multiple hops exist between a source and destination.

Confidential - For Training Purposes Only

PBR Use Cases

Load sharing



Best path: Most traffic

Alternate path: Some traffic

Cost savings



To ISP1: Flat rate per month

To ISP2: Per-GB charges

QoS



Data traffic

VoIP traffic

Provider selection



10 Gbps link: Normal service

100Gbps link: Premium service

MOD 2- 71

By default, a routing device uses its routing table to determine which path to take when reaching a destination. There are times, though, where a company needs to override the default routing behavior based on traffic types or network conditions. The figure in the slide summarizes the use cases.

Load sharing: Imagine that you have two links for redundancy for an OSPF network, but one link has a better cost. The higher cost link is available, but is not used, unless the better cost link fails. Configure PBR to use the secondary link for certain traffic and thus implement a limited form of load balancing.

Cost savings: Suppose that you have two Internet connections via two different providers: ISP1 charges a flat rate and ISP2 charges a flat rate with a usage cap. Once you reach the cap, you are charged per gigabyte sent or received. Use PBR to ensure that even if the ISP2 has better metric routes, that ISP is not used unless ISP1 is unreachable.

QoS: Perhaps you have two equal-cost paths for some destination. You want to send delay and jitter-sensitive VoIP traffic on one link, and all other traffic over the other link. You maintain this redundancy while logically separating the traffic. VoIP traffic gets a link all to itself, but you must still configure QoS in case the VoIP link fails.

Provider selection: You might need to control traffic paths based on the source IP address, as opposed to normal destination-based routing. Imagine that you work for a service provider,

offers normal and premium-level services. The backbone has two sets of connections: 10 Gbps and 100 Gbps. Use PBR to ensure that traffic sourced from premium-tier customers traverse the 100 Gbps connections, while normal customer traffic uses 10 Gbps connections.

Confidential - For Training Purposes Only



Lab Activity

Let's review the lab activity that is associated with this part of the module.

Confidential - For Training

Lab Tasks

Lab 02.02 - Wired Routing

Task 1

- Basic OSPF Configuration

Task 2

- Route Redistribution and Filtering using Route Maps

Task 3

- Multi Area OSPF and Route Aggregation between Areas

Task 4

- Optimizing OSPF Failover

Lab 02.03



MOD 2- 73

The slide provides a brief review of tasks performed in this lab. Please see your Lab guide for details of this lab.




Aruba Central Template Groups

Template groups are a powerful way to streamline network administration. This section quickly reviews template group configurations. This topic was covered in detail in the Aruba Campus Access Fundamentals course.

Confidential For Training

Aruba Central Template Group



MOD 2- 75

Templates in Aruba Central refer to a set of configuration commands that you can use to provision devices in a group. You can apply a set of configuration parameters simultaneously to multiple switches in a group and thus automate switch deployments.

Working with Template Groups: Overview

- 1 Create template group
- 2 Create/add new template
- 3 Import configuration
- 4 Create, manage, modify variables
- 5 Add variables w/variable file, upload
- 6 Add conditional statements to config



AOS-CX switch template Group 1

6405_Template

```
#This template is imported from dev
!
!Version ArubaOS-CX FL.10.10.0002
!export-password: default
user admin group administrators pass
module 1/3 product-number r0x40a
profile default
ntp server pool.ntp.org minpoll 4 m
ntp enable

hostname %_sys_hostname%
%if vlan_1_dhcp%
ip dhcp
%else%
ip address %vlan_1_ip%/%vlan_1_mask%
no ip dhcp
%endif%
```

Variables file: download, modify, upload

_sys_serial	_sys_lan_mac	modified	_sys_hostname	vlan_1_static_ip	ip_gateway
SG90KM7I	f8:60:f0:5a:01:00	Y		192.168.0.10/24	192.168.0.1

MOD 2- 76

First you create a template group. This is a placeholder for templates, etc. This new template group of course contains no templates. So, you create a template and add it to the group. The template name cannot contain spaces or special characters other than dash or underscore. Specify the switch model, part name, and version the template will apply to.

Note: If multiple templates in a group could apply to a given managed switch, Central applies only the most specific template for that switch: the template with the most explicitly defined characteristics (model, part name, version) matching that switch.

Next, import a configuration for the template. Just choose a switch for which this template should be applied, and Central imports its configuration. For consistency, this configuration is applied to other appropriate switches.

Of course, each switch's name and certain other parameters must be unique, so you create, manage, and modify variables to suit your needs. Variables with a name starting with “_sys_” are auto-generated by Central whenever a device is added to a template group, and include the device serial number and MAC address; these two variables cannot be modified.

To add new variables, you must download the group variables file, modify it outside of Central, then re-upload to add the new variables and any modified values to the group.

Variables with a name starting with “_sys_” are auto-generated by Central whenever a device

is added to a template group, and include the device serial number and MAC address (these two variables cannot be modified and are mandatory)”

_sys_hostname - Maintains unique host name

_sys_lan_mac - Indicates system MAC address

_sys_serial - Serial number of the switch

Confidential - For Training Purposes Only

Add Variables: Example

	A	B	C	D	E	F	G	H
1	_sys_serial	_sys_lan_mac	modified	_sys_hostname	vlan_1_ip_address	vlan_1_use_dhcp	vlan_1_static_ip	ip_gateway
2	SG90KM700S	f8:60:f0:5a:01:00	Y	"	Y		192.168.0.10/24	192.168.0.1
3								
4								

Line 1: Names	vlan_1_ip_address	vlan_1_use_dhcp	vlan_1_static_ip	ip_gateway
Line 2: Values	Y		192.168.0.10/24	192.168.0.1

Add new variables: Enter names and values in columns E, F, G, and H



MOD 2- 77

You learned that each device assigned to the group will have its own line in the variables file, with the serial number and MAC address being the first two columns displayed. The third column contains a modified variable that is only utilized for downloaded CSV variables files, and is set to N by default. When you make changes to a downloaded variables file, this value must be changed to Y for any changes to variables to be applied to that device when the file is re-uploaded after changes are made.

For example, you might add the following new variables to the file by entering their names in columns E, F, G, and H in row 1: `vlan_1_ip_address`, `vlan_1_use_dhcp`, `vlan_1_static_ip`, `ip_gateway`. Next, set the following values in line #2 for each new variable:

`vlan_1_ip_address`: Y

`vlan_1_use_dhcp`: leave blank

`vlan_1_static_ip`: 192.168.0.10/24

`ip_gateway`: 192.168.0.1

Change the value of the modified variable for line #2 to Y, then save the file.

Conditional Statements

Selectively apply portions of a template based on the specific value contained in the variable

if, else, endif

```
hostname %_sys_hostname%
%if vlan_1_dhcp%
ip dhcp
%else%
ip address %vlan_1_ip%/%vlan_1_mask%
no ip dhcp
%endif%
```



MOD 2- 78

Aruba Central configuration templates may contain variables and conditional statements (if, else, endif) in order to apply the same template to multiple managed switches without applying unsupported features or duplicate values (such as IP addresses). To use these functions, enclose the variable or conditional statement in '%' characters, as in the examples shown.

Note: An %endif% statement must be present after every %if% statement in the template; any mismatch between the number of %if% and %endif% statements will be flagged as a template syntax error and the affected template will not be applied.

Conditional Example

The screenshot displays a network configuration template editor. On the left, a code editor shows a configuration snippet for interface vlan 1. The code includes conditional logic using `%if` and `%endif` statements. A red box highlights the conditional block for `interface vlan 1`. On the right, a panel titled 'Variables in Template' lists the variables used in the template: `_sys_hostname`, `vlan_1_ip_address`, `vlan_1_use_dhcp`, and `vlan_1_static_ip`. A red box highlights the first three variables. The interface also includes a 'TEMPLATE' tab, an 'IMPORT CONFIGURATION AS TEMPLATE' button, and a 'Hide Variables List' link.

```
347 ip address 192.168.50.6/32
348 ip ospf 1 area 0.0.0.0
349 interface vlan 1
350 %if vlan_1_ip_address = Y%
351 %if vlan_1_use_dhcp%
352 ip dhcp
353 %else%
354 ip address %vlan_1_static_ip%
355 no ip dhcp
356 %endif%
357 %else%
358 no ip dhcp
359 %endif%
360 interface vlan 2
361 ip address 20.1.1.3/24
```

Variables in Template

- `_sys_hostname`
- `vlan_1_ip_address`
- `vlan_1_use_dhcp`
- `vlan_1_static_ip`

The example shown in the slide will perform the following actions when the template is saved, and the resulting configuration is generated to be applied to the switch:

- The `vlan_1_ip_address` variable is checked for the string value Y; if present, continue to next line:

If any value is present for the variable `vlan_1_use_dhcp`, configure `ip dhcp`.

If `vlan_1_use_dhcp` is blank, configure the static IP address defined in `vlan_1_static_ip` and `no ip dhcp`.

- If `vlan_1_ip_address` is blank or any value other than Y, skip to the second `%else%` block and configure `no ip dhcp`.



Knowledge Check

Self-check on key learning points



It's time for a knowledge check to review what you have learned.

Confidential - For Training Purposes Only

Question #1

Which statements below accurately describe VSX and VSF?

- a. Both VSX and VSF share layer 2 links across physical members
- b. Only VSF uses shared management, control, and routing across physical chassis
- c. The solution that uses a shared control plane offers improved resiliency over the solution that uses a dual control plane
- d. A VSX pair can serve as an active-active first-hop gateway for endpoints
- e. You can improve VSX resiliency by configuring the Spanning-Tree Protocol (STP) between the VSX pair and access-layer switches

a

Knowledge Check

MOD 2- 81


+

a, b, d

Confidential - For Training Purposes Only

Question #2

Match each VSX component with its appropriate description.



Knowledge Check ✓



MOD 2- 82



Inter Switch Link (ISL)

Connects VSX peers, improves resiliency

Keep alive mechanism

Helps identify a split-brain condition

Linkup delay

Helps to avoid dropped traffic after a VSX switch reboots

VSX LAG

Leverages the VSX system-mac address

Question #3

Your deployment uses an OSPF network with four areas: 0, 1, 2, and 3. This network also connects to several external networks: some via BGP, and some via OSPF. You would like to minimize routing table size for OSPF internal routers in Area 1. Which statements below accurately describe valid options for you to consider in this scenario?

- a. Configure the internal routers in area 1 to block routes from other areas
- b. Configure your Area Border Routers (ABR) connected to Area 1 as a stub area
- c. Configure your Area Border Routers (ABR) connected to Areas 2 and 3 as a stub area
- d. Configure your Area Border Routers (ABR) connected to Area 1 to summarize routes that exist in Area 0
- e. Configure your Area Border Routers (ABR) connected to Area 1 to summarize routes that exist in Area 1

a

Knowledge Check

MOD 2- 83

+

b, d

Confidential - For Training Purposes Only

Question #4

Your OSPF network is running fine, except that routers are doing equal-cost load balancing over 1 Tbps and 100 Gbps links. What should you do to alleviate this issue?

- a. Disable equal-cost load balancing on all routers that have this issue
- b. Upgrade your 100 Gbps links to 1Tbps links
- c. Configure the link speed of the 100 Gbps to 1 Tbps
- d. Increase the reference bandwidth on all routers in the deployment

a

Knowledge Check

MOD 2- 84

+

d

Confidential - For Training Purposes Only

Question #5

One way to enable separate Virtual Routing and Forwarding (VRF) routing tables to share certain routes between them is to configure one physical interface to be a member of both VRFs.

- True
- False

q

Knowledge Check

MOD 2- 85

+

false

Confidential - For Training Purposes Only

Question #6

When using template groups how can you selectively apply portions of a template?

- a. Use conditional statements with appropriate values in a variable file
- b. Use separate CLI configuration scripts per group member
- c. Use separate variable files for per group member
- d. Download unique variables for each group member

a

Knowledge Check

MOD 2- 86

+

a



Lab Activity

Let's review the lab activity that is associated with this part of the module.

Confidential - For Training

Lab Tasks

Lab 02.03 - Campus Wired with Central

Task 1

- Onboard switch to Central with ZTP

Task 2

- Managing Edge Switch using Template Group

Task 3

- Migrate Aggregation Switches to Aruba Central

**Module 3 –
Wireless
Infrastructure
with Aruba
Gateways**



MOD 2- 88

The slide provides a brief review of tasks performed in this lab. Please see your Lab guide for details of this lab.


Rev. 23.11



Wireless Infrastructure with Aruba Gateways

Implementing Aruba Campus Access

Welcome back folks! This is Module 3 – Wireless Infrastructure with Aruba Gateways. In the Aruba Campus Access Fundamentals course you learned how to set up a campus design using standalone APs managed via Central. These APs used bridged mode to transport the traffic: the APs would receive wireless connections, authenticate them, encrypt/decrypt frames, if necessary, filter frames with firewall policies, and locally forward the traffic. This module will introduce the use of using Aruba APs and gateways together.



Objectives

- Use Aruba Central groups to efficiently manage, monitor, and maintain APs
- Deploy gateways to improve scalability and security
- Leverage gateway clusters to improve resiliency and performance

After completing this module, you will be able to use Aruba Central groups to efficiently manage, monitor, and maintain APs. You will also be able to deploy gateways to improve scalability and security, and leverage gateway clusters to improve resiliency and performance.

After completing this module, you should be able to:


- Use Aruba Central groups to efficiently manage, monitor, and maintain APs.
- Deploy gateways to improve scalability and security.
- Leverage gateway clusters to improve resiliency and performance.

AOS 10 AP & Central Groups

Let's begin by exploring AOS10 APs and Aruba central groups. This is a good refresher from the material found in the Aruba Campus Access Fundamentals course.

Confidential For Training Purposes

Aruba AOS 10 Architecture: Responsibilities



MOD 3- 4

Aruba AOS 10 is a new architecture that relies on Aruba Central for WLAN management, control and services. In this new architecture, the APs are responsible for:

Facilitating user and device authentication (the AP will act as an authenticator).

Forwarding user traffic to its destination (traffic can be bridged or tunneled to a gateway).

Encrypting and decrypting the user traffic

Forwarding telemetry information to Central including AP name, MAC address, serial number, AP model, CPU usage, uptime, mesh mode, LED status, tx channel, tx power, radio band, noise floor, SNR, ESSID, etcetera.

In an AOS 10 architecture, Aruba Central is responsible for:

WLAN Management: Device configuration and firmware management

WLAN Control: ClientMatch and AirMatch

WLAN Services: Key Management Services, AirGroup, RAPIDS, UCC, and etcetera


Using AOS10 implies the communication between Central and APs is important, since now Central runs part of the control plane. Therefore, using redundant Internet links with different

ISPs is highly recommended. In module 12 we will discuss about the impact of losing completely connectivity with Central.

AOS10 APs do not run any cluster nor elect a Conductor (previously known as Master) to run a Virtual Controller instance (this was referred to as Instant Access Point or IAP mode). In this new architecture Aruba Central replaces the VC role running part of the control plane, where some of these important tasks include ClientMatch, RAPIDS, device configuration, adaptive radio management, and etcetera.

Confidential - For Training Purposes Only

Central Groups



MOD 3-5



A group in Aruba Central is a primary configuration element that functions as a container for device management, monitoring, and maintenance. Groups enable administrators to combine devices with identical configuration requirements and manage these devices efficiently by using either a UI-based configuration workflow or CLI-based configuration template.

Consider that you have 10 sites across your country with a single AP at each site. A new WLAN needs to be configured to support contract users. The administrator has two approaches for this task. First configure individually each one of the APs or use the group configuration approach. The first option does not scale well and consumes too much time on the other hand a group is a faster mechanism since the administrator will only configure the WLAN one time and this configuration will be pushed to all the members of that group.

Option two has you place ten APs in the same group in Central. Then, create the new WLAN in that group in Central. Central will then push down the new WLAN to all ten APs!

Group Level vs. Device Level Configuration

Group level configuration

Refers to common configuration

contractor
WLAN

Device configuration **overrides**
group configuration

Device level configuration

Refers to per-device configuration for

AP name: **AP-43**
Location: **NY**

AP name: **AP-8**
Location: **CAL**

AP name: **Lobby-AP**
Location: **Florida**



MOD 3-6



As we have explained previously, the group level configuration in Central refers to common configuration that is pushed to devices. For example, a Contractor WLAN will be configured for all sites.

Aruba Central supports device level configuration as well, which is used to configure specific or unique values to each device in the group. To clarify this concept, imagine that a unique AP name parameter will be used on each of these 10 sites: administrators can individually configure these parameters from Central on each device. Device configuration, by its nature, does not scale well. In case a configuration conflict exists, the device level configuration takes precedence and overrides the group level configuration.

In conclusion, we can affirm that group level configuration should be used for all common configuration tasks for multiple devices and therefore this configuration mechanism scales, and that the device level configuration should be used for unique configurations on each device and overrides the group level parameters.

Default and Custom Groups

default group		
AP	Switch	Gateway
WLANs employee contractor	VLAN 10-20	VLAN 110-120
		

- System group, contains all new devices

Corp-Branches group		
AP	Switch	Gateway
WLANs employee IOT	VLAN 90-99	VLAN 190-199
		

- Custom user-defined groups
- Move devices here manually or use pre-provisioning

Move device to new group - inherits new group config



MOD 3- 7

Aruba Central uses a system group named default. When a new (factory-default) device communicates with Central this device is automatically mapped to the default group, the device then receives the configuration of this group. This group (default) facilitates the possibility to configure a device without intervention of the administrator. This process is known as Zero Touch Provisioning (ZTP).

Central allows a Custom group definition, in this case the administrator must manually move devices to the Custom group. A device can only be mapped to a single group at the time. Every time a device is moved to a new group, this device will wipe out its current configuration and it will be rebooted.

Note: Device level configuration overrides are retained when you move the device to a different group.

The figure in the slide shows how a single group in Central is used to configure APs, switches and gateways.

Group Persona for AOS 10 APs

The screenshot displays the Aruba Central interface for creating a group. The left sidebar shows the 'Organization' tab selected. The main content area is titled 'Groups' and includes a 'Network Structure' section with a '1' in a box. Below this, the 'Install Manager' section shows 'Site installations in progress' and 'Authorized installers'. The 'Architecture for access points and gateways in this group' section shows 'ArubaOS 10' selected. The 'Network role of the access points in this group' section shows 'Campus/Branch' selected. The 'Group will contain' section shows 'Access points' selected. The 'Configure using templates' toggle is turned off.

The figure in the slide shows how to define architecture and persona during group creation. While creating a group, you can define the current group's architecture and persona settings as preferred settings for adding subsequent groups. For subsequent groups, you can either automatically apply the preferred settings or manually select settings for the new group. Based on the device persona you select in a group; the device configuration page displays only particular device tabs for that group. For example, if a group has only access points persona assigned to it, then the device configuration page for that group displays only the access point tab.

Group Persona for AOS 10 Gateways

The screenshot displays the Aruba Central interface for configuring a new group. The left sidebar shows the 'Organization' tab selected. The main content area is divided into three sections:

- Section 1: Group Overview** (indicated by a red circle with '1') shows the group name 'P2401-GW-Group' and a count of 0 unprovisioned devices.
- Section 2: Group Configuration** (indicated by a red circle with '2') shows the group will contain:
 - ☐ Access points
 - ☒ Gateways
 - ☐ Switches
- Section 3: Architecture and Network Role** (indicated by a red circle with '3') shows:
 - Architecture for access points and gateways in this group:**
 - ☒ ArubaOS 10
 - ☐ ArubaOS 8
 - Network role of the gateways in this group:**
 - ☒ Mobility
 - ☐ Branch
 - ☐ VPN Concentrator

The bottom right corner of the interface shows 'MOD 3-9' and a red plus sign.

Personas (roles) are shown in the figure in the slide, and are described below:

- **Mobility:** Gateways provide Overlay WLAN and (or) wired (UBT) LAN functionalities. This persona applies only to the AOS10 architecture.
- **Branch:** In this persona, gateways provide Aruba SD-Branch (LAN + WAN) functionality. This persona applies to both AOS10 and AOS8 SD-Branch architectures.
- **VPN Concentrator:** The VPN concentrator (VPNC) persona is only available when configuring a gateway only group, the VPNC group cannot contain APs or switches.

Deploying Aruba Gateways



This section introduces a new skill: deploying Aruba gateways using Aruba Central!

Confidential - For Internal Use Only

When to Use Gateways

More than 5000
clients

More than
500 APs

Tunneled
WLANs

Network
simplification

A key element in



AOS10 architecture

Dynamic
segmentation

Enhanced
mobility

Microbranch
deployments

RADIUS
proxy



MOD 3- 11

Since the control plane resides now in the cloud (Aruba Central) you might question whether you still need gateways (formerly known as controllers). But gateways are still a key element in the AOS 10 architecture. They support tunneled and mixed SSIDs, allowing you to dramatically reduce the number of VLANs spread across the campus, with traffic inspection in a single location.


AOS10 also uses gateways for dynamic segmentation, allowing for authentication and traffic inspection switches to tunnel users to a gateway. Plus, AOS 10 needs gateways to deploy microbranch APs, which terminate their IPsec tunnels on gateways over the internet, to provide remote access to small branches.

But what about network simplification? How do gateways help with that?

Network Simplification

Must likely enable user VLANs on:

- Every AP-connected access switch
- Each AP port




User VLANs likely span entire network



Must only enable user VLANs on:

- Gateway
- Gateway ports on core (agg) switches




User VLANs do not span entire network

MOD 3-12

Typically, each AP supports multiple SSIDs, each mapped to a different wired VLAN. Without gateways, each AP must locally forward user traffic to its destination. This means that you must configure trunk ports on every AP-connected switch port in your deployment. You must ensure that each trunk link supports the appropriate VLANs for each AP. You could be lazy, and simply allow all VLANs to traverse all trunk links, but this is in bad form. You would have unnecessary traffic traversing all trunk links. Plus, this increases potential attack vectors from bad actors. In addition to all these details, keep in mind opportunities for human error. You have the issue of VLANs that span the entire network: this increases the size of your broadcast domains.

With gateways, all user traffic is tunneled to the controller. This simplifies your AP-connected switch port configuration. Plus, user VLANs do not span the entire network. Nice!

Gateway Default Port Config



MOD 3-13

Zero-touch provisioning (ZTP) is a fast and convenient way to onboard a new or existing gateway into Central—no configuration required from the installer. For ZTP, the gateway must be connected to a switchport that is configured with an untagged VLAN, which must provide DHCP addressing and Internet access. You can use any port on a gateway for ZTP, except GE 0/0/1.

For one-touch provisioning (OTP), you can use the console port or connect your PC to GE 0/0/1. This port has DHCP server enabled and has a local IP address configured: 172.16.0.254.

Provisioning Gateway with ZTP

1 Connect gateway to switch: untagged, DHCP internet access

2 Gateway resolves Activate FQDN, sends request →

← Activate sends firmware, Central FQDN

3 Gateway reboots, communicates with Central

4 Firmware upgraded as needed, reboot

5 Gateway contacts Central, gets config



Internet

6 Gateway is operational

HPE GreenLake



Aruba Activate



Aruba Central

MOD 3- 14

Here are the steps that occur when provisioning an Aruba gateway using ZTP:

Connect a ZTP capable port on the gateway to a switchport. This switchport must be an untagged (access) VLAN that provides DHCP and Internet access. Recall that all ports on the gateway support ZTP except GE 0/0/1.

Upon receiving a DHCP response, the gateway will resolve the Activate FQDN and communicate with Activate for provisioning:

- If the Gateway is new and has not been previously provisioned, Aruba Activate will push a Central enabled firmware upgrade and reboot the gateway.
- Activate will provision the gateway with the FQDN for the assigned Central instance.

After booting to a Central enabled firmware and provisioned with the FQDN, the Gateway will communicate with Central.

The firmware version defined in firmware compliance for the group will be enforced and an upgrade pushed if necessary; once an upgrade is completed the gateway will reboot


The gateway will initialize using the specified AOS 10 version, contact Central, and be configured based on the configuration applied in Central to the device.

Once the configuration has been successfully applied, the gateway will be up and operational

in Central either using the staging port or the configured uplink port(s).

Confidential - For Training Purposes Only

Provisioning Gateway with OTP



MOD 3- 15

Static Aruba Activate is a one-touch provisioning (OTP) option, used to provision gateways that require static addressing or PPPoE authentication. With OTP, you must use a serial console port or web-browser to supply minimum information to the gateway in order to permit initial communication with Activate and Central.

To use a web-browser, you must connect a computer to the gateway on the GE 0/0/1 Ethernet port, which provides a DHCP address for local access.

Serial Console OTP

```
Auto-provisioning is in progress. It requires DHCP and Activate servers
Choose one of the following options to override or debug auto-provisioning...
'enable-debug'      : Enable auto-provisioning debug logs
'disable-debug'     : Disable auto-provisioning debug logs
'full-setup'        : Start full setup dialog. Provides full customization
'static-activate'   : Provides customization for static or PPPOE ip assignment

Enter Option (partial string is acceptable): static-activate
Enter Controller VLAN ID [1]: 1
Enter Uplink port [GE 0/0/0]:
Enter Uplink port mode (access|trunk) [access]:
Enter Uplink Vlan IP assignment method (static|pppoe) [static]:
Enter Uplink Vlan Static IP address [192.168.1.1]: 172.16.1.101
Do you want to configure dynamic port-channel (yes|no) [no]:
...
Do you wish to accept the changes (yes|no)
```

MOD 3-16

Here are the steps involved with OTP via the serial port console connection on an Aruba gateway:

Connect to the gateway's serial console port and power on the Gateway. Once booted, you see the initial provisioning screen.

Type static-activate then press <Enter> to start the process; choose the options as appropriate for the required uplink type, either static or pppoe, then provide the required information.

Once the initial provisioning is complete the gateway will resolve the Aruba Activate FQDN and communicate with Activate for further provisioning:

- If the gateway is new and has not been previously provisioned, Aruba Activate will push a Central enabled firmware upgrade and reboot the gateway.
- Activate will provision the Gateway with the FQDN for the assigned Central instance.

After booting to a Central enabled firmware and provisioned with the FQDN, the gateway will communicate with Central.

The firmware version defined in firmware compliance for the group will be enforced and an upgrade pushed if necessary; once an upgrade is completed the gateway will reboot.


The gateway will initialize using the specified AOS 10 version, contact Central, and be configured based on the configuration applied in Central to the device.

Once the configuration has been successfully applied, the gateway will be up and operational in Central.

Confidential - For Training Purposes Only

Web Based OTP

<https://172.16.0.254>



The image shows the 'Aruba Controller Setup' web interface. At the top, the Aruba Networks logo is visible. Below it, a progress bar shows three steps: 'Provisioning Method' (current), 'Activate', and 'Summary'. The main section is titled 'How will the controller connect to Activate/Central?'. It has two radio buttons: 'Static IP Address' (selected) and 'PPPoE'. Below this, there are several input fields: 'VLAN ID' (1), 'Interface' (0/0/0), 'Port mode' (Access), 'IP address' (10.1.3.100), 'Network mask' (255.255.255.0), 'Default gateway' (10.1.3.1) with '(optional)' text, and 'DNS server' (10.254.1.21) with '(optional)' text. A large diagonal watermark 'For Training Purposes Only' is overlaid on the form. In the bottom right corner, it says 'MOD 3-17'.

Here are the steps involved with OTP via web browser connection to an Aruba gateway:

Connect a computer to the GE 0/0/1 Ethernet port on the Gateway; an IP address will be offered by DHCP in the 172.16.0.0/24 network; open a web browser and navigate to <https://172.16.0.254>, proceeding past the warning for the invalid SSL certificate.

Select By connecting to activate/central then click Next.

Select Static IP Address or PPPoE as the connection method and then provide the required information; the below example provisions a gateway to use the GE 0/0/0 port and a static IP address; a gateway running AOS 10 has additional options for a trunk port and port-channel.

Verify the information is correct then click Deploy and Reboot.

Once the initial provisioning is complete the gateway will resolve the Aruba Activate FQDN and communicate with Activate for further provisioning:

- If the gateway is new and has not been previously provisioned, Aruba Activate will push a Central enabled firmware upgrade and reboot the gateway.
- Activate will provision the gateway with the FQDN for the assigned Central instance.

After booting to a Central enabled firmware and provisioned with the FQDN, the gateway will communicate with Central

The firmware version defined in firmware compliance for the group will be enforced and an upgrade pushed if necessary; once an upgrade is completed the gateway will reboot.

The gateway will initialize using the specified AOS 10 version, contact Central, and be configured based on the configuration applied in Central to the device.

Once the configuration has been successfully applied, the gateway will be up and operational in Central.

Confidential - For Training Purposes Only




Initial Gateway Setup

Now that you've deployed your gateways, you are ready to perform their initial setup.

Confidential - For Internal Use Only


Gateway as L2 or L3 Device

Layer 2 mode



Typical for campus deployment

Layer 3 mode



Typical for branch gateway deployment


MOD 3- 19

You can configure the gateway as a switch with VLANs, or as a router with IP addresses and routes. When the default gateways (DG) for VLANs 10-12 are on an uplink L3 router, the gateway will use L2 switching to move user traffic to the upstream L3 router.


When the L3 switch defines VLANs on the gateway with IP addresses, they will act as the DGs for VLANs. Therefore, all user traffic will route via gateways to the uplink.

Aruba typically recommends that you deploy your gateway as a L2 switch in a campus environment. Layer 3 deployment typically applies to a branch gateway environment.

Setup Port-Channel / LACP



Aruba recommends setting up
port channel at the device level
Port-channel = LAG



Aruba recommends setting up a port channel (LAG) at the device level. This allows the gateway to be moved to another Central group without losing the port channel configuration and possible broken connectivity. A gateway port channel is equivalent to a LAG on the switch.

Important: Make sure to select the Trust checkbox in the Central configuration!

Setting NTP

The screenshot shows the Aruba Central interface for a Gateway (GW1). The left sidebar contains navigation options: Overview, WAN, LAN, Device, Clients, Applications, Security, Alerts & Events, Audit Trail, Tools, Reports, and Maintain. The main content area is titled 'Gateway' and has tabs for System, Interface, Security, and Routing. Under the 'System' tab, there are sub-tabs for General, Admin, and Certificates. The 'General' sub-tab is active, and under 'Basic Info', the 'Clock' option is selected. A table titled 'NTP servers' is visible at the bottom, with columns for IP ADDRESS/FQDN and BURST. A plus icon is at the bottom right of the table.

The screenshot shows the 'Add NTP Server' configuration page. It includes the following fields and options:

- IPv4/IPv6/FQDN: IPv4 (selected)
- IPv4 address: 10.254.1.21
- Burst mode: ☐
- Authentication key: (empty text field)
- Source interface: 1 (selected)
- NTP server VLAN: Choose an option- (dropdown menu)
- Use NTP authentication: ☐
- Time zone: United States: America/Detroit (UTC-0...)

To setup system time on the gateway, you can use NTP. Navigate to your Central group (or device). From there click the Config icon and navigate to System -> General -> Clock. Enter the NTP server IP address and time zone.

MOD 3- 21
+




Introduction to Gateway Cluster

Improve resiliency with gateway clusters. The clustering technology was introduced in ArubaOS 8 with the mobility controllers and is still one of the main reasons customers choose Aruba gateway solutions over other vendors or using standalone Aruba APs in bridge mode. This section provides a quick introduction to clustering. You will learn more about this powerful feature later in the course.

Gateway Cluster

Seamless roaming
Client/AP load balancing
Hitless failover
Client state sync
Ease of deployment



MOD 3- 23

A gateway cluster combines multiple Aruba gateways operating as a single entity to provide high availability and service continuity to the WLAN clients in a network. Gateway clusters provide full redundancy to access points (APs) and WLAN clients in the event of a failover. AOS 10 supports gateway clusters and provides the following features and benefits:

Seamless roaming: when a client roams between APs within the WLAN boundary, the clients remain anchored to the same gateway in a cluster for a seamless user experience.

Client/AP Load Balancing: if there is excessive workload among the devices, the gateways in a cluster balance AP and client loads to ensure faster connectivity to clients.

Hitless Failover: when a gateway device fails, APs and endpoints failover to another gateway in the cluster without any service disruption.

Client State Sync: synchronizes high-value client sessions such as voice, video, FTP and IGMP (used for IPv4 multicast groups) and MLD (used for IPv6 multicast groups) between active and standby gateways in a cluster. This allows the connected devices to failover to the standby gateway seamlessly. However, AOS 10 only supports this synchronization for the first failover. For subsequent failovers, it will not synchronize the session to the recent standby; instead, it synchronizes a new session.

Ease of deployment: a gateway cluster is automatically formed when assigned to a UI group

in the Network Operations app without any manual configuration.

Confidential - For Training Purposes Only



Gateway Cluster: Design Considerations

Homogeneous cluster

- All nodes are of the same gateway series & model
- 7200 series support 12 nodes
- 9000 and 7000 series support 4 nodes
- AP load should be 50% of cluster capacity

Heterogeneous cluster

- Cluster of Gateway same series different model
- 7200/7000 combination supports 4 nodes
- AP load should be lowest value of worst-case scenario load
- Not recommended



AOS 10 supports multi-version between AP and GW

MOD 3- 24

There are two types of clusters you can implement:

Homogeneous

Heterogeneous

The following two sections will cover these cluster types.

Note: Unlike AOS 8, where all the APs and gateways had to be running the same version of firmware, ArubaOS 10 supports different firmware version between APs and gateways; i.e., the AP runs 10.3.1.1 and gateway is running 10.3.1.2, for example.

Homogeneous Cluster

This type of cluster is built with all nodes of the same Aruba gateway series. An Aruba 7200 series gateway homogeneous cluster supports up to 12 nodes in a cluster, whereas an Aruba 7000 and 9000 series homogeneous cluster supports only four gateway nodes. A homogeneous virtual gateway cluster is also supported with four nodes.

The cluster sizing depends on the number of cluster AP count required to ensure that every AP has a DDG and S-DDG with adequate capacity for all APs to failover. The recommended AP load of this cluster should be half of the total cluster capacity. Therefore, the cluster AP

count should be equal to 50% of the cluster capacity.

For example, if a cluster is made up of four 7220 managed devices, the combined capacity of four 7220 managed devices is 4,096 APs, hence, the AP count would be 2,048.

Heterogeneous Cluster

Allows you to combine different gateway models. A heterogeneous cluster with a combination of 7200 series and 7000 series gateways supports up to four nodes.

Cluster AP size should be equal to the lowest value of either 50% of total cluster capacity or the worst-case scenario load. The worst-case scenario load is the AP load handled by the remaining nodes in a cluster in the event of highest capacity cluster member going down.

The AOS 10.x supports multiple firmware versions between gateways in the same cluster profile by exchanging messages between gateways in the cluster.

Note: Heterogeneous clustering is not recommended by Aruba.

Confidential - For Training Purposes Only

Automatic Cluster Mode

The screenshot shows the configuration page for a gateway device named 'P24-T01-Gateway...'. The left sidebar contains navigation links: 'Manage', 'Overview', 'Devices', 'Clients', and 'Guests'. The main content area is titled 'Gateways' and has tabs for 'System', 'Interface', 'Security', 'Routing', and 'High Availability' (which is selected). Under the 'High Availability' tab, there are sub-tabs for 'Clusters' and 'Redundancy'. The 'Clusters' sub-tab is active, showing an 'Automatic' toggle switch that is turned on. Below the toggle are two radio button options: 'Auto Group' (selected) and 'Auto Site'. An orange box labeled 'Default' points to the 'Auto Group' option. An orange box labeled 'Warning' points to a warning message box on the right. The warning message box contains a red triangle icon and the text: 'Existing auto-group clusters will be dismantled and gateways will be remapped to new site specific auto-clusters. Do you want to save and continue?'. The bottom right corner of the warning box is labeled 'MOD 3-25'.

Gateways

System Interface Security Routing **High Availability**

Clusters Redundancy

Automatic: ☒

☒ Auto Group ☐ Auto Site

Warning

Existing auto-group clusters will be dismantled and gateways will be remapped to new site specific auto-clusters.

Do you want to save and continue?

MOD 3-25

Unlike AOS 8, where you had to perform multiple configuration steps on your Mobility Conductor to establish a cluster, using AOS 10 and Aruba Central, the automatic cluster option is enabled by default, and the Auto Group mode is selected. If you wish to use the Auto Site mode, the system will display a warning message indicating that the existed cluster will be dismantled.



Knowledge Check

Self-check on key learning points



It's time for a knowledge check to review what you have learned.

Confidential - For Training Purposes Only

Question #1

You moved a gateway from the Aruba Central default group to the “East-Branches” group. After this you notice that much of the original configuration for that device has changed. Why? (Select two.)

- a. When you move a device to a new group, it inherits the new group configuration.
- b. When you move a device to a new group, its configuration reverts to the defaults.
- c. When you move a device to a new group, device level configuration overrides are retained.
- d. When you move a device to a new group, missing variables in the variables file can cause issues.

a

Knowledge Check

MOD 3-27

+

a, c

Question #2

Your colleague is asking about the difference between gateway deployments that use zero-touch provisioning (ZTP) and one-touch provisioning (OTP). Which statements below would improve your colleague's understanding? (Select two.)

- a. ZTP is great for when remote installers have little or no knowledge about gateway configuration.
- b. Both ZTP and OTP require that the gateway be connected to a network that provides DHCP services.
- c. You can leverage both ZTP and OTP using the console port.
- d. With OTP, you must use a serial console port or web browser to configure the gateway, just enough for it to communicate with Activate and Central.

a

Knowledge Check

MOD 3-28

+

a, d

Confidential - For Training Purposes Only



Lab Activity

Let's review the lab activity that is associated with this module.

.

Lab Tasks

Lab 03.01- Deploying APs

Task 1.1

- Deploying APs

Lab 03.02



MOD 3- 30

The slide provides a brief review of tasks performed in this lab (3.01). Please see your Lab guide for details of this lab.

Lab Tasks

Lab 03.02 - Deploying Gateways

Task 1

- Prepare the wired network to connect the GWs

Task 2

- Prepare the wired network to connect the GWs

Task 3

- Central Group GW config

Task 4

- Monitor config pushes

Task 5

- Moving a gateway to another group in Aruba Central

Task 6

- Troubleshooting

Lab 03.03



MOD 3- 31

The slide provides a brief review of tasks performed in this lab (3.02). Please see your Lab guide for details of this lab.

Lab Tasks

Lab 03.03 - Gateway Clustering

Task 1

- Review the auto cluster

Module 4 -
Tunneled
WLAN
Architecture



MOD 3- 32

The slide provides a brief review of tasks performed in this lab (3.03). Please see your Lab guide for details of this lab.

Rev. 23.11




Tunneled WLAN Architecture

Implementing Aruba Campus Access



Welcome back everyone! This is module 4: Tunneled WLAN Architecture. In this module you'll learn how to set up and monitor tunneled node connectivity between Aruba APs and gateways.

Confidential For Training



Objectives

- Describe tunnel mode WLAN operation and use cases
- Leverage tunnel mode to improve security and ensure a consistent end user experience
- Assess tunnel mode status and perform basic diagnostics
- Explain gateway clustering operation and advantages
- Assess tunneled client status and connectivity

MOD 4- 2



After completing this module, you will be able to:


- Describe a tunnel mode WLAN operation and use cases.
- Leverage tunnel mode to improve security and ensure a consistent end user experience.
- Assess tunnel mode status and perform basic diagnostics.
- Explain gateway clustering operation and advantages.
- Assess tunneled client status and connectivity.

Tunneled Mode WLAN

Let's begin by exploring what a tunneled mode WLAN is and how it works.

Confidential - For Training

Deployment with Gateways



Forwarding modes

- Tunnel mode
- Mixed mode
- Microbranch mode (VPNC)

AP - gateway communication

- Version independent
- IPsec for control plane
- **Single** GRE for data plane

Authentication and encryption

- AP is the authenticator
- Gateway act as a RADIUS proxy
- APs **always** do encrypt/decryption

MOD 4- 4

Deploying Aruba gateways in the LAN infrastructure enable tunnel and mixed forwarding modes for WLANs. If case gateways are deployed as VPN concentrators, then give the possibility to enable microbranch mode.

Since AOS 10 APs are anchored to Aruba Central and not the gateways, this allows version independency; in other words, APs and gateways does not require to run the same AOS 10 version, although it is best practice to do so. IPsec tunnels between APs and gateways are orchestrated, these tunnels are used to carry control plane traffic (PAPI). On the other hand, the data plane uses GRE tunnels to carry endpoints' traffic. In AOS 10 both clients and APs complete wireless packet encryption and decryption, meaning the user data is carried unencrypted in the GRE tunnel from the AP to the gateway and vice versa. For this reason, Aruba has added the possibility to protect the data plane using GRE over IPsec tunnels.

When 802.1X authentication is utilized in these scenarios, the AP plays the authenticator role and the gateway functions as a RADIUS proxy. This is necessary to ensure gateways are aware of the authenticated user and their details and any changes in policies—RADIUS change of authorization (CoA)—send from the authentication server to the authenticator (AP).

AP to Gateway Connections

Control Plane

- AP-GW Heartbeats
- Cluster information
- RADIUS
- Roaming assistance
- Orchestrated tunnel



IPsec



Data Plane

- Client's traffic
- **Single tunnel for all SSIDs**



GRE



AOS 10 GRE tunnel setup different from previous AOS versions

MOD 4-5

There are two connections between the APs and gateways: one for the control plane and the other for the data plane. The Aruba proprietary Protocol Application Programming Interface (PAPI) comprises the control plane function. This protocol helps AOS 10 APs and gateways to exchange important operation information such as:


- Heartbeats between the APs and gateways.
- Cluster information, which includes the bucket map and node list distribution (you'll learn more about this later in the course).
- RADIUS communication: authentication, authorization, accounting, and CoA information.
- Roaming assistance including session synchronization requests.

One problem with PAPI is that this protocol does not provide privacy nor integrity and therefore is vulnerable to attacks. To secure this connection both Aruba devices negotiate a secure IPsec tunnel. The IPsec tunnels between the AP and Gateway cluster are orchestrated by Tunnel Orchestrator for LAN Tunnels service in Aruba Central.

The second connection relates to the data plane or clients' traffic. Like previous architectures, AOS 10 uses GRE tunnels for this purpose. In AOS 10, a single GRE tunnel is created for the tunnel- and mixed-based SSIDs.

Tunnel Forwarding Mode

WLAN **Employee**
Tunnel mode
802.1X
VLAN 11




MOD 4- 6

The figure in the slide shows a client's traffic is received by the AP and then is redirected (tunneled) to the gateway cluster. Gateways analyze and process client's traffic. From the rest of the network's perspective, the client's traffic is generated on the gateway and not on the AP. As we discussed previously, this solution uses a GRE tunnel between the AP and the gateway; therefore, it is important to permit IP protocol number 47 between these two devices on any intermediate filtering or firewall devices.

Notice that the client receives an IP address in VLAN 11 and this VLAN must be defined and configured on the gateway side, as well should be permitted on the link from the gateway to the campus switches. Note that VLAN 11 is not required on the link between the AP and the switch, since the AP is tunneling the VLAN traffic to the gateway. Basically, GRE is acting as an overlay and providing a logical connection.

Overlay Tunnel Orchestrator



No controller discovery mechanism necessary

MOD 4- 7


The communication between an AP and gateway uses an IPsec tunnel to protect the control plane traffic (PAPI). As we mentioned before the creation of this tunnel is orchestrated. This means that the AP and the gateway establish a communication channel with the Overlay Tunnel Orchestrator (OTO) service in Central.

The OTO service internally negotiates IPsec Phase 1, which implies that this phase does not run directly between the Aruba AP and gateway anymore (unlike in AOS 8). If you try to issue the `show crypto isakmp sa` command on the gateway, for example, you will receive no output. Once Phase 1 is completed, the OTO service shares the key material, tunnel specifications and instructs the AP and Gateway to run IPsec Phase 2 between each other; this means that the `show crypto ipsec sa` command on either device displays an entry for this tunnel.

Note: Unlike earlier AOS versions, AOS 10 doesn't require a gateway discovery protocol or method (where the APs try to find the gateways), since all AP-to-gateway communication is now orchestrated by Aruba Central.

Client Connecting to Tunneled Open SSID

Single GRE tunnel used by all clients connected to AP




With AOS 10 user traffic inside GRE tunnel is **802.1Q tagged!**

MOD 4- 8

The AOS 10 architecture implements a single GRE tunnel that will be used by all wireless clients connecting to the AP (it doesn't matter how many radios or WLANs are broadcasted by the AP). The traffic in this tunnel is sent tagged and the VLAN tag is used to communicate the VLAN ID to the gateway. This will process be explored in more depth later in this module.

Tunnel Forwarding Mode



MOD 4- 9

Unlike in AOS 8, the AP in the new architecture (AOS 10) always encrypts and decrypts wireless traffic for both tunnel and mixed modes. The following sequence helps to understand this process.

1. If WPA3 (Personal or Enterprise) was enabled on the SSID, the wireless client sends an encrypted wireless frame to the AP.
2. The AP receives and decrypts the wireless frame.
3. The AP forwards the unencrypted traffic to the gateway using the GRE tunnel.
4. The response is sent without any encryption from the gateway to the AP inside the GRE tunnel.
5. The AP receives and encrypts the frame to send to the client.
6. An encrypted frame is sent to the client.

As all traffic between AP and gateway is sent unencrypted, this could bring about a security concern. Aruba has implemented a solution for this scenario (GRE over IPsec) that will be discussed in a moment.

AP Tunnel Agent

```
AP-1# show ata endpoint
```

```
ATA Endpoint Status
```

UUID	IP ADDR	STATE	TUN DEV	TUN SPI (OUT/IN)
3afc45c3-002	192.168.101.240	SM_STATE_CONNECTED	tun0	7435200/ba8bea00

```
Total Endpoints Count: 1
```

```
<continuation of the command output>
```

```
ATA Endpoint Status
```

PORT(SRC/DST)	VALID TIME(s)	TYPE	GRE VLANs	HBT(Jiff/Missed/Sent/Rcv)	INNER IP
4500/4500	45567	GRE	1,11	102632/0/86657/83860	10.1.10.1


```
Total Endpoints Count: 1
```



MOD 4- 10

From the AP CLI, you can check the status of the tunnels between the AP and gateway using the show ata endpoint command, shown in the slide. Notice that the type of tunnel created is GRE and there are VLANs supported on this tunnel (1 and 11).

Data Encryption




MOD 4- 11

GRE over IPsec is an optional feature protects the traffic between the AP and the gateway:

1. The wireless client sends an encrypted wireless frame to the AP.
2. The AP receives and decrypts the wireless frame.
3. The AP forwards the now unencrypted traffic to the gateway using a GRE over IPsec tunnel. Since IPsec is employed, effectively this packet is re-encrypted before forwarding.
4. The unencrypted response is sent from the gateway to the AP inside the GRE over IPsec tunnel. Again, since IPsec is employed, this is effectively encrypted before forwarding to the AP.
5. The AP receives encrypted IPsec packet, decrypts it, and then encrypts the frame for the wireless client.
6. An encrypted frame is sent to the client.

Data Encryption Configuration



To configure secure data tunnels between APs and a gateway cluster, complete the following steps for a Central group:

1. In the Network Operations app, set the filter to a group that contains at least one AP. The dashboard context for the group is displayed.
2. Under Manage, click Devices > Access Points.
3. Click the Config icon. The tabs to configure APs are displayed.
4. Click Show Advanced, and click the Security tab. The Security details page is displayed.
5. Enable Data Handling.
6. To enable IPsec tunnel for data traffic, turn on the Data Encryption slider button.
7. Click Save Settings.

The Data Encryption toggle button is disabled by default. When this toggle button is enabled, the AP sends client traffic to gateway via GRE over IPsec. When this toggle button is disabled, the AP sends client traffic to gateway via GRE only.

GRE over IPsec Tunnel

```
AP-1# show ata endpoint
```

```
ATA Endpoint Status
```

UUID	IP ADDR	STATE	TUN DEV	TUN SPI (OUT/IN)
3afc45c3-002	192.168.101.240	SM_STATE_CONNECTED	tun0	7435200/ba8bea00

```
Total Endpoints Count: 1
```

```
AP-1# show ata endpoint
```

```
ATA Endpoint Status
```

PORT (SRC/DST)	VALID TIME (s)	TYPE	GRE VLANs	HBT (Jiff/Missed/Sent/Rcv)
4500/4500	45567	GREoIPsec	1,11	102632/0/86657/83860

```
Total Endpoints Count: 1
```



MOD 4-13

After enabling GRE over IPsec, you will see that the tunnel type has changed to GRE over IPsec when viewed from the AP's CLI (see the example output in the slide). The top part shows GRE-only and the bottom part shows GRE over IPsec.

Connected Client Status

AOS 10 AP CLI:

```
Campus-AP1# show clients
```

Client List

Name	IP Address	MAC Address	ESSID	Access Point	Role
----	-----	-----	-----	-----	----
1c4d7007fd14	172.16.1.8	1c:4d:70:07:fd:14	Guest1	Campus-AP1	Guest1

AOS 10 gateway CLI:

```
(GW1) *#show user-table
```

Users

IP	MAC	Name	Role	Connected To	Profile	Forward mode
-----	-----	-----	----	-----	-----	-----
172.16.1.8	1c:4d:70:07:fd	1c4d7007fd14	Guest1	fc:7f:f1:ce:95:44	Guest1_#1656	dtunnel

User Entries: 1/1

Curr/Cum Alloc:1/2 Free:0/1 Dyn:1 AllocErr:0 FreeErr:0

MOD 4- 14

You can verify status of the connected client on the AP and/or the gateway. The show clients command is applicable to the AP and the show user-table command applies to gateway.

Note: Remember that you can have Central execute these commands from your web browser and send the resulting output back to your web browser: you don't need CLI access to either of these two devices.

Configuring a Tunneled SSID

Access Points

Summary List Config

WLANs Access Points Radios Interfaces Security Services System IoT Configuration Audit

Hide Advanced

CREATE A NEW NETWORK

1 General 2 VLANs 3 Security 4 Access 5 Summary

Name (SSID): P24-T01-Tunnel

> Advanced Settings



MOD 4- 15



To configure a Tunneled SSID use the following steps in Central:

1. Clear the context filter and select the group where APs have been mapped to.
2. Select Devices in the navigation pane and then click on the Config icon.
3. Under WLAN tab, click Add SSID. The WLAN wizard appears.

In the first step of the wizard, you must enter a name for the WLAN or SSID. Usually, default parameters for a WLAN work in most cases; however, an administrator can modify other properties of the WLAN, these options are available under the Advanced Settings button.

Configuring a Tunnel SSID

CREATE A NEW NETWORK

1 General 2 VLANs 3 Security 4 Access 5 Summary

Traffic forwarding mode: ☐ Bridge ☒ Tunnel ☐ Mixed

Primary Gateway Cluster: aos10-gw:auto_gwcluster_186_0

Secondary Gateway Cluster: None

Client VLAN Assignment: ☒ Static ☐ Dynamic

VLAN ID: 11(11)

VLAN(s) assigned to SSID

MOD 4- 16

In the second step of the wizard, you must define the forwarding mode: in this case you select Tunnel. You must specify the primary gateway cluster, but the secondary cluster is optional.

Finally, you must specify the VLANs that you would like to place the traffic. Two options are available—Static and Dynamic.

Static. This option means that the administrator manually defines what is the VLAN-ID used to tag client's traffic. This VLAN must be present on the gateway, the next hop (aggregation or core switches, which act as the default gateway for the VLAN) and the link between the gateway and the aggregation/core switches.

Dynamic. If you use this option, the authentication process dictates what VLAN is used. This option permits to run a personalized VLAN mapping for each user. For example, Alice's laptop runs the authentication process and then the gateway starts tagging her traffic using VLAN 11, which was assigned by the AAA server during the authentication process. A few minutes later Bob's smartphone connects and for him the AP tags his traffic on VLAN 12, also dictated by the AAA server.

Configuring a Tunnel SSID

1 General 2 VLANs 3 Security 4 Access 5 Summary

Security Level:

Security levels

Enterprise Personal Visitors Open

Key Management:

WPA3-Enterprise(CCM 128)

Primary Server:

clearpass

Secondary Server:

-- Select --



MOD 4- 17



There are no differences between the bridged, tunneled and mixed SSID's under the wizard's Security tab. Select the desired authentication method and click Next. Keep in mind, when you select enterprise authentication, the primary authentication server configured here will be replicated on the gateway's configuration within the cluster that you selected on the VLANs tab. In other words, the RADIUS configuration is pushed down to both the APs and the gateways.

Configuring a Tunnel SSID

1 General 2 VLANs 3 Security 4 Access 5 Summary

Access rules

Role Based Network Based Unrestricted

ACCESS RULES FOR SELECTED ROLES

- Deny youtube to all destinations + log
- Allow any to all destinations

This step in the wizard helps administrators to indicate what are the permissions that will be assigned to endpoints that are connected to the WLAN that you created.

Unrestricted: This option allows full access to the network.

Network based: This option permits to create restrictions that will be assigned to all the devices connected to the WLAN regardless of who or what device is connected to the network.

Role based: This option permits to assign individual restrictions to each device connected to the network. Based on user's identity, the AP applies the correct restrictions. This option requires an authentication process.


Clustering Continued



Let's explore gateway clustering in more detail.

Confidential - For Training

Gateway to Gateway Tunnels



- Cluster control plane protected by IPsec
- Tunnels are not orchestrated by OTO
- Use factory device TPM certificate
- Full mesh established for clustering
- Cluster configuration is a part of configuration file




MOD 4- 20

Aruba AOS 10 supports gateway clustering. This mechanism brings several advantages to a network deployment, which includes high availability, seamless roaming, and AP and client load balancing. Cluster members must establish a full connectivity with all other members of the cluster using IPsec tunnels, where these tunnels are used as part of the cluster control plane. It is important to clarify that these tunnels are not orchestrated by the OTO service but built by the gateways themselves.

Cluster IPsec tunnels are created using the Trusted Platform Module (TPM) certificates also known as factory certificates, on the gateways themselves.

Cluster Leader Election

1. Highest priority
2. Highest platform
3. Highest MAC address



Member



Platform: **7030**
MAC Address: **aa**
Priority: **128**

Member



Platform: **7030**
MAC Address: **bb**
Priority: **128**

Leader



Platform: **7030**
MAC Address: **cc**
Priority: **128**



MOD 4- 21


During the cluster creation, the cluster must elect one of the gateways as leader. In the election process, the highest priority value has precedence. If all the gateways are using the same priority, the next criteria is the platform: the node having the highest capable platform is elected as leader. The last criteria is the MAC address, where the gateway with the highest system MAC address wins the election. Who becomes the leader is typically not a concern; however, one of the gateways must perform this role.

Cluster Leader Duties

Assigns DDG and S-DDG

Builds and publishes bucket map to other GWs

Load balance



The leader has some important responsibilities. The first one is to assign a Device Designated Gateway or DDG to each AOS 10 AP. The DDG is needed to achieve multicast functionality in the WLAN. It is important to clarify that in AOS 10, APs do not anchor to a gateway: the APs are anchored to the cloud (Aruba Central). The cluster leader assigns the DDG and S-DDG (standby) for the AP in a round-robin fashion based on the current AP load on all gateways in the cluster.

In the example in the slide, AP1 boots and contacts the cluster. The cluster leader (GW-3) assigns GW-2 as the active DDG and GW-1 as the standby device.

The second responsibility of a leader is to generate and publish a bucket map for the cluster. This is a simple table which is first shared to the gateways and then to the APs or switches (dynamic segmentation). When an endpoint is connected to the WLAN, the AP uses the bucket map to select a gateway from the cluster to tunnel endpoint's traffic. The selected gateway is known as the User Designated Gateway (UDG).

In case the cluster becomes unbalanced, like a cluster gateway member is serving more clients (users) compared to the rest of the members, then a load balancing process is automatically triggered by the leader that will redistribute the clients (users).


DDG Duties

1. Distribute bucket map to APs

3. RADIUS proxy

2. Distribute node list to APs

4. IGMP proxy




MOD 4- 23

The DDG is responsible for publishing a bucket map and node list (a list of members in a cluster) to every device (AP and switches). The DDG selection sends the AP node list and bucket map from the Gateway to the AP. This selection also configures the VLAN multicast table during AP bootstrapping and cluster failover. In addition, the DDG and S-DDG are responsible for performing IGMP proxy and RADIUS proxy for the APs and clients.

Note: In AOS 10.x, the APs do not anchor to a gateway, but rather anchor to the Aruba Central cloud. The gateway identifies the AP when it boots. However, the DDG, which is an AP-designated gateway, is needed to achieve multicast functionality. From the user point of view, the APs would still be aware of the cluster bucket map that is used to direct the user traffic to the user designated gateways (UDG).


Client to Gateway Mapping



MOD 4- 24

For each client connected to AP there will be a user designated gateway (UDG) selected as well as a backup gateway (S-UDG). The AP will send the client's traffic via GRE tunnel to that gateway and that gateway will actively process the client's traffic. If UDG fails, then the client will failover to the standby UDG (S-UDG). With proper configuration this process is hitless and the client will not even notice that there was a failure.

Client Load Balancing




When a client associates to an AP, its Wi-Fi MAC address runs through a hashing algorithm. This produces a decimal number between 0 and 255. The decimal number acts as an index to a mapping table (the bucket table referenced earlier in the module). This table indicates what the client's UDG and S-UDG gateway are.

The cluster leader computes the bucket map and push this to all the APs in the clustering domain. When a client associates to an AP, the AP runs the client's MAC address through a hashing algorithm. The hash result is an index value which is used to select the gateway from the bucket map table. The selected gateway becomes the User Designated Gateway (UDG). For high availability purposes a secondary UDG (S-UDG) is selected as well.

As an example, Client1 connects to the WLAN, the AP uses the hash process to select a UDG, in this case GW-1, where the secondary UDG is GW3. For Client2 the hash process determines the UDG is GW-2 and the secondary UDG is GW-1. Initially the load might not be distributed, but based on more and more clients connecting, the load becomes more evenly distributed for the client connections.

AP to Gateway Cluster Tunnels



A single IPsec & GRE tunnel to every gateway

MOD 4- 26

AOS 10 APs will establish IPsec and GRE tunnels to every gateway in a cluster even though only one of gateways is selected as DDG and another one as S-DDG. This is very different from previous AOS firmware versions.

Cluster VLAN Probing

Trigger: cluster setup, add/remove VLANs


Sent to all members on every VLAN

If VLAN probe succeeds clients will be synced for that VLAN

If VLAN probe fails clients will not be synced for that VLAN

L2 unicast packet (EtherType 0x88b5)

```
interface 1/1/1
vlan trunk allowed 1,11
vlan trunk native 1
!
interface 1/1/2
vlan trunk allowed 1,11,12
vlan trunk native 1
```



MOD 4- 27

Once the cluster members form a full mesh, gateways begin the VLAN probing process. Each cluster member sends a specially-crafted L2 broadcast frame in each configured VLAN on the controller. VLAN probing is bidirectional between every pair of cluster members. If a cluster member receives a probe acknowledgment from its cluster peer, for all configured VLANs, it changes its connection state with that peer to L2-connected. If a probe fails, that VLAN will be automatically excluded from the cluster; that way the cluster remain in L2 connected state.

NOTE: Cluster members probe on every configured VLAN. This is important for configuration and troubleshooting efforts. This was very different from AOS 8, where you had to manually determine which VLANs to exclude and manually exclude this VLANs. With AOS 10, this is done for you automatically!

Cluster Member Connections

VLAN probe successful	VLAN probe unsuccessful
Cluster is fully redundant	Cluster is not fully redundant
APs & clients are fully synced	Only APs are fully synced
Users are fully synced between Active & Standby UDG nodes	Users are not synced between Active & Standby UDG nodes
High value sessions are synced	High value sessions are not synced
Users are not de-authed on a failover, even from active to standby UDG	Users are de-authed on a failover, even from active to standby UDG. Users must do a full 802.1X authentication



MOD 4- 28

There are several differences that exist between L2- and L3-connected cluster members. L2-connected members are fully redundant and enjoy full replication both of AP and client sessions: L2-connected supports stateful failover. L3-connected members, in contrast, only fully replicate AP sessions (stateless failover).

L2-connected cluster members fully synchronize the active UDG and standby UDG nodes. L3-connected clusters cannot synchronize these nodes or high-value sessions.

Failover events do not de-authorize users of L2-connected cluster members, even from Active UDG to Standby UDG. Failover events for L3-connected clusters, however, will de-authorize users. This means that L-3 connected cluster users must perform a complete 802.1X re-authentication after they re-associate to the AP, which means that if the user was running an application(s) during this process, the user would probably have to restart that application. Obviously, based on this information, you want to ensure that the gateways in a cluster are L2-connected!

Cluster Heartbeats


Cluster members exchange heartbeats

PAPI packet routed outside of IPsec tunnel

Round trip delay computed for each request/response

Failure detection based on no response over set time

Last Heartbeat response 69ms ago,
RTD = 0.500 ms



MOD 4- 29

Cluster members also send periodic heartbeat messages to each other, to speed up failure discovery. The heartbeat scheme is dynamic and considers the round-trip delay (RTD) between each pair of cluster members. If a cluster member fails to exchange a heartbeat packet, it will be considered down, which will trigger a cluster re-balance of APs and clients to the remaining cluster members that are up.


Cluster Load Balancing

Homogeneous cluster



- Equal number of buckets always assigned to cluster members

Heterogeneous cluster



- Buckets are assigned in the ratio of members capacity

Rebalance when cluster size changed (GW join or GW split)

No rebalancing on stable cluster



MOD 4- 30

Gateways complete AP load balancing during the initial setup. In either a homogeneous or a heterogeneous cluster, the APs are load balanced in a round-robin manner among gateways depending on the platform AP capacity. In a homogeneous cluster, an equal number of buckets is assigned to each gateway, and APs are equally distributed evenly across the cluster. In a heterogeneous cluster, tokens are assigned based on the member capacity, meaning that APs will be balanced based on the system load on each gateway. System load is compared by the percentual usage of the system (the number actual connections versus the number of connections supported by the gateway model).

The APs are distributed to offload multicast handling evenly across cluster peers. Only active and standby DDGs are configured for the multicast tunnel to the user VLAN in the data path.



Client Monitoring

Let's keep tabs on those clients and examine how you can verify your clustering operation in Aruba Central.

Confidential - For Training

Tunneled Clients

CLIENTS

ALL

470.37 KB (189.77 KB | 280.60 KB)

All3

Connecting0

Connected2

Failed0

Offline1

Blocked0

Wireless3

Wired0

Remote0

CLIENTS

Client Name	Status	IP Address	VLAN	Connected To	SSID/Port	Tunneled
Batmans-iPad	Connected	10.1.11.3	11	AP-1	P24-T01-Tunnel	Yes
AlvarosMBP1640	Connected	10.1.11.5	11	AP-1	P24-T01-Tunnel	Yes



MOD 4- 32



Use the context filter in Central to narrow your scope, and then in the Navigation pane, select Clients to view your connected clients. You can see a column labeled Tunneled, indicating that tunneled mode is used when set to Yes.

Tunneled Client Details

CLIENT DETAILS Actions ▾ Go Live

DATA PATH

CLIENT

USERNAME	8c8590766c95
HOSTNAME	AlvarosMBP1640
IP ADDRESS	10.1.11.5
GLOBAL UNICAST IPV6 ADDRESS	--
CLIENT CATEGORY	Computer
CLIENT OS	Mac OS
CLIENT TYPE	Wireless
MAC ADDRESS	8c:85:90:76:6c:95
LINK LOCAL IPV6 ADDRESS	fe80::14b1:f033:436d...
CLIENT FAMILY	Apple Mac
CONNECTED SINCE	Feb 08, 2022, 18:41:29

NETWORK

VLAN	11	VLAN DERIVATION	VSA
AP ROLE	P24-T01-Tunnel	AP DERIVATION	RADIUS
GATEWAY ROLE	P24-T01-Tunnel	SWITCH ROLE	--
SEGMENTATION	OVERLAY		
AUTH SERVER	192.168.101.240	DHCP SERVER	10.1.11.254
TUNNELED	Yes	TUNNELED ID	0

CONNECTION

CHANNEL	128 (80 MHz)	BAND	5 GHz
CLIENT CAPABILITIES	802.11ac		
CLIENT MAX SPEED	1.30 Gbps		
LEDs on ACCESS POINT (AP-1)	○○○Blink LEDs		

MOD-4-33

When monitoring a specific tunneled user (clicking a specific client name from the previous slide), notice that the Aruba Central graph shows data path at the top of the window: the client connected to an SSID, which is mapped to an AP, that is connected to a switch, and lastly to a gateway. In addition, notice the overlay tunnel label between the switch and the gateway, indicating that the user's traffic is being tunneled to the gateway.

In the Network tile, notice the Tunneled flag is set to "Yes," confirming that the user is being tunneled. You'll also see that, under the Gateway icon in the topology, this is the Gateway name that is elected to be the User Designated Gateway (UDG).



Knowledge Check

Self-check on key learning points



It's time for a knowledge check to review what you have learned.

Confidential - For Training Purposes Only

Question #1

Which statements below accurately describe tunneled mode WLANs? (Select three.)

- a. Two tunnels are created between each AP and the gateway: one for the control plane and one for the data plane
- b. Both control and data plane functions are handled by the Protocol Application Programming Interface (PAPI)
- c. By default, only control plane traffic is protected by IPsec

a

Knowledge Check

MOD 4- 35

+

a,c

Confidential - For Training Purposes Only

Question #2

Members of a gateway cluster communicate with each other over IPsec tunnels.

- True
- False



Knowledge Check

MOD 4- 36



true

Question #3

Each AP forms a single set of tunnels from itself to one member of the gateway cluster: the device designated gateway (DDG).

- True
- False



Knowledge Check

MOD 4- 37



false

Confidential - For Training Purposes Only

Question #4

If a cluster VLAN probe is successful, then the cluster is fully redundant, and only APs are fully synced.

- True
- False



Knowledge Check

MOD 4- 38



false

Confidential - For Training Purposes Only

Question #5

Regardless of how you deploy gateway clusters, an equal number of *buckets* are always assigned to cluster members, and therefore APs are equally distributed across the cluster.

- True
- False



Knowledge Check

MOD 4- 39



false

Confidential - For Training Purposes Only



Lab Activity

Let's review the lab activity that is associated with this module.

Confidential - For Training

Lab Tasks

Lab 04.01 - Deploy Tunneled WLAN

Task 1

- Review the wired network

Task 2

- Create PSK Tunneled WLAN with the GW cluster

Task 3

- Review the Configuration

Task 4

- Verify the Operation of the Tunneled WLAN

Task 5

- Configure GRE over IPsec

Lab 04.02



MOD 4- 41

The slide provides a brief review of tasks performed in this lab (4.01). Please see your Lab guide for details of this lab.

Lab Tasks

Lab 04.02 - Tunneled WLAN Cluster Operation

Task 1

- Review the Cluster status

Task 2

- Cluster Bucket Map

Task 3

- Load Distribution and Failover

**Module 5 -
Secure
Enterprise
WLAN**



MOD 4- 42

The slide provides a brief review of tasks performed in this lab (4.02). Please see your Lab guide for details of this lab.

Rev. 23.11




Secure Enterprise WLAN

Implementing Aruba Campus Access

Welcome back. This is module 5, which is about securing enterprise WLANs in a campus environment.

Confidential - For Training



Objectives

- Leverage RADIUS, 802.1X, and EAP knowledge to deploy more secure WLAN solutions
- Use your knowledge of role-based ArubaOS firewalls and role derivation to deploy more consistent, easier-to-manage security policies
- Enhance your overall security stance with solutions based on dynamic authorization and CoA capabilities

This module will really take your ability to secure the enterprise WLAN to a new level. After completing this module, you will be able to leverage RADIUS, 802.1X, and EAP knowledge to deploy more secure WLAN solutions in a campus environment. You will also be able to use your knowledge of role-based ArubaOS firewalls and role derivation to deploy more consistent, easier-to-manage security policies. And you will be able to enhance your overall security stance with solutions based on dynamic authorization and CoA capabilities.

The objectives of this module include the following:

Leverage RADIUS, 802.1X, and EAP knowledge to deploy more secure WLAN solutions.

Use your knowledge of role-based ArubaOS firewalls and role derivation to deploy more consistent, easier-to-manage security policies.




Enhance your overall security stance with solutions based on dynamic authorization and CoA capabilities.

Introduction to Authentication

Let's begin this module by providing an introduction to authentication for WLANs.

Confidential - For Training

WLAN Security Options

		Open	WPA-Personal	WPA-Enterprise
Authentication		None	PSK, passphrase	802.1X
	Optional MAC-Auth			
Encryption		Open: None	WPA: TKIP Deprecated	WPA: TKIP Deprecated
		OWE: AES	WPA2/WPA3: AES	WPA2/WPA3: AES



MOD 5- 4

Every WLAN operates in one of three broad categories, regardless of the additional authentication types you layer on it. 802.11 WLANs can be open, where users need not submit authentication credentials to connect. The client requests a connection, and the AP permits that connection—fairly simple. Traditionally, open networks offer no encryption. But today's Opportunistic Wireless Encryption (OWE) now provides Advanced Encryption Standard (AES) encryption for open networks, as you will soon learn.

Some WLANs are secured with Wi-Fi Protected Access (WPA), in one of two flavors.

WPA Personal requires users to enter a passphrase: a shared secret for all authorized users to connect. The passphrase is sometimes called a pre-shared key (PSK) and is used as the foundation to encrypt and authenticate traffic.

WPA Enterprise is recommended, due to its superior, 802.1X-based authentication support. With 802.1X, a client authenticates to a backend authentication server using individual credentials - a username/password and/or a digital certificate.

Both the Personal and Enterprise flavors support authentication and encryption. WPA comes in three versions, and the version determines the supported encryption algorithm. WPA originally used Temporal Key Integrity Protocol (TKIP), based on the RC4 encryption algorithm. The first implementation of WPA and TKIP are now deprecated.

WPA2 and WPA3 both support AES, but WPA3 can impose different requirements for the AES key length and mode of operation. Note that later versions can also have backward compatibility with earlier versions. For example, a WPA2 WLAN could permit clients to use TKIP or AES; however, for security, it should only allow AES.

The WPA versions have additional differences, as you soon learn. For now, just understand that the user experience for all three versions is similar. The user enters a passphrase to gain access to a WPA-Personal WLAN, but submits individual credentials to gain access to a WPA-Enterprise WLAN.

Finally, note that you can implement MAC-Auth together with any of these options. However, the focus for this module is the base security options.


Note: Wired Equivalent Privacy (WEP) was a pre-WPA security standard. Static WEP used a passphrase to authenticate users and derive a key used for encrypting data. Dynamic WEP was used 802.1X for authentication and a key established during the 802.1X exchange to encrypt data. However, both forms of WEP are proven very insecure and should never be used in today's networks.

Confidential - For Training Purposes Only

Authentication Overview


Supplicant

802.1X-capable
client software




Authenticator

Controls initial access



Authentication Server

Grants / denies access



802.1X/EAP

RADIUS



MOD 5- 5



Let's look at a typical port-based authentication scenario. You have endpoints that must connect to an access switch or AP to gain appropriate network access. For tight security, these endpoints often run a supplicant—software that enables use of 802.1X network authentication protocol in conjunction with an Extensible Authentication Protocol (EAP). In this context, the switch or AP is called the authenticator: the device that controls initial network access. In many solutions, the AP forwards these authentication messages on to its gateway (GW), which acts as the authenticator.

It is as if the authenticator says, “should I grant access to this device? I’ll ask the authentication server, which is responsible for making the decision to grant or deny access”. Therefore, the switch or AP/gateway harvests credentials carried in the 802.1X message, creates a RADIUS message, and sends it to the authentication server. RADIUS is both an authentication service running on some server, and it is a protocol used to communicate between authenticator and authentication server, used specifically for end user access control to network services and applications. To control administrative access to the network components themselves, like switches, routers, controllers, a TACACS+ service is often used. Many vendors, however, also use RADIUS to implement access control to the networking components themselves.

The RADIUS server might have a local data store of usernames, passwords, and group memberships, especially in a smaller deployment. In this case, the server checks this data store, validates credentials and sends “access granted” or “access denied” messages to the


authenticator, which in turn notifies the supplicant.

However, for larger, more scalable deployments, you typically have several authentication servers (for redundancy) all accessing a common set of backend servers (typically a directory service) that provide a common data store. This is often a Microsoft server running Active Directory (AD), for example. This allows clients to use single sign-on: the same credentials you use to access AD is also used to access the network itself, which is very convenient for the users. Some deployments might use a Lightweight Directory Access Protocol (LDAP) database, and you might also have a Certificate Authority (CA) service for even tighter endpoint and end user authentication.

Confidential - For Training Purposes Only

AAA Services

Authentication Who are you?	Validates credentials based on password/certificate
Authorization What can you do?	Controls authenticated user capabilities. Permits/Denies access
Accounting What did you do?	Auditing, reporting, tracking user activity



Common network access control (NAC) solutions are based around running AAA services on the authentication Server. So, what are the three A's of AAA services?


First is authentication: who are you? The RADIUS or TACACS+ protocol/service performs authentication based on your credentials. Perhaps your username is Maria, with a password of Secret123. If those are valid credentials, access is granted. The user might alternatively submit a digital certificate that includes a valid username and is signed by a trusted certificate authority (CA).

OK, now that you're on the network, what can you do? That's Authorization. User Maria is a member of the marketing group, so she can access all standard business applications, plus the marketing application and database. Other resources are off limits for Maria. You might define a policy to authorize Maria to fall into a certain VLAN, with certain quality of service (QoS) settings, which the switch port dynamically applies

We have controlled who can access the network, and what they can do. Now we need accounting: what did they do? So, we have auditing, reporting, and tracking of user activity: when and from where they logged in. When they logged out, and what resources they accessed during that time. One thing we're trying to eliminate here is "plausible deniability". If the log file indicates that you accessed an HR database at 2:00am, it is difficult to deny that fact.

Introduction to RADIUS

Remote Authentication Dial In User Service




MOD 5-7

Remote Authentication Dial-In User Service (RADIUS) is a key component of nearly all network security solutions. It is both a service, running on a RADIUS server, and a standardized protocol used to communicate with that server. This enables the vital security components listed here, as defined in RFCs 2865 and 2866.

RADIUS implements AAA services, support many authentication mechanisms and is extensible. Let's look at how RADIUS is involved in typical WLAN client authentication scenarios.

RADIUS Access-Request

Access-Request Client authentication/authorization for a connection attempt



MOD 5- 8

If the network has RADIUS-based authentication enabled and there is a client that attempts to connect that network, then the network device such as an AP (authenticator) will send a RADIUS Access-Request message to the RADIUS server. This message initiates a RADIUS-based authentication process. It contains number of attributes that provide additional information about the client and the network client is trying to access. Some attributes are generic such as NAS-Port-Type, which indicates if the connection is wired or wireless or the Calling-Station-ID, which contains the client's MAC address. There might be also Vendor Specific Attributes (VSAs) such as Aruba-Essid-Name, which is an Aruba attribute that describes the SSID name of the network or Aruba-location-Id, which contains the name of the AP.

RADIUS Access-Accept

Access-Accept

Connection attempt is authenticated and authorized
May provide instructions necessary to connect user

Client connected



Access-Accept

Aruba-User-Role=employee
Aruba-User-Vlan=11



MOD 5- 9

A RADIUS Access-Accept message is sent by the RADIUS server if the client is allowed to access the network. This means that the authentication process was successful and that the RADIUS policy permits network access. This message can optionally contain authorization attributes that provide restricts the supplicant's access, like applying an Aruba user role or a VLAN, for example, that should be applied to the supplicant's access.

RADIUS Access-Reject

Access-Reject Connection attempt is NOT authenticated and authorized

Client denied access



Access-Reject



- Authentication not completed
- Client rejects server certificate
- Client failed authentication
- Policy denies access




MOD 5- 10

An Access-Reject message is sent by RADIUS server in the following situations:

- The client did not complete the authentication process.
- The client rejected server certificate.
- The client failed authentication (for example, the wrong password was entered).
- The RADIUS network access policy denies network access.

After receiving this message, network device (authenticator) will not allow the client to access the network.

Client Connecting to Open/PSK SSID



The AOS 10 architecture utilizes the RADIUS protocol even if RADIUS-based authentication is not enabled in the WLAN configuration.

Take a look at the example in the figure illustrated in the slide. This WLAN SSID has Open-Mode enabled, with no encryption and authentication active. Immediately after client associates to the SSID Aruba AP will send RADIUS Access-Request message to the gateway. The gateway acts as a RADIUS server in this scenario. The Access-Request message contains the client (supplicant) MAC address as well as gateway network profile and other attributes that will be explained latter in this course. Since this network is Open with no authentication enabled, the gateway will simply reply with RADIUS Access-Accept message. This message contains attributes such as the Aruba-User-Role and Aruba-User-VLAN. Now the AP knows what role to assign to the client as well as what VLAN to use when sending client traffic through the GRE tunnel. Note that the gateway is responsible for client VLAN assignment on the AP for an open SSID.

Client Known to Gateway

```
(GW1)*# show user
```

This operation can take a while depending on number of users. Please be patient

Users

```
-----  
      IP           MAC           Name           Connected To       Forward mode  
-----  
10.251.1.50  1c:4d:70:07:fd:14  1c4d7007fd14  fc:7f:f1:ce:95:44  dtunnel
```

User Entries: 1/1

Curr/Cum Alloc:1/16 Free:1/15 Dyn:2 AllocErr:0 FreeErr:0



MOD 5- 12

Since the gateway receives the RADIUS request message sent from AP after the clients associates to the WLAN, this client is now known to the gateway. You can display information about all connected clients and known clients using the show user gateway command. In the output, dtunnel refers to decrypt tunnel: the GRE tunnel is not protected with IPsec.


802.1X Authentication



802.1X authentication is the most common method for restricting internal corporate access. This section introduces 802.1X.

Confidential - For Training

802.11 with 802.1X Connection Steps




Before you look at each EAP mode in more detail, it will be useful for you to review the steps in establishing an 802.11 WLAN connection. An AP periodically advertises beacons to identify the WLANs that it supports. The beacons include information about the network name, or SSID, and the security options. A client typically sends a probe request for the specific SSID to which it wants to connect to the BSSID MAC address of the associated radio. If the AP supports that SSID, with compatible options, it will send a probe response back to the client.

The client then sends an 802.11 authentication request. 802.11 authentication often operates in open mode, which means that the AP permits all clients to pass 802.11 authentication. The real authentication will occur later, as you will see.

Next, 802.11 association occurs. The client sends an association request, and the AP responds. If the WLAN is using Enterprise mode, 802.1X authentication follows the 802.11 association. The AP acts as the 802.1X authenticator. The 802.1X process can involve many packets, but in the end the client is hopefully authenticated.

The final step for all WLANs that support encryption is an exchange called the 4-way handshake. The client and AP use this handshake to derive the keys used to provide secure communications for the data; this handshake might also provide implicit authentication, as you will see. After the handshake, the 802.11 connection is up, and the client can send and receive data on it. Now the client can send DHCP request to obtain an IP address and then access network resources.

Client Connecting to 802.1X SSID



Once you enable RADIUS based authentication in the WLAN SSID profile and select tunnel mode for the forwarding operation, the gateway will act as a RADIUS proxy. The AP still acts as the authenticator and will generate the initial RADIUS Access-Request message. This message is then sent through the IPsec tunnel to the gateway, which is acting as RADIUS proxy and will forward this message to the RADIUS server (configured by the administrator in the WLAN profile). After receiving the Access-Accept from the RADIUS server the gateway, it will then forward the Access-Accept to the AP via the IPsec tunnel. This message contains Aruba-User-VLAN attribute, which will allow AP to forward the traffic through GRE tunnel with appropriate VLAN tag.

Key Management Service

Aruba roaming

Facilitates 802.11r and OKC clients

Neighbor AP list

Provided by AirMatch

Maximum 128 APs in the list (mobility domain)

Station synchronization

PMK or R1 key synchronized with neighbor APs

Machine auth state, VLAN, User Role shared to neighbor APs

PMK lifetime is hardcoded to 8 hours

Max 1,000 station records can be cached on each AP




MOD 5- 16

The key management service facilitates fast roaming for networks that use 802.11r and OKC. This service dynamically provides a list of neighbor APs (target APs) for each AP-client association; in other words, it dynamically manages the mobility domain. The neighbor AP list is provided by the AirMatch service via Aruba Central; a maximum of 128 APs are allowed in the list. The criteria for a neighboring AP is that the path loss should be less than 150dBm.

Another task of the key management service is the station synchronization. This is a process where the PMK or R1 key is shared with neighbor APs, this process also includes the Aruba user role, VLAN, and machine authentication state. The PMK lifetime is hardcoded to 8 hours and a maximum of 1,000 station records can be cached on each AP.

Key Management Service: Workflow



The following steps describe the key management workflow

1. A client associates and authenticates with the AP1.
2. AP1 cache client's information and sent it to the key management service in Aruba Central. The information includes the user's VLAN, the Aruba user role, and the R0 key in case 802.11r is used or PMK for OKC.
3. The key management service receives from AirMatch a list of all AP1's neighbors.
4. The key management service then generates R1 keys for AP1's neighbors.
5. Finally, the key management service shares R1 key or PMK along with VLAN and Aruba user role to target APs.

Configuring 802.1X SSID

The screenshot displays the configuration interface for an 802.1X SSID, specifically the Security tab. The interface features a navigation bar at the top with five tabs: 1 General, 2 VLANs, 3 Security (highlighted), 4 Access, and 5 Summary. Below the navigation bar, the Security Level is set to Enterprise, indicated by a slider and four buttons: Enterprise, Personal, Visitors, and Open. The Key Management is set to WPA3-Enterprise(CCM 128). The Primary Server is set to clearpass, and the Secondary Server is set to -- Select --. The interface also includes a search icon in the bottom left and a plus icon in the bottom right.

1 General 2 VLANs 3 Security 4 Access 5 Summary

Security Level: Enterprise Personal Visitors Open

Key Management: WPA3-Enterprise(CCM 128)

Primary Server: clearpass

Secondary Server: -- Select --

There are no differences on the bridged, tunneled and mixed SSID's security tabs. For 802.1X, select Enterprise using the security level slide and click Next. Keep in mind, when you select enterprise authentication, the primary authentication server configured here will be replicated at the gateway(s) within the cluster you selected from the VLANs tab.




Firewall Roles and Role Derivation

Aruba user roles and firewall policies were introduced in the Aruba Campus Access Fundamentals course and associated HPE Press book. This section expands upon that coverage.

Confidential For Training

Role-Based ArubaOS Firewall



Same network subnet
Same destination and application
Different user = different handling



MOD 5- 20

The role-based ArubaOS firewall can enforce different rules based on the user identity and assigned role, rather than simply relying on information found in packet headers. For example, the firewall can restrict contractors to sending web traffic to limited ranges of IP addresses while it allows managers' traffic to flow more or less restricted.

With less sophisticated firewalls, you would try to create such behavior by ensuring that contractors and managers receive IP addresses in different ranges. You would then create different access control rules based on the source IP addresses. However, such an approach makes planning user connections more complex: you need more VLANs/subnets and a way to get users in the right VLAN every time. With the ArubaOS firewall, you can place all users in the same subnet(s), but easily create different rules for the different groups of users.

The ArubaOS firewall determines the user role associated with a device as the device authenticates and accesses the network. It might assign this role entirely on its own, or it might have help from a AAA server like Aruba ClearPass using RADIUS.

By sending all internal user traffic through the role-based ArubaOS firewall, you can change from a model in which you protect a single perimeter to a model in which you place every user within their own unique and properly customized firewall.

What is a Role ?

Role: A group of settings that controls client traffic

Access control rules

Filter traffic based on L2-4 packet headers
Control dest. IP addresses and TCP/UDP apps accessed by clients

Application rules

Control client access to specific apps and web site categories
Enabled by DPI and WebCC

Bandwidth Controls

Rate-limit client traffic

Client portal redirection

For client connections enforced by captive portal authentication

Other settings

Role-based VLAN override, etc.


MOD 5- 21

A role is a group of policy settings that controls client access by bundling multiple settings that enable the ArubaOS firewall to control the client's traffic. Some of the most important settings for a role are access control rules, grouped together in firewall policies. These rules filter traffic based on layer 2-4 packet headers: to control the destination IP addresses and TCP/UDP applications accessed by clients assigned to the role.

However, the ArubaOS firewall has far more capabilities than plain access control rules. A role's policies can also include application rules, which let you control client access to specific applications—such as social media or gaming—or specific web site categories, based on URL or reputation. These policies are enabled by deep packet inspection (DPI) and web content classification (WebCC). The policies are the focus of this training. However, you should also be aware that the role can implement bandwidth controls to rate limit clients' traffic. And, if the client connection is enforcing captive portal authentication, the role can further specify how the client is directed to the captive portal.

Finally, roles define a variety of other settings. For example, a role-based VLAN overrides the VLAN assigned in another way such as through the WLAN settings. For wireless users it is best practice to assign the VLAN through the WLAN itself and leave the role-based VLAN non-specified.

Aruba Firewall Roles



MOD 5- 22

The Aruba Firewall is an identity based-firewall that assigns user-roles to clients pre- and post-authentication process. These roles contain parameters such as the VLAN ID, captive portal profile, session ACLs, and bandwidth contracts that dictate the access rights. The role basically defines the characteristics and restrictions of the entire network session for a user.

The session ACLs (also known as policies) are comprised of rules that contain matching criteria and actions. Traffic being processed by policies in a role is compared against the rules of the first policy top to bottom. If there is match against the conditions of any rule then the action is performed (either permit, deny, and any of the extended actions supported by the firewall). If there is no match, then the traffic is compared against the rules of the following policy, if there is match then the action is enforced, otherwise the firewall keeps comparing. If there is no match, then an implicit deny all at the bottom of the last policy will match traffic and drop it. This means that there is no implicit deny at the bottom of every policy but only at the bottom of the very last one.

You can imagine rule 1 of the Internet policy saying, “if this is Internet access, then permit it” and rule 2 saying “if this is corporate access, then deny it”. If traffic matches a rule, then the firewall takes an action or actions, and it does not examine any other policy rules. Possible rule actions include allow traffic, deny traffic, and apply QoS policy to the traffic.

Note: authenticated is an example of a default Aruba user role that permits all traffic. It has an allow_all policy associated. When using the WLAN wizard, uses this for successful

authentications unless you change it.

Confidential - For Training Purposes Only

Aruba Identity-Based Firewall


Peter
172.16.56.19
Role **Finance**



Robert
172.16.56.20
Role **Sales**



Susan
172.16.56.22
Role **Admin**



Aruba-User-Role

- **Finance:** policy1, policy2
- **Sales:** policy1, policy3
- **Admin:** policy2, policy4




MOD 5- 23



In the Aruba solution, user roles are assigned when the client associates to the network by using a role derivation process. This process can be as simple as assigning a default role to all clients connected to the same SSID, or as sophisticated as considering the user's attributes on the AAA server or back-end directory server, the client's connected AP, the time of the day, or other contextual data related to the connection. It is all up to you.

In the example shown in the slide, Peter, Robert, and Susan are connected the same AP, SSID, and subnet, yet each receives a different user role. In bridge mode, it is the AP that enforces these policies. In tunnel mode, it is the gateway the performs the enforcement.

Centralized and Consistent Security




MOD 5- 24

Here's another example of how user-roles allow you to centralize your security policies. Peter associates to the top-most AP. He is assigned the Finance role and his traffic bridges to subnet 10.2.30.0. As Peter moves from one AP to another, he may be mapped to a different subnet, but his assignment to the finance role is unaffected.


This can also apply to wired connections. Aruba switches use dynamic segmentation and create GRE tunnels to a dedicated gateway. The network treats the wired user like any other wireless client. This ensures that you have central control over a consistent security policy.

Firewall Roles in Tunneled Mode



When configuring a tunneled SSID profile using WLAN wizard in Aruba Central, role configuration is pushed to both the AP and the gateway. This means the client will have their role on both devices and will be a subject of two firewall policies, one on AP and one on the gateway. Either firewall can filter the user traffic. If packet is denied on the AP – it will be blocked. If the packet is denied on the gateway, it will be also blocked. For the traffic to pass it must be allowed on both the AP and the gateway

Role Derivation



User associates to an SSID

Initial role assigned

User derived rule (if present)

User authentication

Server derived role

Authentication default role

DD 5-26

The role derivation process assigns user roles to clients during certain connection phases, starting with association and then authentication. First the user associates to an SSID, where the user is automatically assigned the initial role configured in the AAA profile mapped to the SSID. The initial role is less useful for WLANs where 802.1X is used, but is vital for Guest SSIDs that use Captive portal authentication. The initial role defines what the user is allowed to do before actually authenticating.

Next, the process checks for a user-derived rule. User rules are used to override the initial role in pre-authentication state. Historically, if a device needed unique access requirements, you would put it in a different SSID. However, this can be problematic if you have lots of different types of devices with special initial access requirements, since it would require a lot of SSIDs to do this. However, with user rules, all these devices can be placed in the same SSID, but different initial access policies can be applied to them. Use this when you must base the role on user location, connected SSID, VLAN, MAC address, DHCP options, or encryption methods. A typical use case for this is a PSK protected network, where legacy IoT devices connect, such as bar code scanners or old IP phones. User rules are configured in advanced and attached to AAA profiles.

If the network uses no form of authentication, then the process stops here. Otherwise, an authentication process follows: either layer 2 MAC-auth or 802.1X authentication or layer 3 captive portal.

During authentication, a server-derived role can be applied, if configured on the AAA server. That is, the authentication server can return a standard RADIUS attribute that the AP/gateway matches against a rule contained in the server group, or you can use an Aruba Vendor Specific Attribute (VSA)—either `aruba-user-role` or `aruba-cppm-role`—that can automatically assign the role.

Lacking any server-derived roles, the network assigns the authentication method's default role that is defined for the WLAN.

Confidential - For Training Purposes Only

AAA Profiles: What Do They Define?

Authentication options

Authentication server group

Accounting server group


Gateway role derivation



MOD 5- 27

AAA profiles define authentication, authorization and accounting policy settings. These profiles allow you to setup authentication, authorization and accounting options, like the initial and default role assignments, for example. You can enable or disable the accounting server as well as select what RADIUS server should be used to authenticate clients. Optionally you can also specify the use of change of authorization (defined in RFC 3576), which we'll talk about later in the module.

Network Profile Communicated to AP & GW



When you create a WLAN using the Aruba Central wizard, the AAA network profile configuration is pushed to the both the AP and the gateway. The name of that profile corresponds with WLAN profile name. The AP receives AAA network profile name from Central and associates that with WLAN. The gateway receives the AAA network profile definition with all the settings. When the client connects to WLAN, the AAA network profile name is communicated to the gateway using the RADIUS VSA attribute Aruba-Network-Profile. Therefore, the gateway is aware what AAA network profile to use for authentication, authorization and accounting.

AAA Profile Configuration

The screenshot shows the AAA Profile Configuration interface. The top navigation bar includes 'Access Points' and 'Gateways' tabs. The 'Gateways' tab is selected, and the 'Security' sub-tab is active. The 'Role Assignment (AAA Profiles)' section is highlighted. The main area displays the configuration for the 'P28-Employee_#1668204011789_142#' profile. The configuration includes the following fields:

- Initial role: logon
- MAC Authentication Default Role: guest
- 802.1X Authentication Default Role: P28-Employee
- Set username from dhcp option 12: ☐
- L2 Authentication Fail Through: ☐
- Multiple Server Accounting: ☐
- User idle timeout:


To edit or review the AAA profile created by WLAN configuration wizard, navigate to your gateway group and click the Config icon. From there navigate to Security > Role Assignment (AAA Profiles).

Role Derivation: Server Derived Roles

Gateway network profile

AAA profile: **Corp_#16689_146#_**
 802.1X default role: **Corp**
 802.1X server group: **RADIUS_1**
 Default VLAN: **11**

SSID: Corp




In this first example shown in the slide, the client connects to the Corp WLAN. The AP generates a RADIUS Access-Request message and sends it to gateway. The gateway, acting as RADIUS proxy, forwards this request to the RADIUS server configured in AAA gateway network profile.

After successful authentication RADIUS server sends Access-Accept message with no role assignment attributes. The AAA gateway network profile has the Corp role configured as the 802.1X default role and it has Default VLAN set to 11. Therefore, the gateway adds the Aruba-User-Role=Corp and Aruba-User-Vlan=11 attributes to Access-Accept and sends it to AP. The AP now knows what role and VLAN to assign to the client. The client gets the Corp role assigned to it on both the AP and the gateway, all client traffic in GRE tunnel will be VLAN 11 tagged.

Role Derivation: Server Derived Roles

Gateway network profile


AAA profile: **Corp_#16689_146#_**
 802.1X default role: Corp
 802.1X server group: **RADIUS_1**
 Default VLAN: **11**



In the second example shown in this slide, the administrator has modified the role configuration when setting up WLAN network using WLAN wizard and has created a role called Employee. This role will be pushed to both AP and the gateway. Now, when the client connects to Corp WLAN, The AP generates the RADIUS Access-Request message and sends it to gateway. The gateway acting, as RADIUS proxy, forwards this request to the RADIUS server configured in AAA gateway network profile.

After successful authentication, the RADIUS server sends an Access-Accept message that contains the Aruba VSA attribute Aruba-User-Role=Employee. This attribute overrides the default role configured in AAA gateway network profile. The gateway retains the Aruba-User-Role=Employee and adds the Aruba-User-Vlan=11 attribute to the Access-Accept and sends it to the AP. The AP knows what role and VLAN to assign to the client and what VLAN to use when tunneling it via GRE. The client gets the Employee role assigned on both the AP and the gateway and all client traffic in the GRE tunnel will be VLAN 11 tagged.

Role Derivation: Server Derived Roles



In the last example shown in the above slide, the administrator has created role called Contractor that is pushed to gateway only. This is possible if you do not use WLAN wizard to configure roles and can present a possible problem, as you will soon see.

The client connects to the Corp WLAN. The AP generates a RADIUS Access-Request message and sends it to the gateway. The gateway, acting as RADIUS proxy, forwards this request to the RADIUS server configured in AAA gateway network profile. After successful authentication, the RADIUS server sends an Access-Accept message that contains the Aruba VSA attributes Aruba-User-Role=Contractor and Aruba-User-Vlan=12. These attributes override the default role and default VLAN configured in AAA gateway network profile. The gateway retains the Aruba-User-Role=Contractor and Aruba-User-Vlan=12 and sends an Access-Accept message to AP. The AP doesn't have role Contractor configured (since you didn't use the wizard and made the change for just the gateway), therefore, it will assign the WLAN default role Corp to the client. For GRE tunneling AP will use VLAN 12 sent by the gateway. At the same time, Gateway has role Contractor configured and so the client will have this role assigned on the Gateway.

Important: This example shows that you could have different roles assigned to client on the AP (default) and the gateway (Contractor). This can lead to confusion and troubleshooting issues since the AP and gateway are implementing two different roles. Therefore, best practices is to modify the user roles in the WLAN wizard, where you know that Central will push the configuration to both the APs and associated gateways.

Configuring Roles

The screenshot displays the 'Create a New Network' wizard with the 'Access' tab selected. The progress bar shows steps: 1 General, 2 VLANs, 3 Security, 4 Access (current), and 5 Summary. Under 'Access rules', there are three radio buttons: 'Role Based' (selected), 'Network Based', and 'Unrestricted'. Below this, there is a list of roles on the left and a table of access rules on the right. The role list contains 'P28-Employee' with a trash icon. The access rules table has one rule: 'Allow any to all destinations'. An 'Add Role' dialog box is open, showing a text input field with 'Employee' and 'OK'/'Cancel' buttons. The '+ Add Role' button at the bottom left of the role list is highlighted with an orange box.

Create a New Network

1 General 2 VLANs 3 Security 4 Access 5 Summary

Access rules

Role Based Network Based Unrestricted

Role

P28-Employee

Access Rules For Selected Roles

Allow any to all destinations

+ Add Role 1 Role(s) + Add Rule

Add Role

Role: Employee

Cancel OK



MOD 5- 33

It is recommended to configure roles using the WLAN wizard so that the role gets pushed to both the AP and the gateway. In Access section of the WLAN wizard, create new roles by clicking the + sign. Specify the role name and configure firewall rules that will be associated with that role.

Role Assignment Rules

Create a New Network

1 General 2 VLANs 3 Security 4 Access 5 Summary

Access rules

Role Based Network Based Unrestricted

Role

P28-Employee

Access Rules For Selected Roles

Allow any to all destinations

New Role Assignment Rule

Attribute: Operator: Integer: Role:

Acct-Authentic equals

Acct-Authentic

Acct-Delay-Time

Acct-Input-Gigawords

Acct-Input-Octets

Acct-Input-Packets

Acct-Interim-Interval

Acct-Link-Count

This field is required

Cancel Save

Aruba VSA attributes require no role assignment rules

MOD 5- 34

You can assign roles using custom attributes which you can select from the dropdown list in the Access section. Assigning roles using Aruba-User-Role attribute doesn't require any rules, since this will be process by default by Aruba's APs and gateways.

Server Rules

Access Points

Gateways

System

Interface

Security

Routing

High Availability

Config Audit

Roles

Policies

Aliases

Applications

Apply Policy

Auth Servers

Role Assignment (AAA Profiles)

SELECTED GROUP TYPE

Mobility Gateway

Summary

List

Config

Basic Mode

P28-Employee_#1668532776849_155#_auth_svg	1	--	--	2	
P28-Employee_#1668532776849_155#_cp_svg	1	--	--	0	

+

Server Group > P28-Employee_#1668532776849_155#_auth_svg

Drag rows to re-order

ATTRIBUTE	OPERATION	OPERAND	TYPE	ACTION	VALUE
Filter-Id	equals	emp	string	set role	Employee
Filter-Id	equals	corp	string	set role	Corp

MOD 5- 35

Any role assignment rules configured in WLAN wizard will be pushed to gateway as Server Rules. You can review or modify these rules if you navigate to the gateway group context > Config > Security > Auth Servers. Select the RADIUS server group and click the Server Rules tab.

Aruba Gateway Role-Based Firewall Policies

Role: contractor

Policy: **global-sacl**

global-sacl rules apply to every role in the system
Traditionally used to enforce organization wide policies

Policy: **apprf-contractor**

apprf-role name rules apply only to this one role
Used for application-based policy enforcement on this one role

Policy: **contractor**

role name rules apply only to this role
Used for policy rules-based enforcement on this one role

Policy: **Optional 1**

Optional policies may be shared between multiple user roles
Normally used for policy rules that apply to multiple roles

Policy: **Optional 2 or more**



MOD 5- 36

You are now ready to look at the role-based firewall policies in more detail. The ArubaOS firewall implements a different set of policies for each role. Each policy contains a series of rules or access control entries (ACEs), which filter traffic for users assigned to that role.

At the top of every role's list of policies is the global-sacl policy, in which you can place any rules that you want to apply to every user and device, no matter what the role (this would apply to all users, including guests, employees, contractors, printers, security cameras, and etcetera). By default, this policy is empty (no rules). In the slide, this policy has been modified and includes a rule for allowing DHCP traffic. This is just an example of the type of rule that you can put in this policy. If your corporate security policy prohibited the use for telnet by everyone, the easiest way to enforce that policy is to add a deny rule in the global-sacl. This policy is always processed first and you cannot change its order.

The next policy in every role is related to AppRF. It includes application rules that are specific to the user role. It is empty by default, and it is always the second policy that is applied. Its use is optional.

When you create a role in the Web UI, the role has a policy with the same name as the role automatically created for it. This is the default policy for access control rules that are specific to the role. It is empty by default, and it is always the third policy that is applied. Its use is optional.

In addition to adding rules to the default role-specific policy, you can also add other policies to the role. The AubaOS comes with multiple pre-configured policies, but you can also easily create your own policies.

Confidential - For Training Purposes Only

Components of a Session Rule

IP Version	Source	Destination	Service	Action	Extended Action
ipv4 ipv6	Any User Host Network Alias	Any User Host Network Alias	Protocol TCP port range UDP port range Service (pre-programmed well-known ports)	Permit Deny Reject (deny+ ICMP) NAT Redirect	Mirroring Quality of Service (QoS) More
Match				Action	Extended action

Examples:

ipv4	any	any	srv_dhcp	permit
ipv4	user	alias MyDNS	srv_dns	permit
ipv4	user	10.1.10.0/24	srv_https	permit

MOD 5- 37

Now let's examine the components of a rule in more detail. In the slide, you are looking at access control session rules. A session rule begins with the IP version, either IPv4 or IPv6.

For source and destination, you can specify a specific IP address (host) or a network. Because the rules are policy-based, the easiest way to create most rules is usually to specify user for the source. User applies to any IP address assigned to a client in the client table that is assigned to this role. In some cases, you will need to apply the rule to any IP address, rather than user. For example, a client needs to send DHCP traffic before receiving an IP address.

You can also specify an alias. An alias is a named object, which includes one or more host or network addresses. Aliases also make rules simpler to define and more flexible.

In addition to specifying the destination for an access control rule, you can also specify the service, which is defined by TCP or UDP destination port. You can define those manually or select a named service. You can also specify the service based on IP protocol.

Finally, you define the action. The basic actions are permit and deny. You can also specify reject, which denies the traffic and sends an Internet Control Message Protocol (ICMP) reject. More advanced actions include network address translation (NAT), in which the Mobility Controller forwards the traffic and translates its address, and redirection.

The gateway can also add extended actions on top of the main action. For example, you can specify mirroring for traffic that matches the rule—perhaps mirroring the traffic to an intrusion detection system (IDS) for further analysis. Refer to ArubaOS user guides or more advanced training for more information on other extended actions.

As you can see, while these rules might look like traditional ACL rules at first glance, they can do much more than filter packets.

The bottom of the figure shows a few simple examples of rules. The first permits any DHCP traffic. The second permits DNS traffic between clients assigned to the role and an alias named MyDNS, a DNS server. The third permits HTTPS traffic between clients assigned to the role and servers in 10.1.10.0/24.

Important: Unlike many products in the market placed, the policy is bi-directional. It affects traffic originated by the user as well as affecting traffic originated from somewhere else and sent to the user.

Confidential - For Training Purposes Only


Aliases

Network alias Networks or host IP addresses

- Pre-defined
- Admin-defined
- Examples: Domain_controller, Internal-network

Service alias Service, defined by protocol or TCP/UDP destination ports

- Pre-defined
- Admin-defined
- Examples: svc_icmp, svc_https, svc_ftp



The screenshot shows the ArubaOS Gateway configuration interface. The 'Security' tab is selected under the 'Gateways' section. The 'Aliases' sub-tab is active, displaying a table of network aliases. The table has three columns: NAME, ITEMS, and DESCRIPTION. Pre-defined aliases include 'any', 'auth-facebook', 'auth-google', 'controller', 'localip', and 'mswitch'. A plus sign at the bottom indicates the option to add new aliases.

NAME	ITEMS	DESCRIPTION
any	1	--
auth-facebook	3	--
auth-google	2	--
controller	1	--
localip	1	--
mswitch	1	--



MOD 5- 38

Aliases are very useful, so let's take a moment to examine them in more detail. ArubaOS supports two types of aliases:

- Network
- Service

Network aliases can be used for the source or destination address(es) of an access control rule; they represent one or more networks or host addresses. The ArubaOS comes with several pre-defined aliases. For example, controller is an alias that includes any IP address of a gateway controlled by this solution. User and any are very useful aliases, which you will examine in more detail in a moment. Admins can also define their own aliases such as Domain_controller or Internal-network. Here is more complete list of pre-defined destination aliases is:

Mswitch—the primary IP address of the Mobility Conductor

Controller—The IP address of the Mobility Controller

Vrrp_Ip—The VRRP address used by the Mobility Controller

Any—Any IP address

User—IP address of a user assigned to this role

localIP—the local IP address of a Remote AP

Auth-google--An IP address used with Google authentication

Auth-facebook—An IP address used with Facebook authentication

Wificalling-block—A component of the UCC service for voice. This alias enables you to block voice calls.

Service aliases represent a service or application based on IP protocol or TCP or UDP port list. The ArubaOS comes with many pre-defined service aliases for well-known applications. For example, `svc_icmp` corresponds to the ICMP protocol, and `svc_https` corresponds to TCP port 443. If the company has custom applications that use defined ports, you can create your own aliases that similarly select traffic by IP protocol or TCP or UDP destination port or port range.

Confidential - For Training Purposes Only

Benefits of Aliases

- Self-documenting rules
- Better manageability and scalability with dynamic updates

Role Finance

Policy internal

Rule user 172.16.0.0/16 any permit

Static

Permits access to 172.16.0.0/16

Role Admin

Policy building

Rule user **Internal-network** any permit

Dynamic

Permits access to 172.16.0.0/16 and **10.107.0.0/16**

Network Alias

netdestination **Internal-network**

network 172.16.0.0 255.255.0.0

network 10.107.0.0 255.255.0.0



MOD 5- 39



Aliases have several benefits. First, aliases are self documenting. For example, when you specify Internal-network rather than 172.16.0.0/16 in a rule, other admins can instantly understand the intent of the rule. In addition, aliases make the solution more manageable and scalable. You can specify multiple subnets and ranges within an alias, so that you can create one rule rather than multiple rules.

Aliases also update rules automatically. In this example, the Internal-network alias originally included 172.16.0.0 255.255.0.0. At that point, the policies in the Finance, Sales, and Admin roles all permitted users to access 172.16.0.0/16. However, the rule in the Finance role is static. It never updates. This makes the policy harder to manage when changes occur. For example, the network might expand into the 172.17.0.0/16 space. With aliases, you can easily add that network to the Internal-network alias. The rules that use this alias now automatically update, permitting sales and admin users to access 172.16.0.0/16 and 172.17.0.0/16.

Trusted vs Untrusted Traffic

Untrusted Traffic

- AAA process applied
- Receives a firewall role
- All wireless traffic always untrusted

Any alias always applies



User alias applies with IP Assignment

Trusted Traffic

- No AAA process
- No firewall role

Any alias always applies



MOD 5- 40

It is important to understand the difference between trusted and untrusted Traffic and how it is referenced in firewall aliases. All wireless traffic is considered untrusted traffic, by default. All untrusted traffic will be processed by the AAA process and assigned a firewall role through one of the role derivation methods discussed earlier.

Trusted traffic, on the other hand, is not authenticated and is not assigned a firewall role.

The any firewall alias will apply to both trusted and untrusted traffic. The user alias, however, only applies to untrusted traffic after the client has received an IP address assignment.

Alias User

Robert - 10.2.55.79
Role IT



Peter - 10.2.55.78
Role Finance



Julie - 172.16.22.22
Role Finance



Frank - 192.168.2.34
Role Finance



Susan - 172.16.56.22
Role IT



Role **IT**
Policy 123

Rule **user** any any permit

Users of role IT
10.2.55.79
172.16.56.22

Role **Finance**
Policy xyz

Rule **user** Internet any permit

Rule **user** server any permit

Users of role Finance
10.2.55.78
172.16.22.22
192.168.2.34

Helps prevent IP spoofing

MOD 5- 41



This example features two roles, each which have rules with the user alias. In this example, Robert and Susan are at different sites. They have been assigned IP addresses in different subnets by DHCP. But both users are assigned to the IT role, and so their IP addresses are both automatically included in the user alias for that role.

Peter is in the same subnet as Robert, but he's assigned to the Finance role, so his client IP address is included in the user alias for Finance rules. Similarly, Julie's and Frank's IP addresses are included in the Finance rules' user alias, even though Julie and Frank are at different sites and in different subnets.


The ArubaOS firewall always applies a role's rules only to traffic from or to clients assigned to that role. Therefore, the user alias can seem like the any alias in some ways. However, using the user alias in rules can help to prevent IP spoofing. For example, if Robert tries to spoof IP address 10.15.10.10, his traffic will not match the "user any any permit" rule, as 10.15.10.10 is not included with the IT users.

As you learn more about the bi-directional nature of the firewall, you will come to understand more differences between the user and any aliases.

Session Rule Handling by Stateful ArubaOS Firewall

Contractor role

Contractor policy
ipv4 user 10.1.10.0/24 svc_https permit




MOD 5- 42

You should also understand that the session access control rules are stateful and bi-directional.

A stateful firewall, like the ArubaOS firewall, can track which packets belong to the same session and handle the traffic as a session, rather than only as individual packets. When a packet that is not part of an existing session arrives, the firewall checks that packet against its rules and permits or denies it. For a permitted packet, the firewall creates a session associated with the traffic. If more traffic in the same session arrives, including return traffic, the firewall permits that traffic. If you have set up a rule to permit clients to access a server on TCP port 443, as in this example, you do not have to create another rule to permit the server to access clients.

The same holds true for a rule that permits UDP DNS traffic from the client to the DNS server. When the client sends a DNS request to the server defined in the MyDNS alias, the firewall permits the DNS response from that server without needing another rule for the return traffic.

Bi-directional Rules: "user" vs "any"




**User applies with
IP Assignment**



10.2.55.79

```
ip access-list session pingTest2
user any svc-icmp permit
```



Any applies



10.254.1.100



MOD 5- 43




Now that you understand how the stateful firewall works, you can better understand the difference between the any and user aliases.

When you create a permit rule with the user alias as a source address, you are not just preventing spoofing. You are also indicating that the controlled client must initiate the traffic flow associated with that rule. When a permit rule includes the any alias as a source address, on the other hand, a device on the outside of the firewall can also initiate the session.

Here you can see the differences between the any and user aliases. The top rule has any for source the and permits ICMP (pings). A client assigned to the role can ping the server. The reply is permitted by the bi-directional stateful firewall, but would be permitted in any case, because pings are permitted from any IP address. The server can also initiate a ping, and the rule permits this traffic because it permits ICMP from any IP address, within or outside of the firewall.

The bottom rule, on the other hand, uses user for the source of the rule that permits ICMP. The client assigned to the role can ping the server and get a reply, because the bi-directional, stateful firewall permits return traffic. However, the server is not a user and therefore it doesn't match this rule: it is not allowed to initiate a ping.

AppRF 2.0 Enforcement



To enforce application rules, the firewall goes through two phases. Let's examine a situation in which a packet arrives for which no session yet exists. During phase 1, the application detection phase begins. The firewall checks the packet's Layer 3 and 4 headers against all rules in the user's role, beginning at the top. Application rules should be at the top, so, if the packet uses the implicit ports associated with an application rule, it hits the application rule first. The firewall flags the packet as subject to that rule. (If the packet does not match an application rule, the firewall applies the first matching access control rule as normal.)


The firewall lets traffic flagged by an application rule flow temporarily, but it uses DPI to detect the application within the traffic. Phase 2 then begins; the firewall applies the flagged application rule or rules and permits or denies the traffic based on application category or Web category and reputation. Typically, the AppRF firewall will have to see a few packets before it can assign it to an application and then enforce the application policies.

CoA Support

Let's learn about supporting change of authorization (CoA) functions.

Confidential - For Training

Dynamic Authorization



MOD 5- 46


A RADIUS server like ClearPass can make decisions about the correct authorization level for a user or device beyond the initial authentication. For example, ClearPass OnGuard can detect a client's security status and tell ClearPass to disconnect a client that is infected with malware. Or ClearPass might find out after the initial authentication that the client is an IoT device, which has particular security policies associated with it, and thus its user role needs to be changed. Many such use case examples exist.

For ClearPass to enforce the new policies, it must force the client to disconnect and re-authenticate in many cases. It sends a RADIUS change of authorization (CoA) message to the RADIUS client (gateway). By default, it sends the message on standard port UDP 3799. If it is configured to use a different port, you must make sure to match that configuration on the ArubaOS device acting as RADIUS client.

When the ArubaOS device receives a CoA from an authorized dynamic authorization server, it enforces the message and disconnects the client (assuming the latter is defined). The client will typically try to re-authenticate, and ClearPass can then apply its policies to make new decisions about the client's access.

Note: Originally, the RFC for this process referred to it as “change of authorization” and many people still use that term. However, the term dynamic authorization is used today to describe this process and is seen in the current RFC, RFC 3576.

Dynamic Authorization Flow



MOD 5- 47


If CoA action is triggered, then the RADIUS server sends a CoA message to gateway. This message contains the client's MAC address. The gateway then forwards this message to the AP. The AP de-associates the client and sends a CoA acknowledgement message to the gateway. The gateway proxies the acknowledgement message to RADIUS.

AAA User Delete

Client
D0:37:45:72:60:C3




Client
disconnected



MOD 5- 48

If you access the gateway CLI and use the `aaa user delete` command, this will also result in CoA Disconnect message sent from gateway to the AP. The client gets disconnected and the AP sends an acknowledgement to gateway.

Configuring Dynamic Authorization



Access Points Gateways SELECTED GROUP TYPE Mobility Gateway

System Interface **Security** Routing High Availability Config Audit

Roles Policies Aliases Applications Apply Policy **Auth Servers** Role Assignment (AAA Profiles) ✓

All servers

NAME	TYPE
AS1_#CLOUD_AUTH#_	Radius
AS2_#CLOUD_AUTH#_	Radius
c ppm	Radius
--	RFC 3576

New server

IP address: 10.254.1.23

Key: *****

Retype key: *****

Type: RFC 3576 ✓

Cancel Save

MOD 5- 49

In order to configure dynamic authorization (CoA), navigate to the gateway group > Security > Auth Servers and add the RFC 3576 server. Specify server IP address and shared key. This is typically the same server you're using for the other AAA functions. Preferably, though, you would make this change to the AAA profile from the WLAN wizard so that both the AP and gateway receive the same configuration!

Adding RFC 3576 to AAA Profile

The screenshot displays the Cisco configuration interface for a Mobility Gateway. The top navigation bar includes 'Access Points' and 'Gateways' (selected). The 'Gateways' section shows 'SELECTED GROUP TYPE: Mobility Gateway' with 'Summary' and 'List' options. Below this, a tabbed interface includes 'System', 'Interface', 'Security' (selected), 'Routing', 'High Availability', and 'Config Audit'. The 'Security' tab has sub-tabs: 'Roles', 'Policies', 'Aliases', 'Applications', 'Apply Policy', 'Auth Servers', 'Role Assignment (AAA Profiles)' (selected), 'L2 Authentication', and 'L3 Authentication'. The main content area is divided into two panels. The left panel, titled 'AAA Profile', lists several profiles: 'AAA Profile', 'AOS10-Corp_#16683474...', '802.1X Authentication Profile', 'MAC Authentication Profile', '802.1X Authentication Server Group', 'MAC Authentication Server Group', 'RADIUS Accounting Server Group', 'RFC 3576 server' (highlighted with an orange box), and 'XML API server'. The right panel, titled 'RFC 3576 Server', shows a table with the header 'RFC 3576 SERVER' and the IP address '10.254.1.23'. An orange box with a plus sign is located at the top right of this panel, indicating the addition of the server to the profile.




MOD 5- 50

Here is an example of assigning previously created RFC 3576 server to a AAA profile.

Dynamic Authorization Support

- ✓ Available in manual cluster config
- ✓ VRRP instances per cluster member
- ✓ VIPs used as NAS IP
- ✓ CoA/DM messages sent to VIP
- ✓ VRRP Instances auto-increment



MOD 5- 51

Dynamic authorization supports CoA requests in a gateway cluster. RADIUS servers send CoA or disconnect message requests to gateways to dynamically modify the authorization attributes for a connected client session. Administrators can enable the dynamic authorization feature to ensure that the CoA requests are not dropped when gateway nodes change due to load balancing or in the event of a failover.

This feature requires to run a VRRP instance per cluster member. In the figure in the slide, you can see how instance 1 generates a VIP1 and GW-1 is the VRRP master. A second instance must be run for the second cluster member, on instance 2 GW-2 is now the VRRP master.

Dynamic Authorization Configuration

The screenshot shows a configuration page for dynamic authorization. It includes several toggle switches and input fields. Five orange boxes with dashed lines point to specific elements, each with a numbered instruction:

- 1. Enable manual configuration (points to 'Manual cluster configuration' toggle)
- 2. Enable dynamic authorization (points to 'Dynamic authorization (CoA)' toggle)
- 3. Select the VIP for each GW (points to the VRRP IP column in the gateway table)
- 4. Set the VLAN (points to 'VRRP VLAN' input field)
- 5. VRRP ID – leave the default (points to 'VRRP ID' input field)

The 'Cluster name' is set to 'GW-Cluster'. The 'VPN termination' toggle is disabled. Below the toggles is a table titled 'Gateways in P9T11-GW-Group Cluster'.

GATEWAY	VRRP IP	
GW-1	10.1.110.200	
GW-2	10.1.110.201	

Below the table, 'Multicast VLAN' is empty. 'Heartbeat threshold' has 'Default' selected. 'VRRP VLAN' is set to 110. 'VRRP ID' is set to 110.

MOD 5- 52

To support dynamic authorization configuration for a gateway cluster, you must configure a manual cluster, where you define a VRRP configuration with a VIP address for each gateway in the cluster. Dynamic authorization uses a VIP address as the NAS IP on RADIUS requests. The RADIUS responses and CoA packets are sent to the VIP address. If the original gateway has failed, the backup will assume the VIP address and process RADIUS responses.



Knowledge Check

Self-check on key learning points



It's time for a knowledge check to review what you have learned.

Confidential - For Training Purposes Only

Question #1

For WLAN security options, what is a primary difference between WPA-Personal and WPA Enterprise?

- a. WPA Enterprise uses more powerful encryption algorithms
- b. WPA Enterprise uses more powerful hashing algorithms
- c. For WPA Enterprise, TKIP has been deprecated
- d. WPA Enterprise uses a more secure authentication mechanism

a

Knowledge Check

MOD 5- 54



d

Confidential - For Training Purposes Only

Question #2

You are diagnosing a WPA Enterprise authentication issue for a client. You are not seeing any RADIUS packets between the endpoint and the AP, which is serving as the Authenticator. What is the most likely reason for this?

- a. The authentication session is failing before the RADIUS server is involved
- b. RADIUS packets are being blocked by an intervening device
- c. The AP has not been properly configured as an authenticator
- d. This is normal. RADIUS communications are primarily between the AP and the authentication server

a

Knowledge Check

MOD 5- 55



d

Confidential - For Training Purposes Only

Question #3

You are diagnosing an 802.1X-based WLAN authentication issue. All other endpoints are working fine but this one. You see that the endpoint has successfully passed 802.11 authentication, yet the client cannot access the network. What is the most likely cause of this issue?

- a. The client supplicant does not support 802.1X authentication
- b. 802.11 authentication is merely a handshake and is not related to 802.1X authentication
- c. The client is configured to use the wrong hashing algorithm
- d. There is a password mismatch between the authenticator and the authentication server

a

Knowledge Check

MOD 5- 56



b

Confidential - For Training Purposes Only

Question #4

Which of the following are valid components of an Aruba firewall role?

- a. Access control rules that filter traffic based on L2-4 packet headers
- b. Application rules that control client access to specific applications
- c. Bandwidth controls that rate-limit traffic
- d. Client portal redirection
- e. Role-based VLAN override

a

Knowledge Check

MOD 5- 57

+

a,b,c,d,e

Question #5

One advantage of tunneled mode for Aruba WLAN deployments is that end users are only subject to a single firewall policy, applied on the gateway.

- True
- False

Q

Knowledge Check

MOD 5- 58

+

false

Confidential - For Training Purposes Only

Question #6

One advantage of using dynamic authorization with change of authorization (CoA) messages is that clients are not required to disconnect to enforce new policies.

- True
- False

Q

Knowledge Check

MOD 5- 59

+

false

Confidential - For Training Purposes Only



Lab Activity

Let's review the lab activity that is associated with this module.

Confidential - For Training

Lab Tasks

Lab 05.01- Deploy Tunneled Corporate WLAN

Task 1

- Understanding the AAA Profile on PSK WLAN

Task 2

- Configure Corporate 802.1X Tunneled WLAN

Task 3

- Connect with a WLAN client

Task 4

- Monitoring and Roaming Key Distribution

Lab 05.02



MOD 5- 61

The slide provides a brief review of tasks performed in this lab (5.01). Please see your Lab guide for details of this lab.

Lab Tasks

Lab 05.02- Roles and Access Control

Task 1

- User Role Derivation

Task 2

- Use the WLAN Workflow to Apply Access Control

Task 3

- Gateway Controlled Access Control

Task 4

- Gateway Controlled Access Control using the User alias

Task 5

- Configure Dynamic Authorization with the Gateway Cluster

Task 6

- Server Rule based Role Derivation

**Module 6 -
Guest and
Captive Portal**



MOD 5- 62

The slide provides a brief review of tasks performed in this lab (5.02). Please see your Lab guide for details of this lab.

Rev. 23.11




Guest and Captive Portal

Implementing Aruba Campus Access

Confidential For Training

This is module 6—Guest and Captive Portal. Guest and captive portal access was introduced in the Aruba Campus Access Fundamentals course and corresponding HPE Press book. This module will expand upon that information.



Objectives

- Describe guest access and captive portal options
- Explain the captive portal authentication process
- Configure guest WLANs
- Leverage MAC caching to improve the user experience for certain devices

MOD 6- 2



After completing this module, you will be able to:

- Describe guest access and captive portal options.
- Explain the captive portal authentication process.
- Configure guest WLANs.
- Leverage MAC caching to improve the user experience for certain devices.

Introduction to Guest Access

Let's begin with an introduction to guest access.

Confidential - For Training

Aruba Guest Access Solution

Single infrastructure

For internal & guest

Firewall rules for


Access control (User Role)

User tunneling

When needed

Captive Portal

Authentication



MOD 6- 4

Aruba allows for a simple collapsed architecture for your company's WLANs and Ethernet connectivity. You get differentiated access based on user and device characteristics. This is the basis for several guest access features. Guest access is often configured as a software option. No new hardware is required for basic guest access beyond the Aruba gateway and/or APs used for the internal WLAN.

In contrast to some vendors, where you must configure the LAN to add a VLAN for guest access at every LAN switch where an AP connects, you can add Aruba's user-centric networks as an overlay on the existing wired LAN. AP traffic is directed via secure tunnels straight to the gateway, where an integral stateful firewall maintains strict segregation between different traffic classes.

Internal traffic gets access to core LAN and corporate resources, while guest traffic travels through a secure tunnel to a gateway in the DMZ, typically, and then on to the Internet. Or you can let the AP route traffic to the gateway or directly to the Internet. You configure this "split tunnel mode" for an AP group.

Captive portal login screens and web forms for administration are served directly from the mobility gateway or you can use the Aruba ClearPass server for more sophisticated guest access solutions, like credit card processing, access code authorizations, property management systems, advanced forms and web page development, and more.

Captive Portal Options

Cloud Guest

- Splash page hosted in Aruba Central
- Managed by Aruba Central
- Can be customized

External Captive Portal

- Splash page hosted on external server
- Customization depends on server
- Example : ClearPass guest




MOD 6- 5

There are two options to deploy a guest network with captive portal authentication:

- Cloud Guest: with the cloud guest service from Aruba Central, a network admin can easily implement captive portal authentication in a multisite environment with consistency and central management. The splash page is customizable and includes great features like guest self-registration and social login.
- External captive portal: with the external captive portal option, you may use an authentication solution that's not part of Aruba Central or that does not run on the gateways and APs. An example is the Aruba ClearPass Guest and Policy Manager, which can provide guest captive portal authentication with many customization options and scalability.

Single Gateway Cluster Architecture


- ✓ Single gateway cluster for employees and guests
- ✓ The same GRE tunnel for both types of clients
- ✓ Client separation provided by VLANs and firewall



MOD 6- 6

You can terminate guest clients on the same gateway cluster as employees. This is referred to as a single gateway cluster architecture. APs send guest user traffic and employee user traffic to the gateway cluster using the same GRE tunnel, typically using different VLANs, though. In this solution guest clients are kept separate from employees by assigning them a dedicated guest user role which is not allowing access to the corporate network. In addition to the role-based firewall feature, you can keep guest users isolated by assigning them a dedicated VLAN (for example, VLAN 11 for employees and VLAN 13 for guests).


Multi-Cluster Architecture



MOD 6- 7

A multi-cluster solution involves dedicated a gateway cluster for employees and guests. In this architecture AP forwards employee user traffic to corporate gateway cluster and guest user traffic to guest gateway cluster, likely isolated in a DMZ. This solution represents enhanced security as guest traffic is isolated from the corporate network by using a dedicated gateway cluster for guests.


External Captive Portal



Supported guest solution in combination with tunneled mode WLAN MOD 6- 8

The Aruba tunnel WLAN mode supports external captive portal authentication. In this scenario guest user will be redirected to a splash page hosted on an external server such as ClearPass Guest. HTTP(s) is the required communication protocol between gateway cluster and the ClearPass server. In addition to this, the gateways must be able to send RADIUS Access-Request messages to ClearPass, which could be using plain RADIUS or RADSEC encrypted communication.

Cloud Guest




Cloud Guest supported in bridge mode

MOD 6- 9

Aruba Central provides Cloud Guest functionality, which allows splash page and guest accounts to be hosted in the Aruba Central cloud. This scenario requires APs to be able to access Aruba Central using HTTPS and RADSEC. At this time Cloud Guest is supported in bridge mode only.

Guest Client Experience: High-Level Overview



MOD 6- 10

From a guest user point of view, getting network access is a simple process. First, the client connects to the guest SSID. When the user opens up a browser and tries to navigate to a website, the browser connection will be redirected to the captive portal page. Depending on the captive portal configuration, client is now required to enter guest user credentials, register an account, or simply accept the terms and conditions for network access. After completing these steps, Internet access should be possible. Let's take a look at what happens behind the scenes during that process.

Note: Even though captive portal access is most typically used for guest access, it can also be used for internal access. Many companies use a captive portal to onboard new devices and employees.

Guest Authentication Process



A good understanding of the guest auth process can improve your deployment and diagnostics skills.

Confidential - For Training

Captive Portal Auth Requirements

Security mode for guest WLAN

- Typically: Open
- Optionally: Enhanced open
- Captive portal authentication

Initial guest access is limited

- Must allow: DHCP, DNS, HTTP(S) to CP server
- Everything else denied
- Pre-authentication role

After authentication access change

- More complete access
- Post-authentication role




MOD 6- 12

In order to support guest users, you need to create a guest WLAN profile. Aruba recommends always using the WLAN configuration wizard to do this. The guest SSID usually operates in open mode and, as an option, you could enable enhanced open (OWE).

For the authentication mode you should select Visitor in the wizard, which applies to utilizing captive portal authentication. In addition to that you, must make sure that the guest clients are redirected to the splash page after they connect to guest SSID. To make this possible, clients initially must receive limited network access that allows DHCP, DNS and enforces web redirection. This is provided using a guest pre-authentication role.


After successful web authentication guest user role must change to post-authentication role which is less restrictive and allows Internet access. If you use WLAN configuration wizard these roles will be automatically configured and pushed to both the AP and the gateway.

Captive Portal Auth Process: Phase 1



During phase 1 guest client connects to guest SSID, which triggers AP to send RADIUS Access-Request message to the Gateway. This RADIUS message contains Aruba-Network-Profile. During phase 1, the guest client connects to the guest SSID. This triggers the AP to send a RADIUS Access-Request message to the gateway. This RADIUS message contains the Aruba-Network-Profile attribute that matches the AAA profile name on the gateway. In that AAA profile, the initial role is set to the pre-authentication role, which will grant limited network access and assign the guest VLAN to client. If you create the WLAN profile using wizard, then all the necessary roles and profiles are configured automatically on both the APs and gateways.

Captive Portal Auth Process: Phase 2




AOS 10: AP does web redirection

MOD 6-14

Next step is that the client opens up a browser and tries to connect to a website. The pre-authentication (initial) role contains a policy that enforces web redirection to the splash page hosted on the external server, like ClearPass. Please notice that in AOS 10, the AP is responsible for that web redirection. The client gets splash page displayed.

Note: In AOS 8 and earlier, it was the controller (gateway) that performed the redirection.

Captive Portal Auth Process: Phase 3



After the client enters their credentials, the username and password are sent by the client browser in an HTTPS POST message to a special host which is located, by default, at `securelogin.hpe.com`. That message is captured by the AP which can extract username and password entered by guest user. This information is now used by the AP to issue RADIUS Access-Request message that gets sent to the gateway.

This message contains the layer 3 (L3) authentication profile name that exists on the gateway and was created by the WLAN configuration wizard. This L3 authentication profile contains instructions on how to perform web authentication and what external RADIUS server(s) to use. The gateway forwards the RADIUS Access-Request message to an external server and the authentication process begins. If authentication is successful, the RADIUS server sends an Access-Accept that may contain additional attributes. After receiving the Access-Accept from the RADIUS server, the gateway forwards this message to the AP and adds the post-authentication role name as the `Aruba-User-Role` attribute. The AP updates the connected client role to typically a less restrictive one via the received attribute from the gateway. At this point guest user can access the Internet.


Captive Portal Auth Process: Phase 3

L3 authentication profile

L3 auth. profile: **Guest_#1656553...**


Default role: **Guest**

Server group: **Radius_1**



In the slide you can see the details of the RADIUS messages exchanged between the AP, gateway, and the RADIUS server. Take a close look at the Aruba-Network-Profile name communicated by the AP to the gateway in the Access-Request message. Notice that after successful authentication, the gateway will send the post-authentication role to the AP, changing the limited pre-authentication role to a more complete and less restrictive post-authentication role.

Using "securelogin.hpe.com"



MOD 6-17

How do APs know what HTTPS POST message to capture when the client enters connects to the web?

The Aruba AP monitors DNS requests generated by the clients. When a factory default AP sees a request for securelogin.hpe.com, it will intercept the request and return the AP's own IP address. The behavior and identity for securelogin.hpe.com are dependent on the common name in the HTTPS certificate that is applied on the Aruba AP. securelogin.hpe.com is the common name identity contained in the HTTPS certificate on the AP, by default. For example, when the AP sees the request for securelogin.hpe.com, it replies with its own default IP address. This certificate is signed by public Certificate Authority and is pushed to APs from Aruba Central by default.



Guest WLAN Configuration

Let's now cover how to configure those guest WLANs!

Confidential - For Training

Guest WLAN Setup

The screenshot shows the 'Guest WLAN Setup' wizard in the network management interface. The interface has a top navigation bar with 'Access Points' and 'Gateways' tabs. Below this is a sub-navigation bar with tabs for 'WLANs', 'Access Points', 'Radios', 'Interfaces', 'Security', 'Services', 'System', 'IoT', and 'Configuration Audit'. The 'WLANs' tab is selected. On the right, there are links for 'Summary', 'List', and 'Config'. A 'Hide Advanced' link is also present.

The main content area is titled 'Create a New Network' and features a progress bar with five steps: 1 General, 2 VLANs, 3 Security, 4 Access, and 5 Summary. The 'Security' step is currently active and highlighted with a green arrow.

Under the 'Security' step, there is a 'Security Level' section with a slider. The slider has four positions: Enterprise, Personal, Visitors, and Open. The 'Visitors' position is selected and highlighted with an orange box.

Below the 'Security Level' section is the 'Access Network' section. It contains three fields: 'Type', 'Captive Portal Profile', and 'Primary Server'. The 'Type' field is set to 'External Captive Portal' and is highlighted with an orange box. The 'Captive Portal Profile' field is set to 'Captive-Portal' and is highlighted with an orange box. The 'Primary Server' field is set to 'clearpass' and is highlighted with an orange box.

A large diagonal watermark reading 'Confidential - For Training Purposes Only' is visible across the center of the interface.

To configure a guest WLAN, start with WLAN configuration wizard. In the Security section, select Security Level: Visitors. Set the Type to External Captive Portal – this requires a captive portal profile to be created. Select from the list or add RADIUS server as Primary Server. Optionally define a secondary server.

External Captive Portal Setup

External Captive Portal-New

Name:	Captive-Portal
IP or Hostname:	guest.company.com
URL:	/guest/AOS10-guest.php
Port:	443
Use HTTPS:	☒
Captive Portal Failure:	Deny Internet ▼
Server offload:	☐

MOD 6-20

Next, create a captive portal profile – enter a name for the profile the IP address or hostname of your external captive portal server. Specify redirection URL (captive portal splash page location on the external host) and the Port. For more secure authentication, select Use HTTPS. The name of the profile created at this stage is significant because the wizard will create a pre-authentication role with the same name. This role gets pushed to both the AP and gateway.

Post and Pre-Authentication Role

Networks > Configuration - AOS10-Guest

General VLANs Security Access Summary

Access rules

Post-authentication role

Role Based Network Based Unrestricted

Role

AOS10-Guest

Captive-Portal

cppm

Access Rules For Selected Roles

Deny any to network 10.1.0.0/255.255.0.0

Allow any to all destinations

Pre-authentication role

+ ADD ROLE ASSIGNMENT

Assign Pre-Authentication Role:



Captive-Portal

MOD 6- 21

In the wizard's Access section, you can adjust the post-authentication role; for example, only allow access to the internet and deny access to the corporate network. Notice at the bottom that pre-authentication role is automatically selected and the name of that role matches the captive portal profile name.

Pre-Authentication Role on the GW

Access Points Gateways SELECTED GROUP TYPE Mobility Gateway Summary List Config

System Interface **Security** Routing High Availability Config Audit Basic Mode

Roles Policies Aliases Applications Apply Policy Auth Servers Role Assignment (AAA Profiles) ✓

Captive-Portal > Policy > captive-portal Rules ⓘ Drag rows to re-order

PRIORITY	IP VERSION	SOURCE	DESTINATION	SERVICE/APPLICATION	ACTION
1	IPv4	any	captive-portal_guest.company.com	svc-https	permit
2	IPv4	any	captive-portal_guest.company.com	svc-http	permit
3	IPv4	any	any	svc-dhcp	permit
4	IPv4	any	any	svc-dns	permit
5	IPv4	any	any	svc-icmp	permit

MOD 6-22

You can review the pre-authentication role on the gateway. Navigate to gateway group > Config > Security > Roles. Notice that the pre-authentication role in the slide matches the captive portal profile configured when using the wizard. Take a look at the rules created by wizard for that role: they provide limited network access, such as DHCP, DNS, ICMP and HTTP(s) to external splash page server.

Pre-Authentication Role on the AP

The screenshot displays the Cisco Meraki configuration interface for Access Points. The 'Security' tab is active, showing a list of roles on the left and a detailed view of the 'Captive-Portal' role on the right. The 'Captive-Portal' role is highlighted with an orange box. The 'Access Rules For Selected Roles' section lists several rules, each with a status indicator (green dot for enabled, grey dot for disabled), a priority arrow, a checkmark, and edit/delete icons. The rules are:

- Enforce Captive Portal (Enabled, Priority 1)
- Allow https on domain guest.company.com (Enabled, Priority 2)
- Allow http on domain guest.company.com (Enabled, Priority 3)
- Allow dhcp to all destinations (Enabled, Priority 4)
- Allow dns to all destinations (Enabled, Priority 5)
- Allow icmp to all destinations (Enabled, Priority 6)

You can also review pre-authentication role on the AP. Select an AP group context and navigate to Devices > Access Points > Config > Security > Roles. The name of the pre-authentication role is the same as the captive portal profile created earlier using the wizard. Since the AP is responsible for the web redirection, you will find access rules that provide limited access and enforce that redirection to the splash page.

Pre-Authentication Role is the Initial Role

The screenshot displays the Cisco Meraki configuration interface for a Mobility Gateway. The top navigation bar includes 'Access Points' and 'Gateways' (selected). The 'SELECTED GROUP TYPE' is 'Mobility Gateway'. The main menu includes 'System', 'Interface', 'Security' (selected), 'Routing', 'High Availability', and 'Config Audit'. The 'Basic Mode' toggle is visible. The 'Roles' tab is active, showing 'Role Assignment (AAA Profiles)'. The left sidebar lists 'AAA Profile' and 'AOS10-Guest_#1669150621170_155#_'. The main content area shows the configuration for the selected profile, with the 'Initial role' set to 'Captive-Portal' (highlighted with an orange box). Other settings include 'MAC Authentication Default Role' and '802.1X Authentication Default Role' both set to 'guest', and checkboxes for 'Set username from dhcp option 12' and 'L2 Authentication Fail Through'.

MOD 6-24

The WLAN configuration wizard will set the pre-authentication role as initial role in the AAA Profile associated with the guest network. This means that every unauthenticated guest client will initially be assigned this role.

Post-Authentication (Guest) Role

The screenshot displays the Cisco Meraki configuration interface for a Mobility Gateway. The top navigation bar includes 'Access Points' and 'Gateways' (selected). The 'SELECTED GROUP TYPE' is 'Mobility Gateway'. The left sidebar shows tabs for 'System', 'Interface', 'Security' (selected), 'Routing', 'High Availability', and 'Config Audit'. The main content area is titled 'L3 Authentication' and shows a list of 'Captive Portal Authentication Profile' entries. The profile 'AOS10-Guest_#1668436...' is selected. The configuration details for this profile are shown on the right, including 'Default Role' (AOS10-Guest, highlighted with an orange box), 'Default Guest Role' (guest), 'Redirect Pause' (10 sec), 'User Login' (checked), and 'Guest Login' (unchecked). A watermark 'Confidential - For Training Purposes Only' is visible across the page.

L3 Authentication

Captive Portal Authentication Profile : AOS10-Guest_#1668436310799_146#_

Default Role: AOS10-Guest

Default Guest Role: guest

Redirect Pause: 10 sec

User Login: ☒

Guest Login: ☐

MOD 6- 25

You can review the post-authentication role if you access the L3 authentication profile associated with the guest network. This role will be pushed from the gateway to the AP after the client successfully authenticates.




MAC Caching

This section introduces the MAC caching feature that is commonly used with guest access. Use MAC caching is used to improve certain end user experiences.

Confidential for Trial


The Concern With Captive Portal



MOD 6- 27

The captive portal process is simple for users to use, since it launches a pop-up web page and then login with their initial connection. This works very well for devices that connect to the network and stay active. However, most battery powered devices disconnect from the network when they go to sleep, like smart phones. When these lightweight devices wake up and reconnect to the network, they must authenticate to the guest network again, which means executing the captive portal again and can be very frustrating to users.

The Better Solution



A better solution is to combine a MAC authentication request and a captive portal request. This combination provides for a better user experience than just captive portal alone. You will execute the captive portal on the initial connection and then subsequent connections will execute a MAC authentication.

When the lightweight devices wake up and reconnect to the network, the AP will execute a MAC authentication and the user devices will gain access to the network without the need for the user to negotiate the captive portal the second time.

MAC authentication still affords control, as the users MAC address is associated to their guest account. If either the guest account or MAC address is disabled, the MAC authentication will fail. This method is not without risk, as MAC addresses can be spoofed. You should only use this method for true guest access, which does not expose any internal resources.

MAC Authentication Setup

Networks > Configuration - AOS10-Guest

General VLANs **Security** Access Summary

Security Level: 

Enterprise Personal Visitors Open

✓ **Advanced Settings**

Captive Portal Proxy Server IP: 

Captive Portal Proxy Server Port: 

MAC Authentication: ☒

Use IP for Calling Station ID: ☐

MOD 6- 29

MAC caching requires enabling MAC authentication when configuring the WLAN profile. When using the configuration wizard, in the Security section enable MAC Authentication.



Knowledge Check

Self-check on key learning points



It's time for a knowledge check to review what you have learned.

Question #1

Which of the following statements accurately describe the Aruba guest access solution?

- a. A single architecture supports both employee and guest access
- b. It is often configured as a software option
- c. Guest traffic is tunneled to a mobility gateway located in the core or distribution layer of your network
- d. Captive portal login screens can be served directly from the mobility gateway



Knowledge Check

MOD 6- 31




a,b,d

Confidential - For Training Purposes Only

Question #2

Match each captive portal option for guest networks with their appropriate descriptions.



Knowledge Check ✓



MOD 6- 32



Cloud guest

Nice when you need an easy guest solution that is part of Aruba Central

Native support for splash page customization

External captive portal

Great when you require a solution that is not part of Aruba Central, like Aruba ClearPass Policy Manager (CPPM)

You need more involved customization options and scalability

Question #3

Your organization prefers that you implement a guest solution where employee-guest separation is provided at the gateway cluster. Which architecture should you use?

- a. Single-cluster architecture
- b. Three-tier architecture
- c. Multi-cluster architecture
- d. Two-tier architecture with a DMZ



Knowledge Check

MOD 6- 33



Question #4

With an Aruba guest access solution, which services or protocols must you allow to provide for initial guest access, before authentication?

- a. DHCP
- b. HTTP/HTTPS access to the DMZ
- c. Dynamic DNS (DDNS)
- d. DNS
- e. HTTP/HTTPS access to the ClearPass server



Knowledge Check

MOD 6- 34



a, d, e

Confidential - For Training Purposes Only



Lab Activity

Let's review the lab activity that is associated with this module.

Confidential - For Training

Lab Tasks

Lab 06.01 - Overlay Guest WLAN with CPPM Guest

Task 1

- Verify a ClearPass Guest page

Task 2

- Configure WLAN profile with ClearPass Guest Splash page

Task 3

- Test ClearPass Guest access

Task 4

- Guest Authentication with ClearPass MAC Caching

Task 5

- Web Redirect for Corporate User

**Module 6 -
Wireless
Authentication
for IOT**



MOD 6- 36

The slide provides a brief review of tasks performed in this lab. Please see your Lab guide for details of this lab.

Rev. 23.11




Wireless Authentication for IoT

Implementing Aruba Campus Access

Confidential For Training

This is module 7 – Wireless Authentication for the Internet of Things (IoT). Here you'll learn how to securely connect IoT devices to your campus network



Objectives

- Describe MPSK benefits and operation
- Deploy MPSK WLANs to improve security for IoT deployments

MOD 7- 2



As IoT continues to gain traction, your ability to accommodate them becomes more vital. After completing this module, you will be able to describe multi-pre-shared keys (MPSK) benefits and operation. You will be able to deploy MPSK WLANs to improve security for IoT deployments.

The objectives for this module are:

- Describe MPSK benefits and operation.
- Deploy MPSK WLANs to improve security for IoT deployments.

MPSK Introduction

We begin this module with an introduction to MPSK.

Confidential - For Training

MPSK Overview

Challenges

- IoT devices often lack 802.1X support
- PSK is the only supported mode
- All clients in WLAN share same password
- All clients get same role on Aruba AP & GW

Benefits

- Multiple PSK keys in the same WLAN
- Enhanced security
- Ability to assign different roles/VLANs
- ClearPass facilitated device registration



MOD 7- 4

Some companies have headless devices that do not support 802.1X authentication; therefore, these devices must connect to WLANs in WPA/WPA2/WPA3 Personal mode. However, the devices might also lack support for WPA3, so WPA2-Personal, with all of its flaws, is the only option for them. Aruba provides multi pre-shared keys (MPSK) to enhance security for these devices. In MPSK mode, multiple keys are used in the same WLAN; therefore, security is enhanced. When the device connects to MPSK WLAN, the key used during authentication process can be used to assign different roles and/or VLANs. You can also integrate MPSK with ClearPass for the authentication function!

PSK Deployment Options

Single PSK

- All devices use same key

When key compromised **ALL** connected devices must be reconfigured

Multiple PSK Local

- Group of devices use dedicated key
- Key management in Aruba Central
- Optional use of MAC authentication

When key compromised, **GROUP** of devices must be reconfigured

Multiple PSK ClearPass

- Every device, unique key
- Requires device MAC registration
- Optional key based on device class
- Key mgmt in ClearPass
- Requires MAC auth

When key compromised, **ONLY ONE** device must be reconfigured (when using per device unique key)




MOD 7-5

This figure in the slide summarizes all PSK deployment options:

- Single PSK is the least secure option since all devices use the same key. If the key is compromised it must be changed and ALL devices need to be reconfigured.
- Multiple PSK local mode is more secure as there are only group of devices that use the same key. If key is compromised it must also be changed by the devices that use that key.
- Multiple PSK with ClearPass is the most secure option. In this scenario you can have a unique key per device, so if the device becomes compromised it will only affect single device.

Multiple Pre-shared Keys



MPSK only possible with WPA2 encryption

MOD 7-6



In MPSK mode, AP has multiple keys configured. When client connects, the client will use its configured key to authenticate. The AP will then work down the list to check if it has matching key. Once a matching key is found AP and client complete authentication process and the client gets connected to that WLAN. The core advantage of MPSK is that traffic generated by the bar code scanner is encrypted using different key than traffic generated by the camera thus security is enhanced.

Note: MPSK is only possible if the network uses WPA2 Personal encryption mode.

MPSK Roles & VLANs

PSK : **scanner!123**

Role: **scanner**

VLAN: **25**



PSK : **camera!123**

Role: **cctv**

VLAN: **35**



PSK : **sensor!123**

Role: **IoT**


VLAN: **45**



SSID: IoT
MPSK



APs



VLAN 25

VLAN 35

VLAN 45

Gateway cluster



MOD 7- 7

Example PSKs for training purposes **ONLY!** Always use complex keys in production environment

In the example in this slide, the key configured on client device is used as a differentiator, allowing you to assign different VLANs and roles to clients based on the key that is used. All clients authenticated using the scanner!123 key will be assigned to VLAN 25 and the role scanner; at the same time, clients using the camera!123 key will be assigned to VLAN 35 and the role cctv.

This solution provides flexibility to customize the network access for IoT devices, even though they all connect to the same WLAN. Plus, you can don't have to use a separate SSID for each device type: you can place them in one SSID, making a much more efficient WLAN solution.


MPSK Authentication Process



Let's explore the how the MPSK authentication process works.

Confidential - For Training

MPSK with ClearPass



The ClearPass MPSK solution requires device registration. A user registers the device MAC address in ClearPass and enters the PSK key for that device. When the device connects, the AP sends a MAC authentication request to ClearPass, which will look up the device MAC address and PSK key in its database (note that this information can also be stored in an external database that ClearPass accesses). After successful authentication, ClearPass sends an Access-Accept message with a VSA attribute that contains the PSK for the device. This way, the AP dynamically receives the correct PSK that will be used to generate the wireless encryption keys.

The slide shows a high-level overview of how the solution typically works:

Before attempting to connect the headless device, a user connects to the ClearPass Self Service Registration page using a different connection. On that page, the user authenticates and specifies the device's MAC address. ClearPass assigns the device a passphrase, or PSK, and displays that passphrase to the user.

The user then configures the client device to use the specified passphrase.

The client can now connect to the WPA2 Personal WLAN. From the client's point of view the connection proceeds like a typical WPA2 Personal connection.

However, at the point at which the client connects to WLAN, the AP sends MAC authentication request and gateway forwards it to ClearPass. ClearPass responds with the

passphrase assigned to the client in an Aruba-MPSK-Passphrase VSA. This is then forwarded by the gateway to the AP. The Aruba AP then performs the 4-way handshake using the PMK derived from the passphrase individual to the client device.

Confidential - For Training Purposes Only

Local MPSK



Local MPSK profile

Name	Key	Role
Scanner	scanner!123	iot
Camera	camera!123	cctv

AAA profile IoT_#1565656....


User Derivation Rules

- if mpsk-key-name = **Scanner** -> role **iot**
- if mpsk-key-name = **Camera** -> role **cctv**

MOD 7- 10

With local MPSK, you create a MPSK profile, where you define a PSK key name, the actual key, and user role. MPSK key configuration will be then pushed to the AP, while the gateway receives key to role mapping. This mapping is implemented on the gateway using the user derivation rules which maps the RADIUS mpsk-key-name attribute to the user role. User derivation rules are then mapped to a AAA profile specific to the MPSK WLAN that you have created.

Local MPSK Flow



Let's take a look at the details when implementing local MPSK. After the client connects to the MPSK WLAN, the client will use a PSK key for authentication. The AP has multiple keys configured: it will match the client key with the locally defined key. If no matching key is found, client access is denied. If a matching local key is found, the AP will send a RADIUS Access-Request message to the gateway, where this message contains the VSA attribute Aruba-Mpsk-Key-Name. The value of this attribute corresponds with the key name used by the client when connecting to WLAN. The gateway will process the Access-Request against the AAA profile associated with WLAN. This AAA profile has user derivation rules configured that will be compared to the Aruba-Mpsk-Key-Name attribute; and upon a match, the associated user role and VLAN is assigned to the client. The gateway sends an Access-Accept message to the AP that contains the Aruba-User-Role and Aruba-User-VLAN attributes. Now the AP knows what user role and VLAN to assign to the client.

Local MPSK & ClearPass

Combine local MPSK with MAC authentication

- MPSK done locally
- Client MAC used for MAC-auth against ClearPass

Take advantage of ClearPass decision making process


- Enforcement profiles match multiple attributes
- ClearPass profiling eliminates device impersonation
- Incorporate Aruba-MPSK-Key-Name attribute in ClearPass policy



MOD 7- 12

You can combine local MPSK with RADIUS-based MAC authentication. MPSK keys are still configured locally on the AP but the authentication process will be performed against a ClearPass server. When using ClearPass, you can take advantage of its advanced decision-making process. With ClearPass you can match multiple attributes and properties, such as device type, MAC OUI, time-of-day, device location, and etcetera. The ClearPass profiling function lets you protect your network against device impersonation. ClearPass can be made aware of the key used by the device upon connecting to the WLAN thanks to Aruba-Mpsk-Key-Name attribute that will be send from the Gateway to ClearPass. And along with the other information that you can optionally have ClearPass examine, the server can pass the appropriate RADIUS attributes to the gateway for the device's session.

Local MPSK & ClearPass Flow



From the client and AP perspective, there is no difference to the process. The AP sends a RADIUS Access-Request that contains the Aruba-Mpsk-Key-Name attribute. With MAC authentication enabled on the gateway this request is forwarded to ClearPass and is processed by a ClearPass policy.

A ClearPass policy is where you can implement a very complex decision-making process. In the example shown in the slide, the ClearPass server checks the client device MAC OUI and checks the key the client used to connect to the WLAN. If device MAC OUI corresponds with a supported scanner device and the key used by the device is applicable to the scanner type of device, ClearPass will allow access. ClearPass then sends back an Access-Accept message that will contain attributes to assign the user role and VLAN.


Local MPSK Configuration



Now that you understand how MPSK works, let's discuss how to configure it.

Confidential - For Training

Create IoT Roles on the Gateway




The screenshot shows the configuration page for a Mobility Gateway. The 'Gateways' section is selected, and the 'Security' tab is active. Under 'Security', the 'Roles' sub-tab is chosen. The table below lists existing roles:

NAME	RULES	GLOBAL
ap-role	35 Rules	No
authenticated	4 Rules	No
Camera	2 Rules	No
default-iap-user-role	2 Rules	No
default-via-role	3 Rules	No
default-vpn-role	4 Rules	No

A red box highlights the '+' button at the bottom left of the table, indicating where to click to add a new role.

You will need to create the appropriate roles on the gateway. To create the roles for devices connecting to the MPSK WLAN, select the gateway group context and navigate to Devices > Config > Security > Roles. Click the + sign to add the new roles for the IoT or user devices.

Setup MPSK Profiles



MOD 7- 16

Next, you need to create the MPSK key profiles by selecting the AP Group context and navigating to Devices > Access Points > Security > MPSK Local.

Setup MPSK Profiles

MPSK Local

Name : IoT-MPSK

* MPSK local passphrase table should not be empty

MPSK Local Passphrase

Name	Role

MPSK Local Passphrase

Name : Sensor

Passphrase :

Retype Passphrase :

Role : Sensor

Fields do not match

MOD 7- 17

Specify MPSK local profile name. Then for each key in the profile, define the key name, enter the actual key twice, and select the role that will be mapped to that key.

Configure Local MPSK WLAN

The screenshot shows the 'Create a New Network' wizard in the Aruba configuration interface. The 'Security' step is selected, indicated by a green arrow. The 'Security Level' is set to 'Personal' on a slider between 'Enterprise' and 'Open'. The 'Key Management' dropdown is set to 'MPSK Local'. Below it, the 'MPSK Local' dropdown is set to 'AOS10-MPSK-Local'. A link 'To add/edit MPSK local profile' is visible. The interface includes a top navigation bar with 'Access Points' and 'Gateways' tabs, and a bottom navigation bar with various configuration categories like 'WLANs', 'Access Points', 'Radios', etc. A 'Hide Advanced' button is also present.

Access Points Gateways Summary List Config

WLANs Access Points Radios Interfaces Security Services System IoT Configuration Audit Hide Advanced

Create a New Network

1 General 2 VLANs 3 Security 4 Access 5 Summary

Security Level: Enterprise Personal Visitors Open

Key Management: MPSK Local

MPSK Local: AOS10-MPSK-Local To add/edit MPSK local profile

Next, create the WLAN that will use MPSK authentication. In WLAN wizard select Personal as the security option. For Key Management, select MPSK Local and select previously configured MPSK local profile.

As you can see, setting up MPSK is a simple and straight-forward process.

MOD 7-18
+

AAA Profile with User Rules

The screenshot displays the configuration page for an AAA Profile within a network management system. The interface is organized into several sections:

- Top Navigation:** Includes 'Access Points' and 'Gateways' tabs. The 'Gateways' tab is active, showing 'SELECTED GROUP TYPE: Mobility Gateway'. There are also links for 'Summary', 'List', and 'Config'.
- Sub-Navigation:** A row of tabs includes 'System', 'Interface', 'Security' (which is selected), 'Routing', 'High Availability', and 'Config Audit'.
- Left Panel:** A list of configuration items for the selected group. The first item, 'IoT_#1668515269077_1...', is highlighted. Below it are several server and profile entries: '802.1X Authentication Profile', 'MAC Authentication Profile', '802.1X Authentication Server Group', 'MAC Authentication Server Group', 'RADIUS Accounting Server Group', 'RFC 3576 server', and 'XML API server'.
- Right Panel:** The 'Role Assignment (AAA Profiles)' section. It contains several configuration options:
 - '802.1X Authentication Default Role:' set to 'guest'.
 - 'Set username from dhcp option 12:' with an unchecked checkbox.
 - 'L2 Authentication Fail Through:' with an unchecked checkbox.
 - 'Multiple Server Accounting:' with an unchecked checkbox.
 - 'User idle timeout:' with an empty text input field.
 - 'RADIUS Interim Accounting:' with an unchecked checkbox.
 - 'User derivation rules:' with a dropdown menu showing 'iot_#1668515269077_155#_'. This dropdown is highlighted with an orange box.

A large diagonal watermark reading 'Confidential For Training Purposes Only' is visible across the center of the interface. In the bottom right corner, the text 'MOD 7- 19' is displayed next to a small orange plus icon.

You can review the AAA profile associated to WLAN from the gateway group context; then navigate to Devices > Gateways > Security > Role Assignments (AAA Profiles). In the slide, notice the user derivation rules assigned to AAA profile: this is automatically configured by WLAN wizard.

Local User Derivation Rules

The screenshot displays the Aruba configuration interface for a Mobility Gateway. The navigation path is: Gateways > Security > Advanced > Local User Derivation Rules. The 'Advanced' tab is selected, and the 'Local User Derivation Rules' section is expanded. A table shows a rule named 'iot_#1668515269077_155#_' with 2 rules. Below, a 'Rules-set' table shows two rules mapping 'mpsk-key-name' to roles 'Sensor' and 'Camera'.

NAME	NO. OF RULES
iot_#1668515269077_155#_	2

PRIORITY	ATTRIBUTE	OPERATOR	OPERAND	ACTION	VALUE	TOTAL HIT	NEW HIT	DESCRIPTION
1	mpsk-key-name	equals	Sensor	role	Sensor	--	--	--
2	mpsk-key-name	equals	Camera	role	Camera	--	--	--

MOD 7- 20

You can review the user derivation rules from the gateway group context by navigating to Devices > Gateways > Security > Advanced > Local User Derivation Rules. In the slide, notice that the attribute mpsk-key-name is mapped to a user role. Again, the wizard created this configuration when you defined the MPSK WLAN during the wizard process.



Knowledge Check

Self-check on key learning points



It's time for a knowledge check to review what you have learned.

Confidential - For Training Purposes Only

Question #1

You can configure multiple PSKs locally on a group of APs, such that if a key is compromised, only one device must be reconfigured.

- True
- False

Q

Knowledge Check

MOD 7-22

+

false

Confidential - For Training Purposes Only

Question #2

Using ClearPass with multi-pre-shared keys (MPSK), you can take advantage of a more advanced decision-making process, such as matching multiple attributes like device type, MAC OUI, time-of-day, and location.

- True
- False



Knowledge Check

MOD 7- 23



true

Confidential - For Training Purposes Only

Question #3

When you configure MPSK locally for a group of APs, you must also use MAC authentication.

- True
- False



Knowledge Check

MOD 7- 24



false

Confidential - For Training Purposes Only



Lab Activity

Let's review the lab activity that is associated with this module.

Confidential - For Training

Lab Tasks

Lab 07.01 – PSK IOT WLAN

Task 1

- Create MPSK Local Overlay WLAN

Task 2

- Configure ClearPass based Role Mapping for MPSK

**Module 8 -
Mixed Mode
Architecture**




MOD 7- 26

The slide provides a brief review of tasks performed in this lab. Please see your Lab guide for details of this lab.

Mixed Mode Architecture

Implementing Aruba Campus Access

In the Aruba Campus Access Fundamentals course and the corresponding HPE Press book, you learned how to integrate APs into Central without a gateway. And when implementing WLANs, you learned how to use bridge mode for the device access. Then, in module 4, you learned about adding gateways into the campus design and using tunnel mode to have the APs tunnel traffic, via GRE or GRE over IPsec, to the gateways for processing. In this module you'll learn how to combine both of these functions in one WLAN: this is referred to as mixed mode.



Objectives

- Leverage the advantages of Mixed Mode WLANs to improve security for IoT and other devices
- Configure Mixed Mode WLANs

After completing this module, you will be able to:

- Leverage the advantages of mixed mode WLANs to improve security for IoT and other devices.
- Configure mixed mode WLANs.

Aruba Mixed Mode WLAN


Let's first explore what mixed mode WLANs actually are and how they work.

Confidential - For Training

Introduction


Bridge & tunnel mode

2 SSIDs: One bridge, one tunnel



Mixed mode

1 SSID: Bridge & tunnel



Reduced number of SSIDs in Campus

MOD 8- 4

Mixed mode SSIDs enable both bridge and tunnel forwarding modes within a single SSID. Reducing SSIDs on a campus increases WLAN performance by reducing the number of management and beacon frames transmitted. Mixed mode SSIDs, however, only support 802.1X authentication. They should be considered if the design requires both bridged and tunneled traffic within the campus where 802.1X authentication is used. AOS 10 mixed mode, as you will see, is very different from AOS 8's split-tunnel mode.

Mixed Mode

Single SSID

- Forwarding modes: bridge or tunnel
- Decision basis: authentication attributes or MAC address analysis
- Policy based traffic forwarding

VLAN assignment rules

- Default forwarding is tunneled
- At least one rule required for bridge mode

VLAN assignment rule types

- **MAC address analysis:** all auth. types: 802.1X, PSK, Captive Portal, Open
- **Authentication attribute analysis:** 802.1X only



MOD 8- 5


AOS 10 mixed mode dynamically assigns the forwarding mode for each device connected to the SSID. The decision is taken by the AP side and it is based on the authentication results analysis or on the value of the client's MAC address.

VLAN assignment rules are created to control the used of the forwarding mode. By default, there is an implicit rule which forwards all traffic to the gateway: therefore, all traffic is tunneled. To change this default behavior, you need to define at least one rule for bridging traffic.

The use of the authentication result decides the forwarding mode: this typically requires the SSID to implement 802.1X authentication (WPA3 Enterprise, for example). MAC address analysis is also available in all type of authentication (802.1X, PSK, captive portal and open).

Mixed Forwarding Mode

WLAN: Employee
Mixed mode
802.1X




MOD 8-6

In mixed mode the AOS 10 AP decides which forwarding role is applied to each device that connects to the SSID. This decision is based on the authentication result (particularly on the RADIUS attributes received from the AAA authentication server) or in the client's MAC address.


One advantage of mixed mode is that you are not required to broadcast two SSIDs to use two forwarding modes, this represents a RF resource saving for the APs. Mixed mode thus allows the WLAN to perform more efficiently, resulting in a more positive user experience.

Mixed Forwarding Mode Architecture



In mixed mode, it all comes down to what VLAN the gateway assigns to the client connected to the AP. As you learned earlier, AOS 10 APs use a single GRE tunnel for all SSIDs and that GRE tunnel has tagged VLANs enabled. All VLANs that exist on the gateway are allowed on the GRE tunnel. This is independent of the VLANs that are configured in the WLAN profiles. If the client VLAN that the gateway sends to the AP is active on the GRE tunnel, the client traffic will be tunneled. Otherwise, the client traffic will be locally bridged by the AP. Bridged VLANs must be enabled on the switch trunk port where the AP connects. The VLAN plan should have a separate set of VLANs for bridged and tunnel forwarding. In other words, a VLAN that is used in one SSID for tunneled forwarding should not be used in another SSID for bridging.

Mixed Forwarding Mode Flow



MOD 8- 8



Look at the example in the slide. The client connects to the mixed mode SSID. The AP sends a RADIUS Access-Request to the gateway. Depending on AAA profile configuration, that request is forwarded to the RADIUS server. After successful authentication, the RADIUS server sends an Access-Accept message that may contain the Aruba-User-Role attribute to assign the user role to the client. The gateway forwards the Access-Accept message to the AP and adds the Aruba-User-Vlan=45 attribute due to AAA profile default VLAN configuration: the VLAN assignment rules or RADIUS attribute send from RADIUS server. VLAN 45 doesn't exist on the gateway in this example and thus is not enabled on the GRE tunnel: therefore, traffic will be bridged by the AP. The AP will forward user traffic on it's local Ethernet port and add VLAN 45 802.1Q tag for this client.

Note: Make sure that the switch has VLAN 45 enabled on the switch port connected to the AP!

Bridged VLAN vs GRE Tunneled VLAN

Tunneled to the gateway

```
AP-1# show ata endpoint
```

```
ATA Endpoint Status
```

PORT(SRC/DST)	VALID TIME(s)	TYPE	GRE VLANs	HBT(Jiff/Missed/Sent/Rcv)	INNER IP
4500/4500	45567	GRE	1, 15, 25, 35	102632/0/86657/83860	10.1.10.1

```
Total Endpoints Count: 1
```

All VLANs not enabled on the GRE will be bridged

MOD 8-9

You can check what VLANs are enabled and supported on the GRE tunnel by using the `show ata endpoint` command on the AP. As you can clearly in the slide see VLAN 45 is not present on GRE tunnel: if that VLAN gets assigned to a client, the corresponding traffic for VLAN 45 will be bridged.

Mixed Mode Configuration



Now that you know how mixed mode WLANs function, let's learn how to configure them.

Confidential - For Training

Mixed Mode WLAN Setup

The screenshot shows the 'Mixed Mode WLAN Setup' wizard in a network management interface. The interface has a top navigation bar with 'Access Points' and 'Gateways' tabs. Below this is a sub-navigation bar with tabs for 'WLANs', 'Access Points', 'Radios', 'Interfaces', 'Security', 'Services', 'System', 'IoT', and 'Configuration Audit'. The 'WLANs' tab is selected. On the right side of the top bar, there are links for 'Summary', 'List', and 'Config'. A 'Hide Advanced' link is also present on the right. The main content area is titled 'Create a New Network' and features a progress bar with five steps: 1 General, 2 VLANs (current step), 3 Security, 4 Access, and 5 Summary. The 'VLANs' step is highlighted with a green arrow. Below the progress bar, there are four configuration options: 'Traffic forwarding mode:' with radio buttons for 'Bridge', 'Tunnel', and 'Mixed' (selected); 'Primary Gateway Cluster:' with a dropdown menu showing 'Campus1:auto_gwcluster_237_0'; 'Secondary Gateway Cluster:' with a dropdown menu showing 'None'; and 'Client VLAN Assignment:' with radio buttons for 'Dynamic' (selected) and another option. The bottom right corner of the interface shows 'MOD 8-11' and a small red plus icon.

Access Points Gateways Summary List Config

WLANs Access Points Radios Interfaces Security Services System IoT Configuration Audit Hide Advanced

Create a New Network

1 General 2 VLANs 3 Security 4 Access 5 Summary

Traffic forwarding mode: ☐ Bridge ☐ Tunnel ☒ Mixed

Primary Gateway Cluster: Campus1:auto_gwcluster_237_0 ▼

Secondary Gateway Cluster: None ▼

Client VLAN Assignment: ☒ Dynamic

MOD 8-11

To implement a WLAN using the mixed forwarding mode, create a new SSID using Central's WLAN wizard. On the second step of the wizard, select Mixed for the traffic forwarding mode. The system requires that you specify a primary gateway cluster, where the secondary gateway cluster is optional.

Bridged VLAN Setup

The screenshot displays the Aruba configuration interface for VLAN management. The 'WLANs' tab is active, showing a list of VLANs: vlan45, vlan15, vlan25, and vlan35. A dialog box titled 'Add Named VLAN' is open, allowing the user to create a new VLAN. The dialog box contains three fields: 'VLAN Name:' (set to 'vlan45'), 'VLAN:' (set to '45'), and 'VLAN Type:' (set to 'Bridge'). The 'Add Named VLAN' button is highlighted with an orange box. The background shows the 'WLANs' tab with a list of VLANs and a '+ Add Named VLAN' button at the bottom.

VLAN Name:	VLAN:	VLAN Type:
vlan45	45	Bridge

Buttons: Cancel, OK

Background: + Add Named VLAN, 4 Named VLAN(s)

MOD 8-12

You can create bridge VLANs using the wizard or simply enter the Bridge VLAN ID manually when setting up WLAN the profile. Bridge VLANs that you create will not be pushed to the gateway nor will be enabled on the GRE tunnel.

Adding VLAN Rule

New VLAN Assignment Rule

Attribute: Operator: String: VLAN Type: VLAN:

Aruba-User-R equals IoT-AC Bridge vlan45(45)

Aruba-Template-User
Aruba-User-Role
Aruba-User-Vlan
Aruba-User-Group
Aruba-WorkSpace-App-Name
Authentication-Sub-Type
Authentication-Type
CHAP-Challenge
Callback-Id
Callback-Number
Class
Client MAC Address
Connect-Info
Connect-Rate

OK

You could enforce bridge forwarding by sending the Aruba-User-Vlan attribute from the RADIUS server: just make sure that the bridge VLAN doesn't exist on the gateway. This would be the most common way of implementing mixed mode.

However, VLAN rules are not really required to enable bridging, but they are simply one way to do it.

To create a new VLAN assignment rule, click + Add Rule under VLAN Assignment Rules, which will open the New VLAN Assignment Rule window.

1. Select an attribute from the Attribute drop-down list.
2. Select the Operator drop-down list, depending on your policy criteria. Enter a string in the String text box.
3. Select a VLAN ID from the VLAN drop-down list.
4. Click OK to save the changes.

As you can see, if you don't have an AAA server to implement mixed mode, but using an AAA server is more commonly done.

Default VLAN

Access Points

Gateways

Summary

List

Config

WLANs

Access Points

Radios

Interfaces

Security

Services

System

IoT

Configuration Audit

Hide Advanced

VLAN Assignment Rules

Default VLAN: 15

If Aruba-User-Role equals iot-AC assign bridge VLAN: vlan45

If Aruba-User-Vlan equals 45 assign bridge VLAN: vlan45

+ Add Rule

3 Rule(s)

MOD 8-14

You can define a default VLAN that will be assigned if none of the rules are matched. Based on whether this VLAN exists on the gateway or not will determine if the traffic for this VLAN is tunneled or bridged.

Mixed Mode Server Rules

The screenshot displays the Aruba configuration interface for a Gateway. The top navigation bar includes 'Access Points' and 'Gateways' (selected). The 'SELECTED GROUP' is 'Mobility Gateway'. The main menu has tabs for 'System', 'Interface', 'Security' (selected), 'Routing', 'High Availability', and 'Config Audit'. The 'Basic Mode' indicator is present. Below the main menu, there are sub-tabs: 'Roles', 'Policies', 'Aliases', 'Applications', 'Apply Policy', and 'Auth Servers' (selected). The 'Auth Servers' section shows a table of rules for the 'Server Group > AOS-Mixed_#1671185175569_155#_auth_svg'. The table has columns: ATTRIBUTE, OPERATION, OPERAND, TYPE, ACTION, and VALUE. Two rules are listed: 'Filter-Id' with value '45' and 'Aruba-User-Role' with value 'vlan25'. A 'Drag rows to re-order' handle is visible on the right. The bottom right corner shows 'MOD 8-15'.

ATTRIBUTE	OPERATION	OPERAND	TYPE	ACTION	VALUE
Filter-Id	equals	iot	string	set vlan	45
Aruba-User-Role	equals	employee	string	set vlan	vlan25

The VLAN assignment rules you define in WLAN wizard get implemented on a gateway as RADIUS server rules. You can review these rules if you select the gateway group context and navigate to

Devices > Gateways > Security > Server Group > Server Rules.

Mixed Mode Security Options

The screenshot shows the 'WLANs' configuration page for 'Access Points'. The 'Security' tab is selected in the wizard steps (1 General, 2 VLANs, 3 Security, 4 Access, 5 Summary). Under 'Security Level', a slider is positioned at 'Enterprise'. Below the slider, a warning message is displayed in a red-bordered box: 'Personal, Visitors and Open modes are disabled since attribute of vlan and role assignment rules have items other than mac-address'. The 'Key Management' dropdown is set to 'WPA3-Enterprise(CCM 128)'. The 'Primary Server' dropdown is set to '- Select -' and is marked as mandatory with a red error message 'This field is mandatory.'.

Access Points Gateways

WLANs Access Points Radios Interfaces Security Services System IoT Configuration Audit

1 General 2 VLANs 3 Security 4 Access 5 Summary

Security Level:

Enterprise Personal Visitors Open

⚠ Personal, Visitors and Open modes are disabled since attribute of vlan and role assignment rules have items other than mac-address

Key Management: WPA3-Enterprise(CCM 128) ▼

Primary Server: - Select - ▼ + This field is mandatory.

If you select mixed mode in the WLAN wizard and use RADIUS attributes in VLAN assignment rules, security options will be limited to Enterprise 802.1X. The only way to see other security options available is to create VLAN assignment rules based on client MAC addresses only.



Knowledge Check

Self-check on key learning points



It's time for a knowledge check to review what you have learned.

Confidential - For Training Purposes Only

Question #1

Which of the following options accurately describe Aruba mixed mode WLANs?

- a. They enable both bridge and tunnel forwarding modes on a single SSID
- b. They increase performance by reducing the number of management and beacon frames transmitted
- c. They can analyze attributes for guests using WebAuth and for employees using 802.1X authentication
- d. The Aruba gateway decides which VLAN and forwarding role is applied to each device

a

Knowledge Check

MOD 8- 18
+

a,b

Question #2

You have deployed a mixed mode WLAN in one of your offices that involved VLANs 1, 5, 10, 15, 20, 25, 30, and 35. Based on the output below, which VLANs are bridged?

AP-1# **show ata endpoint**

ATA Endpoint Status

PORT (SRC/DST)	VALID TIME(s)	TYPE	GRE VLANs	HBT (Jiff/Missed/Sent/Rcv)	INNER IP
4500/4500	45567	GRE	1, 15, 25, 35	102632/0/86657/83860	10.1.10.1

- a. 1, 5, 10, 15, 20, 25, 30, and 35
- b. 1, 15, 25, and 35
- c. 5, 10, 20, and 30
- d. 5, 10, 15, 20, 25, 30, and 35

a

Knowledge Check

MOD 8- 19
+

c



Lab Activity

Let's review the lab activity that is associated with this module.

Confidential - For Training

Lab Tasks

Lab 08.01 -Configure bridged VLAN for tunneled WLAN

Task 1

- Employee WLAN with mixed mode

Task 2

- RADIUS based VLAN Assignment

Task 3

- Custom RADIUS attribute in VLAN rule

**Module 9 -
Gateway
Cluster
Deployments**



MOD 8- 21

The slide provides a brief review of tasks performed in this lab. Please see your Lab guide for details of this lab.


Rev. 23.11



Gateway Cluster Deployments

Implementing Aruba Campus Access

You were introduced to gateway clustering in module 3. In module 4 you learned how clusters operate and how to monitor them. In this module you'll learn about the different deployment options related to clustering as well as how to configure auto and manual clusters.



Objectives

- Describe and compare gateway cluster deployment options
- Configure gateway clusters
- Monitor gateway clusters and assess clustering status
- Explain secondary cluster redundancy

After completing this module, you will be able to:

- Describe and compare gateway cluster deployment options.
- Configure gateway clusters.
- Monitor gateway clusters and assess clustering status.
- Explain secondary cluster redundancy.


Cluster Deployment Options

Let's begin this module by exploring gateway clustering options to ensure you choose the right cluster deployment for your scenario.

Confidential for Training

Gateway Cluster Review

- Seamless roaming
- Client/AP load balancing
- Hitless failover
- Client state sync
- Ease of deployment



MOD 9- 4

A gateway cluster combines multiple Aruba gateways operating as a single entity to provide high availability and service continuity to the WLAN clients in a network. A gateway cluster combines multiple Aruba gateways operating as a single entity to provide high availability and service continuity to the WLAN clients in a network. Gateway clusters provide full redundancy to access points (APs) and WLAN clients in the event of a failover. AOS 10.x supports gateway clusters and provides the following features and benefits:

Seamless roaming: when a client roams between APs within the WLAN boundary, the clients remain anchored to the same gateway in a cluster for a seamless user experience.

Client/AP Load Balancing: if there is excessive workload among the devices, the gateways in a cluster balance AP and client loads to ensure faster connectivity to clients.

Hitless Failover: when a gateway device fails, APs and endpoints failover to another gateway in the cluster without any service disruption.


Client State Sync: synchronizes high-value client sessions such as voice, video, FTP and IGMP (used for IPv4 multicast groups) and MLD (used for IPv6 multicast groups) between active and standby gateways in a cluster. This allows the connected devices to failover to the standby gateway seamlessly. However, AOS 10.x only supports this synchronization for the first failover. For subsequent failovers, it will not synchronize the session to the recent standby, instead, it synchronizes a new session.

Ease of deployment: a gateway cluster is automatically formed when assigned to a UI group in the Network Operations app without any manual configuration.

Confidential - For Training Purposes Only

Single Cluster Deployment

- ✓ **Single** gateway cluster for all WLANs
- ✓ Per each cluster member tunnels




VLANs and firewall ensure client separation

MOD 9- 5

In a single cluster deployment, all tunneled WLANs terminate on the same gateway cluster. APs establish a single GRE tunnel to all of the gateways in that cluster. VLANs and strict firewall rules provide separation between clients.

Multi-Cluster Deployment

- ✓ **Dedicated** gateway cluster for WLAN
- ✓ Multiple tunnels to each gateway cluster



GW cluster and firewall ensure client separation

MOD 9- 6

In a multi-cluster deployment, also known as multi-zone, different WLANs terminate on different gateway clusters. This deployment is recommended if you want to isolate gateways used to terminate different types of users such as employees and guests. Client separation is provided by the firewall but also dedicated gateway cluster.

Auto-Cluster

- ✓ Simplifies cluster setup
- ✓ New gateways automatically joined
- ✓ Two modes: auto-group and auto-site





MOD 9- 7

Aruba Central and AOS 10 introduce an automated cluster formation feature. This option has two sub-modes auto-group and auto-site. As the name refers, the auto-group option will automatically create a cluster of the gateways as long as all members belong to the same Aruba Central group. The auto-site option does a similar process: as long as all gateway members are mapped to the same site. A manual cluster creation option is supported to either select the members of the cluster or modify an already created cluster.

Note: The automated options are mutually exclusive: you use one or the other.

Auto-Group

- Members of the same **GROUP** automatically clustered
- Cluster name automatically generated




MOD 9- 8

In auto-group mode all gateways that are members of the same Central group will be join to the gateway cluster. The name of the cluster is automatically generated, as you can see in the slide.

Auto-Site

- Members of same **GROUP** and same **SITE** auto-clustered
- Cluster name automatically generated



APs must be assigned to the **SAME SITE** as gateways

MOD 9-9

In auto-site mode, the gateways that are assigned to the same group and the same site will be automatically form a gateway cluster.

Note: For APs to terminate on gateway cluster they must be assigned to the same site as the gateways.

Manual Cluster

Manual cluster configuration:



When ?

- Change of authorization redundancy
- MicroBranch AP termination



MOD 9- 10

The manual cluster configuration is the last deployment mode. This mode allows you to enable MicroBranch AP termination and RADIUS change of authorization (CoA) redundancy (CoA was discussed in module 5).

Cluster Configuration



Now that you understand the deployment options and modes, let's talk about how to configure a gateway cluster.

Confidential - For Tra

Automatic Cluster Mode

P24-T01-Gateway...

Manage

Overview

Devices

Clients

Guests

Gateways

System Interface Security Routing **High Availability**

Clusters Redundancy

Automatic: ☒

☒ Auto Group ☐ Auto Site

Warning

Default

If the existing cluster has SSIDs already mapped the change is not possible

Warning

Existing auto-group clusters will be dismantled and gateways will be remapped to new site specific auto-clusters.

Do you want to save and continue?

MOD 9- 12

By default, the automatic cluster option is enabled and the Auto Group mode is selected. If you wish to use the Auto Site mode, the system will display a warning message indicating that the existing cluster will be dismantled.

Note: In case the existing cluster is already used by an SSID then the change of the mode is not possible until you either change the SSID's reference to a different cluster or you delete the SSID.

Per Site Clusters

Access Points Switches Gateways

Gateways 4 Clusters 2

GATEWAY CLUSTERS (2)

	Name	Group	AP Tunnel	Clients	Model												
▼	• auto_gwcluster_site_41_1...	P9T14-GW-Group-2	2	2	A7030												
	<table border="1"> <thead> <tr> <th>Gateway Name</th> <th>AP Tunnel</th> <th>Clients</th> <th>Model</th> </tr> </thead> <tbody> <tr> <td>• GW-4</td> <td>1</td> <td>1</td> <td>A7030</td> </tr> <tr> <td>• GW-5</td> <td>1</td> <td>1</td> <td>A7030</td> </tr> </tbody> </table>					Gateway Name	AP Tunnel	Clients	Model	• GW-4	1	1	A7030	• GW-5	1	1	A7030
Gateway Name	AP Tunnel	Clients	Model														
• GW-4	1	1	A7030														
• GW-5	1	1	A7030														
>	• auto_gwcluster_site_42_1...	P9T14-GW-Group-2	2	0	A7030												

Per site cluster

MOD 9- 13

In the example shown in the slide, you can see how the auto site mode was used and how the system used a representative name: in this case auto_gwcluster_site_41_1 and auto_gwcluster_site_42_1 for the two clusters.

Per Site Clusters

Site	Version	Hitless Failover	Max Gateway Failover
DFW-Branch	10.2.0.0_78933	POSSIBLE	1

Site	Version	MAC Address	IP Address
DFW-Branch	10.2.0.0_78933	00:0b:86:b5:37:c7	10.1.130.100
DFW-Branch	10.2.0.0_78933	00:0b:86:b5:38:a7	10.1.130.101

LAS-Branch	10.2.0.0_78635	POSSIBLE	2



MOD 9- 14

The output in the slide also shows how hitless failover is possible, assuming only a single member would fail (since there are only two gateways in the cluster).

Manual Clusters

Clusters

Redundancy

Cluster mode:

Automatic: ☐

One to one redundancy:

Clusters

CLUSTER NAME	#OF GATEWAYS	SITE NAME
+		

Gateways in P9T14-GW-Group-2 Cluster

GATEWAY
▼ GW-7 GW-6 GW-5 GW-4

1. Disable automated option

3. Edit options

2. Add gateways

Manual cluster configuration: ☒

Cluster name:

Dynamic authorization (CoA):

☐

VPN termination:

☐

MOD 9- 15

The manual mode of creating a cluster allows you to configure features such as dynamic authorization (CoA), VPN termination, multicast VLANs, and heartbeat thresholds. Enabling manual mode requires removing any WLANs mapped to existing auto clusters.

Note: If you enable a manual cluster configuration, newly-added gateways do not automatically join the cluster anymore: you must manually add these devices.



Cluster Monitoring

Now that you understand how to configure the different cluster modes, let's talk about monitoring your clusters.


Confidential - For Tra

Cluster Monitoring

MOD 9- 17

To monitor a gateway cluster, access the Devices page and navigate to Gateways > Clusters. Notice in the cluster table rows, you have lots of information about the cluster, including the number of AP tunnels and clients, the gateway version, and hitless failover and max gateway failure. The hitless failover means that the cluster is synchronizing users and sessions; that way, if one controller goes offline (for a software update, for example) the user might, but shouldn't, be impacted. The max gateway failures indicate the number of gateways that could fail without impacting the network. Once you expand the cluster, you can see detailed information about each gateway in the cluster.

Cluster Overview



MOD 9- 18

When you select a gateway cluster from the list, Central will open the individual cluster summary page. From here, you can see more information about the specific cluster. The cluster client capacity shows the number of clients the cluster supports. When the VLAN mismatch is defined as yes, the gateways in the cluster have a VLAN mismatch configuration. The client capacity graphic shows each gateway's client load in the cluster. Hover your mouse over a time block to get the controller's load information in the selected time period.

Cluster Gateway Details

Client count

Cluster role

GATEWAYS (3)						
Gateway Name	IP Address	Status	Client Capacity (Active Standby)	Model	Role	Version
GW-1	10.1.140.100	Up	3 (1 2)	A9004	Member	10.1.0.0_76631
GW-2	10.1.140.101	Up	1 (1 0)	A9004	Member	10.1.0.0_76631
GW-3	10.1.140.102	Up	4 (2 2)	A9004	Leader	10.1.0.0_76631

GATEWAYS GW-1 ▼				
GATEWAY PEER DETAIL (3)				
Type	IP Address	Status	Role	VLAN Mismatch
SELF	10.1.140.100	-	Member	-
PEER	10.1.140.101	Connected	Member	1
PEER	10.1.140.102	Connected	Leader	1



MOD 9- 19

You will find detailed information for each gateway within the cluster from the Gateway tab, including the number of active and standby clients on each controller, the gateway model, the gateway role, and the software (SW) version.

Cluster Tunnel Monitoring

Summary Gateways **Tunnels**

GATEWAY CLUSTER DETAILS

TUNNEL DOWN - SUMMARY

AP Name	IP Address	Last Connected	Last Key Recd By AP	Last Key Recd By Gateway	Reason
 No data to display right now					

Down tunnels

Up tunnels

GATEWAYS | GW-1 ▼

TUNNEL DETAILS

AP Name	IP Address	SSID	Status	Uptime	Last Key Recd By Gateway
microbranch-ap1	10.1.145.136	P9T14-M-Tunnel	Up	5 Hrs 1 Min 44 Secs	5 Hrs: 1 Min: 44 Secs ago
p9t14-ap1	10.1.145.129	P9T14-Bridge,P9T14-Mixed,P9T14-Tunnel	Up	5 Hrs 42 Secs	5 Hrs: 44 Secs ago
p9t14-ap2	10.1.145.133	P9T14-Bridge,P9T14-Mixed,P9T14-Tunnel	Up	5 Hrs 51 Mins 27 Secs	5 Hrs: 51 Mins: 31 Secs ago
p9t14-ap3	10.1.145.131	P9T14-Bridge,P9T14-Mixed,P9T14-Tunnel	Up	5 Hrs 51 Mins 27 Secs	5 Hrs: 51 Mins: 30 Secs ago

MOD 9- 20

From the Tunnels tab, you will see all the tunnels created between the APs and gateways in a cluster. In the first table, you will find a list of tunnels that are down. In the second table, you will see all the tunnels that are up.


Secondary Gateway



Let's talk about how to set up a secondary cluster to enhance resiliency in your network.

Confidential - For Tra

Secondary Cluster Redundancy



You can enhance network resiliency by adding a secondary gateway cluster. A gateway cluster itself provides a high level of protection: if any gateway fails, then the APs and clients will failover to any other gateway in the Cluster. This failover is referred to as hitless. What if all Gateways in a Cluster fail? If you set up Secondary Gateway Cluster, APs and clients could failover to that Cluster; this failover however is not hitless and some brief downtime is expected (i.e, expect a brief period of disruption).

Note: A secondary gateway cluster cannot be auto-site enabled.

Secondary Cluster Configuration

The screenshot shows the 'Secondary Cluster Configuration' wizard in the 'WLANs' section. The wizard has five steps: 1 General, 2 VLANs (current), 3 Security, 4 Access, and 5 Summary. The 'VLANs' step is highlighted with a green arrow. The configuration options are as follows:

- Traffic forwarding mode: ☐ Bridge, ☒ Tunnel, ☐ Mixed
- Primary Gateway Cluster:
- Secondary Gateway Cluster:
- Cluster Preemption: ☐

The 'Secondary Gateway Cluster' field and its dropdown menu are highlighted with an orange box. The dropdown menu shows the selected value 'Campus2: auto_gwcluster_244_0' and a downward arrow. The text 'MOD 9-23' is visible in the bottom right corner.

To use a secondary gateway cluster, navigate to the VLANs section in WLAN wizard and select a secondary gateway cluster from the drop-down list. With a best practice, the second cluster should be at the same location as the clients and APs; however, this is not required. If the secondary cluster is at a remote location, expect more latency for your traffic flows if a failover occurs.



Knowledge Check


Self-check on key learning points



It's time for a knowledge check to review what you have learned.

Question #1

Match the appropriate descriptions to each gateway cluster deployment type.



Knowledge Check ✓



MOD 9-25



Single cluster

All tunneled WLANs terminate on the same gateway cluster

VLANs and strict firewall rules provide separation between clients

Multi-cluster

Multiple tunnels to each gateway cluster

Different WLANs can terminate on different clusters

Auto group

Members of the same group are automatically clustered

The cluster name is automatically generated

Auto site

Members of the same group and same site are automatically clustered

Cluster name is automatically generated

Question #2

Manual clusters allow you to enable microbranch AP termination and CoA redundancy.

- True
- False

Q

Knowledge Check

MOD 9-26



true

Confidential - For Training Purposes Only



Lab Activity

Let's review the lab activity that is associated with this module.

Confidential - For Tra

Lab Tasks

Lab 09.01 Gateway Cluster Deployments

Task 1

- Multi-zone

Task 2

- Setup Site Based clustering using a single Site

Task 3

- Site Based clustering using a multiple Sites

Task 4

- Site based cluster with Group based backup cluster

Module 10 - Wired Port Access



MOD 9- 28


The slide provides a brief review of tasks performed in this lab. Please see your Lab guide for details of this lab.

Wired Port Access

Implementing Aruba Campus Access

Welcome to module 10: Wired Port Access. In this module you'll learn how to lock down your edge access on your AOS-CX switches. 802.1X, MAC-auth, and captive portal access control are covered.

Confidential - For Training



Objectives

- Describe common wired authentication methods and the advantages of dynamic segmentation
- Configure 802.1X authentication, dynamic segmentation, and RADIUS server groups
- Describe, configure, and assess MAC authentication
- Leverage user roles with RADIUS attributes for efficient, consistent wired security
- Describe the advantages of client-based and device-based authentication/authorization modes
- Accommodate special scenarios by combining MAC and 802.1X authentication and other features
- Implementing captive portal

After completing this module, you will be able to:

- Describe common wired authentication methods and the advantages of dynamic segmentation.
- Configure 802.1X authentication, dynamic segmentation, and RADIUS server groups.
- Describe, configure, and assess MAC authentication.
- Leverage user roles with RADIUS attributes for efficient, consistent wired security.
- Describe the advantages of client-based and device-based authentication/authorization modes.
- Accommodate special scenarios by combining MAC and 802.1X authentication and other features.
- Implement captive portal.

Wired Authentication

The first section of this module provides an introduction to wired authentication using AOS-CX switches.

Confidential - For Training


Port-Based Authentication Use Cases

Secure the edge

Control guest access

Dynamically provision ports

Enable segmentation




One important use case for port-based authentication is to secure the edge: you must prevent unauthorized users from connecting to open switch ports. This may be required for compliance to corporate security policies and government regulations. A second use case is to control guest access. You must ensure that guests are placed in the proper guest VLAN and must log in for legal and tracking purposes. And a third use case that can really help reduce administrative overhead while improving security is to dynamically provision access ports based on a user's identity or device.

Dynamic segmentation establishes least privilege access to applications and data by segmenting traffic based on identity and associated access permissions. This is a fundamental concept of both zero trust and SASE frameworks.

Note: Originally Aruba used the term tunneled node to describe the dynamic segmentation feature. Today, Aruba uses the term dynamic segmentation when referring to the original tunneled node feature.

Challenge: Tedious, Error-Prone Configuration




MOD 10- 5

The original way of implementing port access was to manually configure specific ports for a VLAN information, access control lists (ACLs), and QoS, and etcetera. Obviously, this would be a tedious task and provides less flexibility: you must ensure that everyone connects each device to the correct port. This is error prone. Maybe the orange ports are for accounting users, red ports are for IT admin, Green is for sales, Blue is for access points, purple is for printers, and so on. And it's less flexible if you have devices that sometimes move around the network and you're not sure where they'll be connecting to the access switches.

Today, a different approach is needed: you have a simple collection of minimally configured colorless ports. Upon connecting to a port, users are assigned a role, which Aruba's ClearPass Policy Manager associates with a particular policy. And the role dictates the policies: VLAN information, ACLs, QoS, and etcetera. This provides flexibility and scalability. Let's take a closer look at this approach.

Customizing Authenticated User Access




Beyond the security benefits, port-based authentication helps you to provision the network edge more quickly. You learned about the old days, when you had to configure each port with the appropriate settings for the intended user: tedious and error prone. This can be done with 802.1X, MAC-auth, and/or captive portal.

With 802.1X, the RADIUS server and the switch (authenticator) automatically determine the user identity as well as other context characteristics: who, what, how, where, when. A RADIUS server policy can use this information to automatically configure the appropriate authorization policy settings: VLAN, ACLs, QoS/CoS, rate limits, and more. The network edge becomes more flexible and adaptable.

Because the switch applies these settings to a user session automatically, as dictated by the RADIUS server, these settings are sometimes called RADIUS settings or dynamic settings. For example, a dynamic VLAN assignment is the VLAN assigned to a user session by the RADIUS server when the user authenticates to the network. A dynamic ACL is the (port-based) ACL applied to a user session by the RADIUS server when the user authenticates, and so on. If the interface is disabled or the user disconnects, the port reverts back to its original configuration and is moved to an unauthenticated state. Therefore, if a new device would connect to the port, it would have to authenticate and a new policy would be dynamically assigned to the port.

Solution: Dynamic Segmentation



MOD 10- 7

To simplify policy enforcement in a wireless and wired network, dynamic segmentation can be utilized. Dynamic segmentation allows a switch to act like an AP: just like an AP, it will tunnel traffic to a gateway, where the gateway can use its advanced firewall capabilities to control traffic. Likewise, you can use the same or different policies for users that need to make connections for both wired and wireless access.

With dynamic segmentation, a policy is pushed down to the policy enforcement firewall (PEF) feature on the gateway, and a tunnel forms between that PEF device and the user's connected switch or AP, which permits or denies access to appropriate resources.

Thus, user capabilities are centrally controlled, and VLAN connectivity, security, and QoS are dynamically assigned. Your network is much more secure, with far less initial configuration and management cycles.

Pretty cool solution, right? You'll learn more about dynamic segmentation later in this module.

Common Authentication Methods

802.1X

- Supports bi-directional authentication
- Usernames/passwords and/or certificates
- Commonly used for employees

MAC-auth

- Must manage a MAC address database
- Subject to spoofing, augment with ACLs and fingerprinting
- Commonly used for non-intelligent devices, like IoT

Captive portal

- Commonly used in guest networks
- User device uses a web browser to authenticate
- Commonly used for guests



MOD 10- 8

The authentication process validates a user or device identity in order to permit or deny network access. Typically, you only wish to permit authorized users with the proper network credentials. A MAC address, passphrase, username and password, or certificate can all serve to authenticate users, with credentials pre-configured in an external source, like an LDAP or Active Directory (AD) data store. This is like checking a passport to verify a traveler's identity. Let's briefly review the three common authentication methods:

- 802.1X
- MAC-auth
- Captive portal (commonly referred to as a web portal)

802.1X

802.1X is the most common method used to validate clients on enterprise-class networks, used in conjunction with an Extensible Authentication Protocol (EAP). There are many flavors of EAP, including PEAP, EAP-TLS, EAP-TTLS and others.

You can think of 802.1X and EAP as “sister protocols” that work as a team. EAP uses 802.1X to transport the user credentials between the user device and the NAS. 802.1X happens at Layer 2, where users are authenticated before an IP address is assigned. This is independent of the layer 3 protocol. Most NAS solutions will then forward the EAP information onward to a

centralized AAA server using the RADIUS protocol.

Note: A NAS (network access service) is a common name used in AAA deployments and refers to the authenticator in 802.1X—the networking device that interacts directly to the user device.

802.1X was covered in Module 5.

MAC-auth

With MAC authentication, the network access device (NAS) validates an endpoint MAC address against a list of valid MAC addresses. Endpoint MAC addresses are burned in ROM at the factory, but are copied to RAM during bootup; and this is the MAC address an endpoint adds to Ethernet or 802.11 Wi-Fi frames. This means it is relatively easy for bad actors to change or “spoof” their MAC address to gain unauthorized access.

Since the MAC address in Ethernet and Wi-Fi frames are not encrypted, hackers can use packet capture software to find valid MAC addresses, and then change their MAC address accordingly. Due to these security weaknesses, MAC Authentication is only used for endpoints that cannot support better authentication methods like 802.1X. This includes specialty devices, like some older medical equipment, IoT sensors, and scanners.

To improve overall security with MAC-auth, use ACLs and/or firewall policies to strictly limit what devices can access. You can also use device fingerprinting to examine device protocol information, like DHCP and HTTP payload information, and then use this information to identify additional information about the device, like the product, operating system, and other information. This information can be used to determine which policy should be applied for the device’s access.


Captive Portal

Some systems use layer 3 authentication such as captive portal web authentication or VPN software. Captive portal is commonly used in guest networks to register and restrict their access, as well as onboarding employee devices, like a new desktop, laptop, or smartphone. Captive portal is web-based and thus only requires a web browser on the user device. Given that web is based on the HTTP protocol, only minimal layer 3 access is needed to perform captive portal: DHCP, DNS, and restricted HTTP to perform the captive portal process.

Captive portal was covered previously in module 6.

Confidential - For Training Purposes Only

802.1X Overview



MOD 10- 9

802.1X forces users to authenticate as soon as the connection between the switch port and the client comes up. To restrict an unauthenticated user so that client can send only authentication messages, 802.1X divides the port into two logical ports.

Uncontrolled port

At first, the switchport is an uncontrolled port. This means no process controls whether this logical port is enabled or disabled. It is always active, but only allows 802.1X / EAP packets used for authentication. Successful authentication is required before the controlled port is enabled. With AOS-CX switches you can optionally allow LLDP or CDP traffic on an uncontrolled port: but all other user traffic is denied, by default, until the user device authenticates.

Controlled port


The controlled is typically only enabled upon successful authentication. After authentication, the controlled port can accept all types of traffic, but you can control this with locally-defined access lists or centrally defined policy (like using dynamic segmentation). Both sides of the connection control the port based on their assessment of the authentication state. In other words, no unauthorized traffic flows across the link.

This is merely a formal way to ensure that no ingress switch traffic is allowed until you authenticate (except for the 802.1X/EAP traffic, which is required to perform the

authentication). The uncontrolled (non-authenticated) port does not even allow DHCP or DNS traffic. For example, once authentication happens, then IP addressing information can be assigned.

Confidential - For Training Purposes Only

802.1X Requirements



MOD 10-10

The figure in the slide summarizes the requirements for 802.1X on clients, servers, and switches. Of course, the endpoint must have an 802.1X supplicant, which comes standard in most modern operating systems like Windows, MacOS, iOS, and others. Some organizations may prefer a third-party supplicant for additional features or personal preference. RADIUS clients or NAS devices must support 802.1X of course (all managed AOS-CX switches support 802.1X). And you must have an EAP-compliant RADIUS server:

Aruba ClearPass provides a RADIUS server, as well as other capabilities for monitoring and managing user access. This includes the powerful context features, device profiling, centralized policy management, and more.

You can alternatively use a third-party RADIUS server such as Microsoft Network Policy Server (NPS) or an open-source server such as FreeRADIUS.


Configuring 802.1X



Let's talk about how to configure AOS-CX switches to perform authentication.

Confidential - For Train

Specifying the RADIUS Server



```
radius-server host 10.100.50.50 key plaintext aruba123
[port <port-number>] [timeout <seconds>]
[auth-type {pap | chap}] [acct-port <ACCT-PORT>]
[retries <RETRY-COUNT>] [vrf <VRF-name>]
```

```
ip dns server-address <IP-ADDR> [vrf <VRF-NAME>]
```

OR

```
ip dns host <WORD> <IP-ADDR>
```

MOD 10-12

Before you enable 802.1X on a switch, you must tell the switch to whom to send RADIUS messages and, optionally, how to send them. Use the syntax shown in the figure, which refers to the RADIUS server by its IP address. You can also use the Fully Qualified Domain Name (FQDN), if preferred.

Note: If you define a FQDN name for the RADIUS server, you'll have to define a DNS server to resolve the name to an IP address or use static hostname entry.

The switch sends a RADIUS Access-Request to this IP address whenever it needs to complete any type of RADIUS authentication. By default, the switch uses the RADIUS standard UDP port 1812 for the authentication messages. After the RADIUS server has authorized a user for access, the switch can continue to communicate information about the user session in accounting packets sent to port 1813.

Note: You can also include the option `auth-port` and `acct-port` options in the `radius-server host` command to change the authentication and accounting ports. Other options are documented in the switch Security guide for AOS-CX switches.

Global key or shared secret

The `key` option in the `radius-server host` command sets an individual shared secret for communicating with the RADIUS server. The RADIUS server has the switch IP address configured on it as a RADIUS client (sometimes called NAS or, on ClearPass, a network

device). The shared secret configured for the switch on the RADIUS server must match the key configured for the RADIUS server on the switch. Otherwise, the RADIUS server will not respond to the switch's messages and vice versa.

Setting the key for the individual server is optional if you set a global key. An individual key configured for a host overrides the global key, but if you do not set an individual key, the RADIUS server uses the global key. You might choose to use the global key if you have multiple RADIUS servers that use the same shared secret. The command for setting a global key is:

```
Switch(config)# radius-server key plaintext <secret-key>
```

Defining a RADIUS host with a key will override the global key setting. The syntax to do this is shown in the slide.


Note: RADSEC (RADIUS with IPSec) is supported by AOS-CX switches. To configure RADSEC, see the Security guide for AOS-CX switches.

Confidential - For Training Purposes Only

Configure Dynamic Authorization

```
radius-server dyn-authorization enable  
radius-server dyn-authorization port <number>
```

Accept CoA messages
Change default DM port



Validate CoA config and operation

```
show radius dyn-authorization  
show radius dyn-authorization client
```



MOD 10-13

When you use Aruba ClearPass as the RADIUS server, it is recommended that you add a few settings to ensure correct operation between the AOS-CX the ClearPass RADIUS. When you use Aruba ClearPass as the RADIUS server, it is recommended that you add a few settings to ensure correct operation between the AOS-CX and the ClearPass RADIUS server.

Enable acceptance of CoA messages


Aruba ClearPass can assess endpoints on an ongoing basis and change an endpoint's authentication status or settings. To make these changes, ClearPass sends disconnect messages (DMs) and CoA messages. For the switch to accept these messages, you must enable dynamic authorization for the RADIUS globally on the switch, as shown in the figure. Without this option, some solution components might work, but others will fail.

If dyn-authorization is enabled, the switch accepts these messages on the default UDP port 3799, which is also the default port on ClearPass. You can change that setting, if necessary, in your environment, as shown in the slide. If you are using a different RADIUS server from ClearPass, you should check whether the server needs to send DMs and CoA messages.

Use the show radius dyn-authorization command to verify the configuration and operation of CoA for all servers, and restrict the view to one server with the show radius dyn-authorization client command.

Dynamic Authorization Time Settings

```
radius-server dyn-authorization enable
radius-server dyn-authorization port <number>
radius dyn-authorization client 10.100.50.50 time-window <1-65535>
```



```
ntp server 10.250.1.3 iburst
ntp vrf mgmt
ntp enable
ntp authentication
ntp authentication-key 1 sha1 aruba123 trusted
```

MOD 10-14

When an AOS-CX switch is configured to allow CoA and DM messages, it checks these messages to determine if they are current. The message should include a time stamp for the event that caused the RADIUS server to change the user status. The switch only accepts the message if the time between the time stamp and the current time on the switch falls within its RADIUS time window. The default time window is 300 seconds. You can adjust the time window as shown in the slide with the radius dyn-authorization client command, optionally specifying a VRF.

To ensure that the switch can accept the CoA and DM messages, make sure that the switch has the correct time. It is best practice to configure NTP on the switch and the RADIUS server and have them take their time from the same clock. You must also determine whether the RADIUS server always includes a timestamp with its DM and CoA messages. ClearPass does not. Unlike the older ArubaOS switches, you cannot disable the time window check. Instead, set up NTP, as shown in the slide, to ensure that your switches have the correct time. For the explanation of these commands, see the Fundamentals Guide for AOS-CX switches.

Enabling 802.1X

```
radius-server key plaintext aruba123
radius-server host 10.100.50.50 vrf mgmt
```

← Server key and IP address

1: Use RADIUS with 802.1X

2: Can use any server

3: Enable 802.1X globally

- ① `aaa authentication port-access dot1x authenticator auth-method eap-radius`
- ② `aaa authentication port-access dot1x authenticator radius server-group radius`
- ③ `aaa authentication port-access dot1x authenticator enable`

```
interface 1/1/1
  aaa authentication port-access dot1x authenticator
  enable
  aaa authentication port-access allow-lldp-bpdu
  aaa authentication port-access allow-cdp-bpdu
```

← Enable 802.1X on interface

← Allow LLDP/CDP, BPDUs on secure port




MOD 10-15

To configure 802.1X, first specify the RADIUS server IP address and key. Next, configure RADIUS with 802.1X, specify that it can use any defined server, and then enable 802.1X globally. The first server defined with the `radius-server host` command will be used, by default. Soon you'll see how to restrict the use of particular RADIUS servers by creating a group.

Last, you enable 802.1X on a per-interface basis, as shown in the slide. Now the switch will act as described earlier. It will keep the port closed to all non-EAP traffic and send EAP messages to trigger an 802.1X supplicant on the connected device to authenticate. If the user succeeds in authenticating to the RADIUS server, the switch opens the port for normal traffic. If the user fails, the port remains closed.

If you need to allow LLDP/CDP or BPDU packets on a secure port, add the optional configuration shown in the figure. This configuration might be necessary for devices like a phone, AP, or camera where, LLDP (or Cisco's CDP) must be exchanged for the device to operate correctly prior to authentication occurring on the port.

RADIUS Accounting



```
radius-server key plaintext aruba123
radius-server host 10.100.50.50 vrf mgmt
```

← Server key and IP address

```
aaa accounting port-access start-stop group {radius | <GROUP-NAME>}
aaa accounting port-access start-stop interim <minutes> {radius | <GROUP-NAME>}
```



MOD 10-16

Enable RADIUS for start-stop accounting as shown in the slide. Accounting will track when a user session begins and ends (start-stop).


For more detailed session tracking, you can configure the switch to send periodic accounting messages during the session, with interim option, as shown. The default interval is 60 minutes—in most cases you'll want to make this a smaller interval for more frequent updates to your statistics. Refer to your AOS-CX switch manual's list of network accounting attributes for a complete list of the RADIUS accounting attributes that it supports for 802.1X.

RADIUS accounting takes effect after a user successfully authenticates. The switch then creates a session for the user, defining the network access that the user is provided. If you selected start-stop, the switch sends a start accounting message to the RADIUS server when the session begins. With either stop-only or start-stop, the switch sends a message when the session ends (because the user disconnects or the session times out, for example). These messages include several standard RADIUS attributes that describe the user's identity (User-Name), location (NAS-IP-Address and NAS-Port), and bandwidth consumption (Acct-Output-Packets and Acct-Input-Packets), as well as other information.

RADIUS Server Groups


Default

- RADIUS servers used in order defined
- First server reachable is used
- Referred to as the *global* group



RADIUS server groups

- Group servers based on purpose/function
- AAA feature assigned to a particular group for processing



MOD 10-17

Companies deploy multiple RADIUS servers to provide redundancy. The servers that you add are automatically added to the global RADIUS server group, which is called radius, and is referred to as the global group. The switch uses this group, by default, for all AAA functions for which you have enabled RADIUS. Servers are used in the order they are defined.


You might want to use only a subset of servers for a particular task, instead of all the globally defined ones. You can do this by creating RADIUS groups. As shown in the slide, you might create a group of servers for end user authentication using 802.1X, and another group used for administrative access to the network devices' CLI and GUI. To use a group, a particular AAA function must reference it.

Note: AOS-CX switches support a maximum 28 groups per switch.

Configure RADIUS Server Groups


```
aaa authentication port-access dot1x authenticator radius server-group radius
```

Global group radius




```
aaa group server radius 8021X  
server 10.1.1.1 vrf mgmt  
server 10.1.1.2 vrf mgmt
```

Group 8021X



Group Admin



```
aaa authentication port-access dot1x authenticator radius server-group 8021X
```



MOD 10-18

The top example in the slide shows how to configure RADIUS with 802.1X using the default global group. The switch uses the global radius group to authenticate users who connect to ports with 802.1X enabled on them.

The bottom example shows how to configure a specific RADIUS group. First you define the group and its member servers. Then you indicate that the 802.1X process should use this group to authenticate users.

See which RADIUS servers are part of any group with the command `show radius-server`. See which RADIUS servers are part of the radius group with the command `show aaa server-group radius`.

You cannot edit the radius server group. All globally defined RADIUS servers belong to it; however, you can create specific RADIUS server groups and include only the servers you want for the RADIUS service in question, like 802.1X or MAC authentication, for example. Then reference them in the appropriate AAA command. In the example at the bottom of the slide, the command to do this for 802.1X is: `aaa authentication port-access dot1x authenticator radius server-group`.

How Switches Handle RADIUS Server Groups


Multiple servers in a group

- No load balancing
- Other servers are only for backup

Proper preparation

- Verify switch-server connectivity
- Configure RADIUS servers with NAS IP
- Add servers in desired order

Global group radius



Basic load sharing

```
hostname Switch1  
aaa group server radius 8021X  
  server 10.1.1.1 vrf mgmt  
  server 10.1.1.2 vrf mgmt
```

```
hostname Switch2  
aaa group server radius 8021X  
  server 10.1.1.2 vrf mgmt  
  server 10.1.1.1 vrf mgmt
```



MOD 10-19

When a RADIUS server group has multiple servers, the switch does not load balance requests to different servers. Instead, it sends all requests to the first server in the group that it successfully accesses. It uses the other servers as backups in case it cannot reach the first server.

For this reason, you should prepare before adding RADIUS servers to your switch and enabling 802.1X. Ensure that the switch NAS has connectivity to the RADIUS servers, and the RADIUS servers are configured to accept messages from it. You should then add the RADIUS servers in the order in which you want the switch to use them.

For rudimentary load sharing, specify different servers first on different switches, as shown in the slide. So, the default timeout is 5 seconds and the default retransmit setting is 3. After 20 seconds, the switch determines that the RADIUS server is not available. The switch then tries to reach the next RADIUS server in the list. If the switch cannot reach that RADIUS server either, after 20 seconds, it tries the next server on the list. If the switch cannot reach any of the servers, it tries the first server in the list again.

However, you can configure the switch to wait a set period after the last failed retransmit to a particular server before attempting to communicate with that server again. You might do this to prevent a switch from endlessly attempting to send RADIUS packets during a network outage. This period is called the dead time.

These are global settings that apply to all RADIUS servers and groups. You can adjust the timeout (in seconds—5 is the default) and the number of retransmits (1 is the default), and the dead time (in minutes) with these commands:

```
Switch(config)# radius-server timeout <1-60>
```

```
Switch(config)# radius-server retries <0-5> fail-through
```

By default, if a RADIUS server is reachable and the user fails authentication, the switch uses this result. However, if the user is a visitor from a different location, like a neighboring campus, and their credentials are not stored on the local campus server, but the remote campus server, the user would be denied access by default.

You can enable the fail-through feature, so that if the first RADIUS server is reachable and returns a “failed” result, the switch can proceed and try other RADIUS servers in the list or group. This feature is disabled by default, but can be enabled with this command:

```
Switch(config)# aaa authentication allow-fail-through
```

Confidential - For Training Purposes Only

Selecting RADIUS Server Groups

Use RADIUS server groups for:

- 802.1X
- Captive web port
- Authenticate Telnet, SSH, GUI

Global group radius



Server1



Server2



Server3



Server4

Add servers to other groups

```
aaa group server-group radius 8021X host 10.1.1.1
aaa group server-group radius Portal host 10.2.2.2
```



MOD 10-20

This module focuses on 802.1X authentication. However, as you will learn in other modules, AOS-CX switches can use RADIUS with other types of authentications, such as captive portal, MAC-auth, and the authentication of telnet, SSH, or web UI managers.

To use different RADIUS servers for different types of authentications, you must place those servers in different RADIUS server groups, as shown in the slide. Use the no form of the command to delete the server from the group. Each group can have up to three RADIUS servers.

A RADIUS server can be part of more than one group. If, for example, you add a server that is part of the radius server group to a new group, the server remains part of the radius group, as well. Therefore, you should define the servers that you want to use for the global group before adding the servers for the other groups. Otherwise, the servers that you intend for other groups will take up the slots in the global radius group.

Remember, to use the RADIUS server group, you must choose the type of authentication to which it applies. Here is the syntax for association 802.1X authentications to a specific RADIUS server group:

```
Switch(config)# aaa authentication port-access eap-radius server-group <group name>
```



MAC Authentication (MAC-auth)


Let's move on and discuss MAC-auth for the AOS-CX switches.

Confidential - For Training

MAC-auth Overview

Objective Control access for devices with limited capabilities

- Simple provisioning for printers, IoT devices (VLAN)
- Based on client source MAC address
- Minimal security—easily spoofed
- Second form of authentication, after 802.1X
- Captive portal with MAC-auth



MOD 10-22

MAC authentication (MAC-auth) enables access control for devices with limited capabilities, with simple provisioning for things like printers and IoT devices that may not have an extensive user interface or lack support for 802.1X. In addition to permitting a device network access based on its group, the authentication server can assign the dynamic settings, such as a VLAN ID and QoS priority, to device's session.

These devices simply connect and normal Ethernet frame, which of course includes their own MAC address as the source in those frames. In most cases, the device is probably generating a DHCP request. The switch harvests this source MAC address and sends a RADIUS Access-Request message to the RADIUS server. No client configuration required on the client side.

When you need rigorous security, you should use 802.1X, whenever possible, since it supports a username and password and/or digital certificate authentication at layer-2. Often it is trivial for a hacker to discover MAC addresses used in your network, whether by eavesdropping or accessing a legitimate device. The hacker can then easily spoof an authorized MAC address. MAC-auth might also be more difficult to manage than other forms of authentication, requiring directory or domain administrators to create accounts for each MAC address for each device, whereas 802.1X can use the user accounts and/or certificates already in the directory or domain.

MAC-auth relies on a RADIUS server. There is no need to configure MAC-addresses on each


switch: a RADIUS server or external server contains the MAC database (DB) for centralized control. You can define up to three RADIUS servers for redundancy in case access to the primary server fails. It also means the same credentials are used for authentication, regardless of which switch or switch port is the current access connection into the LAN.

The switch acts as a port-access authenticator using a RADIUS server for authentication, and uses the CHAP and PAP protocols to for sharing the information with the RADIUS server. Inbound traffic is processed by the switch alone, until authentication occurs. Some traffic from the switch to an unauthorized client is supported (for example, broadcast packets) before authentication occurs. And just like with 802.1X, you can make exceptions to LLDP and CDP messages, if required.

Confidential - For Training Purposes Only

RADIUS-Based MAC-auth

1. Switch prohibits traffic to/from unauthenticated MAC addresses
2. Switch submits unauthenticated MAC address to RADIUS server
3. If switch receives a RADIUS accept, the switch allows traffic



MOD 10-23

You can enable RADIUS-based MAC-auth on individual switch ports, to enable the process described here:

The client is initially unauthenticated. The switch does not transmit traffic on the connection. Authentication is not triggered until the client begins to send traffic on its own. This traffic is initially dropped while the switch completes the next steps.

The switch submits credentials to a RADIUS server on the client's behalf. Often these credentials consist of the client's MAC address as the username and password. However, the password can be different from the MAC address.

The RADIUS server authenticates the MAC address against a user account. Only the switch and the server are involved in the exchange of authentication messages, unlike with other port-based authentication methods. The RADIUS server uses its user database—which should have an account for the client's MAC address—and policies to determine whether to grant access. It communicates its decision to the switch in a RADIUS Accept-Access or Reject-Access message. Optionally, depending on the AAA product, the MAC addresses can be referenced in an external DB. ClearPass supports many connection options in this regard.

The switch handles the client according to the RADIUS server's decision.

- If the server accepted access, the switch allows all traffic to and from the client with that MAC address until a filtering or dynamic segmentation policy was received from the

server.

- If the server rejected access, the client remains in the unauthenticated state, and the switch handles its traffic as indicated in step 1.

The switch continues to examine incoming frames. If a new client connects (new MAC address), the switch will authenticate that client. The switch also attempts to authenticate the client for which authentication failed again. By default, the switch suppresses authentication requests for such clients over the quiet period—by default, 60 seconds. It then sends another authentication request. This process will be covered in more depth shortly.

Confidential - For Training Purposes Only

RADIUS-Based MAC-auth Options

Authentication protocol

- PAP: clear-text password
- CHAP: hashed password

```
aaa authentication port-access mac-auth  
auth-method <chap | pap>
```

Credentials

- Username = MAC address
- Password = MAC address

MAC format

- Switch format must match DB

```
aaa authentication port-access mac-auth  
addr-format <delimiter-format>
```



MOD 10-24

Let's talk about the three things you need to consider when configuring MAC-auth on AOS-CX switches.

Authentication protocol

You have two choices for the protocol that the switch uses to send the credentials to the RADIUS server:

Password Authentication Protocol (PAP)

Challenge Handshake Authentication Protocol (CHAP)

You learned about how RADIUS works with EAP in 802.1X. AOS-CX switches can use PAP or CHAP for MAC-Auth. In this case, the switch acts as the authentication client rather than the client itself, generating the PAP request messages. Note that PAP sends passwords (the MAC addresses, in this case) in clear-text, which can pose a security threat. If this is a concern, then you should implement RADSEC.

RADIUS has specific messages for CHAP authentication and most RADIUS servers support this option. You should be aware that the mechanisms of CHAP require passwords to be stored with reversible encryption, which poses a potential security hazard and is not typical for most directories and domains. This limitation would mean, for example, that the directory, domain, or RADIUS administrator who creates the accounts for the clients would have to take

special measures to store the passwords with reversible encryption. (For a Microsoft Windows domain, you can edit the group policy object (GPO) and enable reversal encryption in the computer policies for account settings.)

Configure PAP or CHAP authentication as shown in slide.

Credentials

You must decide the precise credentials that switch will send so that the RADIUS, directory, or domain administrator can set up the device/MAC accounts accordingly. By default, AOS-CX switches send the MAC address for both the username and password. You can set up a global password, which the switch uses for the password instead of the device's MAC address. The switch itself places the password in the authentication request. No configuration is required on the client. The switch continues to use the client's MAC address for the username, so every client requires a user account on the RADIUS server (or external server).

The password option is intended to add a bit of security to the solution. However, keep in mind that hackers can still spoof an authorized MAC address, connect to the switch, and gain access to the network because the switch adds the password at its end. Remember that MAC-auth provides basic security: you should use 802.1X when you require more rigorous security. You can configure a global password for all MAC addresses as shown in the slide.

MAC address format

When you are using MAC addresses for the username, password, or both, you must ensure that the switch sends the credentials with the same format used in the user accounts on the AAA or backend server. You can change the delimiter format as shown in the slide. The default is no-delimiter. The table below shows the options and the syntax for configuring a specific format for AOS-CX switches.

Table 10-1: MAC address formats

Format	AOS-CX switch parameter
aabbccddeeff	no-delimiter
aabbcc-ddeeff	single-dash
aa-bb-cc-dd-ee-ff	multi-dash

aa:bb:cc:dd:ee:ff multi-colon

AABBCCDDEEFF no-delimiter-uppercase

AABBCC-DDEEFF single-dash-uppercase

AA-BB-CC-DD-EE-FF multi-dash-uppercase

AA:BB:CC:DD:EE:FF multi-colon-uppercase

Note: Accounts in Microsoft Active Directory (AD) domains cannot use colons, so choose a format that does not include colons if you use AD.

Confidential - For Training Purposes Only

Configure RADIUS-Based MAC-auth

```
radius-server host 10.1.1.1 key plaintext aruba123
radius-server dyn-authorization enable
!
aaa authentication port-access mac-auth
  addr-format <delimiter-format>
  auth-method <chap | pap>
  radius server-group {radius | <group-name>}
  enable
!
interface 1/1/1
  aaa authentication port-access mac-auth
    enable
```

1 Specify RADIUS server IP and keys

2 MAC-auth credential format

3 Select authentication protocol

4 Specify RADIUS servers/group

5 Enable MAC-auth globally

6 Enable MAC-auth per port



MOD 10-25

Here are the steps to implement MAC-auth on AOS-CX switches:

1. Specify your RADIUS server. In this example, you are using only the global group. These commands specify the server and set the proper options for ClearPass. Remember to set the key to match the secret configured for this switch on ClearPass.
2. Configure the format for the MAC-Auth credentials so that they match the credentials configured on the RADIUS server (or its backend directory). ClearPass often uses the no-delimiter option, which is also the default option on AOS-CX switches.
3. Choose between PEAP and CHAP for the authentication protocol. Remember that CHAP requires passwords to be stored with reversible encryption, which is not the default for many directories.
4. Specify the RADIUS server group.
5. Enable MAC-Auth globally.
6. Enable MAC-Auth on the individual ports that you want to enforce this type of authentication.

You can use the following commands to verify the configuration and operation of MAC-Auth:

`show aaa authentication port-access mac-auth interface [interface name] | all port-statistics`: Display information related to MAC authentication ports on the switch.

`show aaa authentication port-access mac-auth interface all client-status`: This command display MAC authentication clients status on the interface.

`clear mac-auth statistics [interface]`: Clear the MAC authentication statistics associated with the port and all the MAC authentication state machines attached to this port.

`show tech mac-auth [local-file]`: Displays output of the following commands:

- `show aaa authentication port-access mac-auth interface all port-statistics`
- `show aaa authentication port-access mac-auth interface all client-status`
- `show aaa authentication port-access interface all client-status`
- `diag-dump mac-auth basic`

Confidential - For Training Purposes Only

Optional Configuration

```
aaa authentication port-access mac-auth
  addr-format <delimiter-format>
  password plaintext <password>
  quiet-period <1-65535>
  reauth
  reauth-period <0-65535>
!
aaa authentication port-access mac-auth authenticator radius server-group <server-group>
```

1 MAC-auth credential format

2 Global password

3 Adjust timers

4 Specify RADIUS servers/group



MOD 10-26

Here are some optional configuration steps for MAC-Auth:

Configure the format for the MAC-auth credentials so that they match the credentials configured on the RADIUS server or its backend directory. ClearPass often uses the no-delimiter option, which is also the default option on AOS-CX switches.

If you want the switch to send a password, rather than MAC address in the password field, specify the global (fixed) password.

Optionally, you can adjust other settings such as the timers for the quiet period and re-authenticating clients. (If you set the re-authentication period to 0, the switch does not enforce reauthentication.) All of these timers are set in seconds. You can refer to your switch's Access Security Guide for AOS-CX switches to learn about more MAC-auth and its related options.

Optionally specify a server group for MAC-auth.




RADIUS Attributes & Roles

You can accommodate some useful scenarios with RADIUS attributes and roles. This section introduces these concepts.

Confidential - For Training

How the RADIUS Server Sends Dynamic Settings



You define dynamic policy settings on a RADIUS server, like ClearPass. These settings apply to certain users based on their identity and context. The server communicates these settings to the NAS switch in attribute-value pairs (AVPs) within the RADIUS-Accept message. They are called AVPs because each data item has an attribute type and an associated value. The top of the slide shows three IETF AVPs:

Attribute Type 64 defines the type of medium over which this communication occurs, with a value of 6, which means an 802-type “tunnel”

Attribute Type 65 defines the tunnel type, which is set to 13. This is used to indicate that we intend to assign a VLAN ID.

Attribute Type 81 defines the actual VLAN ID (in this case, VLAN 30).

Thus, the purpose of this set of AVPs sent from the RADIUS server to the switch is to override the port’s statically configured VLAN 99 with VLAN 30.

There are standard AVPs defined by the IETF in an RFC and there are custom attributes that are defined specifically by a vendor like Aruba. These are called vendor-specific attributes (VSAs). There are also extended RADIUS attributes.

User Roles Overview

Objectives

- Enables colorless ports
 - Port settings applied when user authenticates
- 802.1X, MAC-auth, captive portal

Benefits

- Applying a policy based on a user-role
- Less reliant on port configurations

Types

LUR

- Locally defined on the switch or pushed from Central
- Supports third-party RADIUS servers
- Preferred type

DUR

- Requires ClearPass
- Roles defined on ClearPass



MOD 10- 29

The benefits of assigning roles to users and/or devices is well known within the Aruba wireless world. Aruba has adapted this feature for use on controlling port access on the AOS-CX switches. Just as with Aruba wireless, you simplify your configuration by grouping policies into a role that can be referenced by many device or user types. Additionally, when ClearPass is incorporated, that who, what, how, where, when context can be used to control user access rights.

You can configure roles locally on the switch using a local user role (LUR). When a device connects and authenticates successfully, an Aruba-User-Role name (VSA) is passed to the switch, and then the switch will apply the Aruba-User-Role to the port to which the device is connecting. Or configure downloadable user roles (DUR) on a ClearPass server. The ClearPass HPE-CPPM-Role VSA is used with HTTPS to transfer the role to the switch.

A role is a role. At a minimum a role dictates the tagged or untagged VLAN assignment, and if the traffic is locally-switched or tunneled back to a gateway (dynamic segmentation). Optionally a role can also assign a policy (ACL/QOS), reauthentication timers, and a captive portal redirection. The same syntax is used whether pre-defined on a switch (local-roles) or downloaded from ClearPass. With local user roles, you must define them on the switch before the role can be applied to a device/user. With DUR, the role configuration is downloaded to the switch.

Important: the term role is used with both the AOS gateways and the AOS-CX switches. The

two should not be confused, since a role on a gateway pertains to firewall, bandwidth, and captive portal policies, while a role on an AOS-CX switch pertains to VLAN, QoS, ACL, and other policies. In other words, roles on both devices apply policies, but the policy types differ.

Note: In the current release (10.4), Aruba doesn't support role download for a gateway (controller) or secondary role to the gateway when implementing dynamic segmentation. The gateway would have to download itself. Also, since the use of Central for centralizing configurations and policies, it is more common today to define the role information in Central and have it push that configuration to the switches. The switches would then use LUR when interacting with an AAA server.

Confidential - For Training Purposes Only

ArubaOS-CX Switch Role

```
Table1-Switch(config)# port-access role Corporate
Table1-Switch(config-pa-role)#
  associate          Associate a captive-portal-profile or policy with this
                      role.
  auth-mode          Configure authentication mode for this Role.
  cached-reauth-period  Configure cached re-authentication period in the role.
  client-inactivity   Configure client inactivity monitor mode for this Role.
  device-traffic-class  Configure device traffic class for this Role.
  gateway-zone        Configure gateway parameters for the Role.
  mtu                 Configure MTU for this Role.
  poe-priority         Configure POE priority for this Role.
  reauth-period        Configure reauth period for this Role.
  session-timeout      Configure session timeout for this Role.
  trust-mode          Configure trust mode for this Role.
  vlan                Configure VLAN mode for this Role.1
```




MOD 10-30

In the slide, you can see all settings that can be enabled on a port using user role. When creating a role, you have to assign it a name. This takes you into a sub-context where you define the role's policies.

Using ClearPass or Third-Party RADIUS

Switch applies settings based on role

- Pre-authentication role
- Post-authentication role



MOD 10- 31

You have learned how AOS-CX switches can accept RADIUS VSAs to customize access for authenticated devices. Some companies prefer to take a different approach. They might use ClearPass to apply user roles to authenticated users based on identities and other criteria. Aruba gateways, acting as RADIUS clients, use these roles to decide how to filter user traffic. Rather than create new profiles with the RADIUS VSAs for dynamic VLANs and ACLs, you might prefer to have ClearPass simply send a role to AOS-CX switches too.

Every client is associated with a role. Roles associate with a set of attributes for authenticated clients (clients with authentication configuration) and unauthenticated clients, applied to each user session. On a switch, roles are always enabled by default.

Examples of user roles are:

Employee—All access

Contractor—Limited access to resources. Each user role determines the client network privileges,

frequency of reauthentication, applicable bandwidth contracts, and other permissions.

Guest—Browse Internet

The VSA that AOS-CX switches accept from RADIUS servers like ClearPass has these 499

properties:

Type: RADIUS vendor-specific for Aruba

Name: Aruba-User-Role

ID: 25

Value: <role name>

Note: The role name is case-sensitive.

You can define the user role locally. This approach moves the control over customizing settings from RADIUS server administrators to infrastructure administrators, which some companies might prefer. The user derived role (UDR) is applied when the roles are enabled. UDR will have the same precedence order as the authentication type (802.1x, MAC authentication). UDRs are locally configured on the switch using the port-access role command. The configuration of UDRs is beyond the scope of this course. See the Security Guide for AOS-CX switches for more information on how to configure the port-access role profile locally.

Confidential - For Training Purposes Only

Client-Based Mode

Authentication / authorization tied to device MAC

2-client limit set

Phone: Static tagged VLAN 30, authenticates

PC: Static native VLAN 20, authentication fails



PC: Static native VLAN 20, authenticates



```
interface 1/1/1
aaa authentication port-access auth-mode client-mode
aaa authentication port-access client-limit 2
```



MOD 10- 32

Client-based mode is recommended for scenarios in which you want the switch port itself to authenticate and control multiple clients. The term client-based and user-based are commonly used here and are interchangeable. This is recommended when the RADIUS server might implement different dynamic settings for different clients connected to the same switch port. Often, you should also enable user-based control even when the port should only connect to a single device to prevent unauthorized devices from piggybacking on the port.

In client-based mode the port authenticates each downstream device, as identified by their unique source MAC addresses. You enable this mode when you set the client limit on the port to any value, including 1. Client mode is the default, with a default client limit of 1. See the configuration example in the slide to change the number of clients that can simultaneously connect to the switch port.

The port creates a separate session for each device (MAC address), and each session has its own authorized status. In other words, the switch permits ingress traffic with a particular source MAC address only if that MAC address's session is authorized. Conversely, it permits egress traffic based on the destination MAC address's authorization status. The limit value indicates how many devices can be authorized.

After you set the limit, the port begins tracking MAC addresses and defines the authorization status and settings for each separately. For example, in the scenario with the computer and VoIP phone shown in the slide, the switch port sends an EAP Request/Identity to each

separate MAC address detected on the port. If the VoIP phone authenticates successfully, but the computer fails, the computer traffic is blocked, but the VoIP traffic is allowed.

Confidential - For Training Purposes Only


Client-Based Mode Example

2-client limit set

PC1: Authenticates, VLAN 30,
NAS filter applied


PC2: Not authenticated

PC3: Not authenticated



PC2: Authenticates, VLAN 20

PC3: Denied, 2-client limit reached



MOD 10-33

Examine the client-based example in the slide, where the client limit is set to two. Initially, only PC1 is authenticated, so the switch drops traffic from PC2 and PC3. Then PC2 authenticates, so the switch forwards PC2 traffic. Now the limit is reached, so PC3 cannot authenticate until another user disconnects.

The example in the slide also illustrates how client-based mode allows each session to have its own dynamic attributes, such as a VLAN and ACL, for example. The dynamic VLAN and dynamic NAS filter only applies to PC1. The switch forwards any untagged, native VLAN traffic sent by the user device in the dynamic VLAN. And when it needs to forward traffic to the device, it forwards the traffic without a tag. If the port supports multiple users, the users can have different native VLAN assignments. If any users do not receive a dynamic assignment, the static native VLAN applies to their traffic.


In the example, above both PC1 and PC2 are authenticated and assigned to different VLANs. The switch recognizes which user is sending traffic based on the source MAC addresses and forwards the traffic in the correct VLAN appropriately.

Note that the switch must flood native VLAN traffic in all VLANs out of the port. This means that devices will receive broadcasts and multicasts for all untagged VLANs (native VLANs) applied to the port. Keep this in mind when you plan multi-client scenarios.

WLAN in Bridge Mode Use Case

Concern: Only authorized devices should connect to a port

Solution: Device-based mode with 802.1X authentication



MOD 10- 34

The concern with some customers is that the access switch might be in an unsecured environment. Suppose a company statically defines switch port policy settings for autonomous APs. Someone unplugs an AP from statically-defined port 1/1/1 and connects an unauthorized device to that port. The unauthorized device would have the same policy settings applied as the AP.

If this was a corporate user, they would not have the correct VLAN settings and access would probably be denied. If this was a bad actor, they might be able to gain unauthorized access to corporate resources, depending on how the port is configured.

The other issue with this scenario is that the AP is authenticating wireless clients using bridge mode, but then if the switch was in client mode, it would authenticate these clients again!

Device-based authentication mitigates these issues. If someone unplugged the AP and plugged in an authorized device in this situation, the unauthorized device would be denied access to the network unless it successfully authenticated. Likewise, since the AP brings up the switch port, it is the one being authenticated and the appropriate policies are then applied to the port. As wireless users connect, since the AP authenticates them and locally bridges their traffic. Given that the device-based mode is used on the switch port, the wireless users don't have to authenticate.


Note: Device-based mode is useful when you know that you'll only have one device

connected to a port, but you need to authenticate access on the port and want to apply policies dynamically to the port based on the authentication credentials (and possibly other information, like fingerprinting).

Confidential - For Training Purposes Only

Device-Based Mode

- No client limit
- Only *one* device needs to authenticate
- Performs port authentication



Setting typically enabled using user roles

MOD 10-35

A switch applies AVP settings based on whether the port operates in 802.1X port-based or client-based mode. The mode also determines how the port establishes authentication sessions.

With device-based/port-based control, the switch establishes a single session for all traffic sent and received on the port. If any user authenticates, the switch permits all traffic without referring to the source or destination MAC address. Untagged traffic (native VLAN), regardless of source MAC, is forwarded in the native VLAN assigned to the port (either statically or through RADIUS, as you have learned).

Often, only a single device connects to the switch port, so all traffic does originate from a single source MAC address. The example in the slide illustrates this situation, where PC1 and PC2 are connected to two different switch ports and are assigned different VLANs based on their 802.1X credentials.

Configure the port mode for device mode as shown in the slide using the `aaa authentication port-access auth-mode device-mode` command.


Note: Client mode is the default mode of the AOS-CX switch ports.

Combining MAC-Auth and 802.1X

Port configuration: 802.1X & MAC-Auth

- First 802.1X. If that times out, then RADIUS MAC-Auth
- Can control precedence

```
interface 1/1/1
  aaa authentication port-access auth-precedence [dot1x mac-auth | mac-auth dot1x]
```



MOD 10-36

A single port can support multiple authentication methods, enabling it to authenticate different types of devices and users. For example, you have devices that use 802.1X authentication, but also some devices that don't support 802.1X and therefore MAC-auth must be used. By default, 802.1X is performed first on the port. If 802.1X times out, the switch assumes that the device doesn't support it. If MAC-auth is also enabled on the port, then the switch will authenticate the device using MAC-auth. You can change the order, per-interface, as shown in the slide, though this is typically not recommended.

In this scenario, each port supports a single device, but you are not sure which type of device will connect to the port. You want any port to support any device type: those that support 802.1X and those that only support RADIUS MAC-auth. Suppose you left it at its default value, and therefore 802.1X is tried first.

When the printer connects, the switch initiates 802.1X authentication, which the printer does not support and so does not respond. 802.1X times out, and the switch tries MAC-auth, which succeeds. The policy defined on the ClearPass server ensures that printer ports are assigned to VLAN 50.

When the PC connects, the switch initiates 802.1X authentication, which the PC supports and so authentication succeeds. The ClearPass policy ensures that employees are assigned to VLAN 20.

Suppose that you change the order to try MAC-auth first. The port uses the device's MAC address to submit an access request, just as it would if MAC-auth were the only authentication method:

If RADIUS MAC-auth succeeds, the device is authorized for network access with the settings returned by the RADIUS server profile.

If RADIUS MAC-auth fails, the user still has the option to pass 802.1X authentication. If the user succeeds, the device is authorized for network access with the settings returned by the RADIUS server.

If both types of authentication fail, the device is prohibited access.

Even if RADIUS MAC-auth succeeds, the user device might initiate 802.1X. (However, the switch itself will not.) If the user then authenticates successfully with 802.1X, the settings returned at the end of the 802.1X process take precedence over any settings applied by RADIUS MAC-auth.

Confidential - For Training Purposes Only

Onboarding Method Concurrent

- ✓ Configures all methods to start concurrently for faster onboarding process
- ✓ Default priority for concurrent onboarding is 802.1x followed by mac-auth and device-profile
- ✓ When concurrent onboarding is enabled, then auth-precedence will be ignored
- ✓ Existing clients will be de-authenticated and freshly onboarded concurrently

```
switch(config)# int 1/1/1  
switch(config-if)# port-access onboarding-method concurrent enable
```



Configures all methods to start concurrently for faster onboarding process. If authentication priority is not configured when enabling concurrent onboarding, the priority will be 802.1x followed by mac-auth and device-profile.

Default priority for concurrent onboarding is 802.1x followed by mac-auth and device-profile.

When enabling concurrent onboarding on the port, existing clients will be de-authenticated and freshly onboarded concurrently.

When concurrent onboarding is enabled, then auth-precedence will be ignored.

If concurrent onboarding is configured, the client will stay in pre-auth role till it gets succeeded by one authentication method or gets failed by all the authentication methods.

When the authentication method with the highest priority fails, the profile of the next successful authentication method is applied.

If all methods fail, the reject or critical role is applied based on the 802.1X authentication failure reason and continues to reauthenticate with the 802.1X method.

Reauthentication will be triggered for all high priority methods and not just the final successful authentication method.

Some RADIUS server may block the client when it receives two requests, mac-auth and 802.1X, from the same client at the same time. This is because the RADIUS server allows only one authentication request. In such cases, concurrent onboarding is not feasible. To prevent such scenarios, configure auth-precedence with auth-priority.

Confidential - For Training Purposes Only

Multi-Domain Authentication

Allows combination of voice and data clients to be authenticated on a port

Only **ONE** voice client and a maximum of **FIVE** data clients are allowed

Any data client above the limit triggers violation



MOD 10-38

Multi-domain authentication allows a combination of voice data clients to be authenticated on a port. When the port is in multi-domain mode, only one voice device is allowed on the port while data device limit is configurable.


By default, only one data device is allowed on the multi-domain enabled port along with a voice device. If second voice device or data device greater than the configured data client limit onboards, this triggers a violation.

Note: Only one voice client and a maximum of five data clients are allowed for authentication.

Multi-Domain Authentication

Customer wants only 1 data device and 1 voice device

Impossible with client-based mode: both devices treated as data devices



```
Switch(config)# interface 1/1/1
Switch(config-if)# aaa authentication port-access auth-mode multi-domain
Switch(config-if)# aaa authentication port-access client-limit multi-domain 1
```

```
Switch(config)# port-access role Voice
Switch(config-pa-role)# device-traffic-class voice
```

MOD 10-39

With multi-domain authentication, you can enforce a policy that would allow only one voice device and multiple data devices. This is not possible in client-based mode as all devices are always treated as data clients. You can enable the multi-domain authentication mode with the `aaa authentication port-access auth-mode` command as shown in the slide.

You can configure maximum the number of data clients supported with the command `aaa authentication port-access client-limit multi-domain <number>`.

To authorize a voice device, the AAA server must be configured to send an Aruba Attribute-Value (AV) `Aruba-Device-Traffic-Class` with a value of 1 or a role to be applied should have device type (`device-traffic-class as voice`), which indicates a voice device. Without this VSA value or device type in the role, the switch treats the voice device as a data device.

Local Authorization Role Options

preauth-role

Assigned before authentication

auth-role

Assigned after authentication
No role assigned by RADIUS

reject-role

Assigned if authentication fails

critical-voice-role

Assigned if voice client auth failed
due to RADIUS server not reachable

critical-role

Assigned if RADIUS server not reachable

MOD 10-40

The figure shows four commonly used local role types you can use for AOS-CX switches when controlling port access. The following sections cover these in more depth.

pre-auth role

The pre-authentication (pre-auth) role allows a device, such as an IP Phone, to have network access before the device is authenticated. The pre-auth role can also be used to allow DHCP, DNS, and the web redirect for captive portal authentications.

The pre-auth role is triggered when a MAC-based client is connected to an Aruba switch before being authenticated by the RADIUS Server. To provide network connectivity for devices, they must be assigned a VLAN. Two new VLANs are created for pre-auth role functionality, one for voice traffic and one for data traffic. Pre-auth role VLANs can be configured on the switch individually or within a user-role. Devices that can be connected to the switch without authentication are divided into two categories:

Devices that send voice traffic.

Devices that send data traffic.

Note: Either one of pre-auth role VLAN (voice and/or data) or a pre-auth role can be configured for a port. However, both a VLAN and role cannot coexist for an interface. Initial traffic on the port is restricted only by ACLs configured for the port or for VLANs or ACLs in

the role.

To define a pre-authentication role, configure the following:

```
Switch(config)# interface <interface-ID>
```

```
Switch(config-if)# aaa authentication port-access preauth-role <role-name>
```

Impact of pre-auth role on existing features include the following:

Unauthenticated devices: Configuring pre-auth role VLAN will change the behavior of unauthenticated devices. Normally, authentication-enabled ports will not provide unauthenticated client any network access until the device is authenticated by the RADIUS Server. With pre-auth role VLAN configured, the client will be assigned to the pre-auth role VLAN until the RADIUS server authenticates the device. Unauthenticated clients will be placed into the VLAN specified in the pre-auth role. After authentication by the RADIUS server, the client will be placed into the VLAN specified by the RADIUS authentication command string or as specified in the RADIUS authentication accept string.

LLDP-Bypass: When LLDP-bypass is enabled on the switch, Aruba APs are not authenticated. Therefore, the pre-auth role VLAN is not applicable.

Bypass using device-identity: Pre-auth role VLAN is not applicable to VoIP devices because they do not need authentication. It is applicable to PCs which need authentication.

ACLs applied on an Interface: If an ACL rule is applied on an interface which is part of a pre-auth role VLAN, traffic coming through that interface will be affected. Traffic will be affected based on the rule in the ACL.

ACLs applied on a VLAN: If an ACL rule is applied on a pre-auth role VLAN, traffic entering that VLAN will be affected. Traffic will be affected based on the rule in the ACL.

Rate-limiting on an interface: If the traffic is rate-limited on an interface as part of a pre-auth role VLAN, the traffic will be impacted. The traffic will be affected based on the rule in the rate-limiting configuration command.

Authenticated or rejected clients: Clients which are either authenticated or rejected by the RADIUS server are given different VLANs. These clients are moved from pre-auth role to new VLANs based on authentication by the RADIUS Server.

MAC pinning: Clients whose MAC addresses are pinned and have undergone authentication will always be treated as authenticated. Pre-auth role VLAN is not applicable in this scenario.

Effect of RADIUS tracking on pre-auth role: If RADIUS tracking is enabled and no RADIUS server is available for authentication, the port will be changed from a pre-auth role VLAN to a critical VLAN. The time taken to move from pre-auth role VLAN to critical VLAN depends on the time it takes for RADIUS tracker to inform the subsystem.

auth-role

When implementing AAA, a AAA server will authenticate the user. If the AAA server doesn't reply with a role name, or the role name send in the AAA request doesn't exist on the switch (case-sensitive), then a default role can be assigned to the user if the following has been configured.

To define a pre-authentication role, configure the following:

```
Switch(config)# interface <interface-ID>
```

```
Switch(config-if)# aaa authentication port-access auth-role <role-name>
```

reject-role

Normally, if a user fails authentication, they will have no access beyond the switch port. Optionally, you can assign a reject (failed authentication) role to the user that would give them minimal network access, perhaps to download software to become compliant in order to connect to the network, like an 802.1X supplicant. To define a reject role, configure the following:

```
Switch(config)# interface <interface-ID>
```

```
Switch(config-if)# aaa authentication port-access reject-role <role-name>
```

critical-voice-role

Specifies the role for a voice client when the RADIUS server is unreachable for authentication during reauthentication period. This is applicable when multidomain authentication mode is enabled with the aaa authentication port-access authmode command.

critical-role

A critical role can be assigned to a user when the switch sends an AAA request(s) to an AAA server(s), but no response is received, perhaps because there is a network issue preventing communication between the two sets of devices. By default, the authentication would fail and the user would be denied access. However, you can override this behavior by defining a critical role for the interface, like this:

```
Switch(config)# interface <interface-ID>
```


```
Switch(config-if)# aaa authentication port-access critical-role <role-name>
```

If the AAA server(s) is not reachable, then the critical role is applied. This allows important devices, like VoIP phones, for example, to stay connected to the network at all times.

Once an AAA server is reachable, however, the device(s) connected to the interface must be authenticated, at which point the critical role is not used.

Confidential - For Training Purposes Only

AOS-CX Switch User Role Assignment



MOD 10-41

AOS-CX switches always apply role-based authorization on ports that implement a port access method such as 802.1X. These switches can assign unauthenticated clients to different roles based on the reason that the client is unauthenticated. If the client never attempts authentication, the switch assigns it to the pre-auth role. If the client attempts authentication, but the server is inaccessible, the switch assigns the user to the critical role. And if the user fails authentication, the reject role applies. These are all types of roles—not role names. If you wanted to use the role, you would create a role such as `limited_access` and apply it to the respective role type on a switch port. All role settings are per-port.

If the user passes authentication, the switches look for the Aruba-User-Role VSA. If this VSA exists and if the switch can find the role locally or download it, the switch applies the specified role. Otherwise, the switch applies the post-auth role. Think of the post-auth role as the default role for a WLAN—it serves as a fallback if the assigned role doesn't exist on the switch.

If the company needs, for example, to support wired access in the case of loss of connectivity with the RADIUS server, you can use the other roles, such as the critical role, to meet the company's needs. Take care to configure settings for the role type that are appropriate for unauthenticated users.



Lab Activity

Let's review the lab activity that is associated with this module.

Confidential - For Training

Lab Tasks

Lab 10.01 - Wired Access Control

Task 1

- Prepare the sw-edge2 for Access Control

Task 2

- Enable MAC Authentication

Task 3

- User Role with Device Based authentication

Module 10 – Dynamic Segmentation



MOD 10-43

The slide provides a brief review of tasks performed in this lab. Please see your Lab guide for details of this lab.




Dynamic Segmentation

Dynamic segmentation, formally known as tunneled node, is a feature that has existed with Aruba switches for many, many years—long before any other vendor even considered a need for a feature like this. Use dynamic segmentation to help automate and streamline switchport configuration for endpoints.

Confidential - For Training

Dynamic Segmentation Overview

- ✓ Consistent wireless and wired network security
- ✓ Centralized role-based policy enforcement
- ✓ Access to Aruba gateway security features
- ✓ Redundant gateway and cluster support



MOD 10-45

The Aruba dynamic segmentation solution enables you to tunnel traffic on a user or device basis, sending the traffic to an Aruba gateway. Technically, dynamic segmentation is a process of dynamically applying policies (roles) to a switch port based on a user's credentials (and possibly fingerprinting), where one of those policies is local switching or tunneling of the user's traffic to a gateway.

The benefits of dynamic segmentation are:

Consistency between wired and wireless network security. The Aruba gateway enforces firewall policies on received traffic - either coming from an AP or a switch.

Centralized role-based policy enforcement. All firewall policies are setup in the Aruba gateway. This dramatically simplifies network management by eliminating complex and static network configuration on each individual switch.


Access to Aruba gateway security features. All traffic can be sent to the gateway for consistent firewalling, fingerprinting, and packet inspection.

Redundant gateway and cluster support. Tunneling traffic to a unique Aruba gateway creates a single point of failure. The dynamic segmentation solution is supported with redundant controllers (in standalone deployment) or with multiple gateways under a cluster. This cluster solution also helps to increase network performance.

Use Cases

Wired/wireless traffic tunneled to a gateway

- ✓ Consistent user experience
- ✓ Centralized, role-based enforcement



Security features applied to wired and wireless traffic

- Stateful firewall
- Deep packet inspection
- Device fingerprinting



MOD 10-46

Aruba ClearPass and port-based authentication implemented on switches provide one method for unifying access policies. Both the gateways and AOS-CX switches can enforce 802.1X authentication to users and apply the policies returned by ClearPass. However, ClearPass cannot simply send the same settings to both the gateways and the switches. For example, ClearPass often simply sends a user role assigned to a gateway, and the gateway has firewall rules associated with that role. On the other hand, ClearPass would need to send a policy with dynamic ACLs to apply similar controls to an AOS-CX switch. Obviously, the gateways can perform application and stateful inspection, which is something the AOS-CX switches cannot.

Dynamic segmentation provides tighter and simpler unification of wired and wireless access. Switches simply tunnel traffic to the gateway, which handles the traffic much as it would handle traffic tunneled from an AP. An employee can connect via their desktop (Ethernet connection), log in, be assigned a role, and receive customized access. Later the employee can connect wirelessly using a laptop in a meeting room, log in, be assigned the same role, and receive the same access.

Do keep in mind that tunneling the wired devices' traffic to the gateway, of course, increases the demands on that device. Architects will need to plan the Aruba gateway solution's capacity and licensing accordingly. A list of general guidelines is provided later in this module.


The bottom line is this: Use the dynamic segmentation feature whenever you want to apply

Aruba gateway-based security/control mechanisms to both wireless and wired traffic. You get unified access control, where the same policies regardless of whether they have a wired or wireless connection.

Note: Dynamic segmentation is supported on both Aruba switching products: AOS-CX and AOS switches.

Confidential - For Training Purposes Only

Operation Overview



ClearPass profiles each device, deems appropriate endpoints as acceptable, and sends accept messages to the gateway. The messages include tunnel information, the assigned VLAN, and a secondary role. In the slide, a user was profiled and assigned the UBT (user-based tunneling) role, which ClearPass sent to the switch; and therefore, their traffic is securely tunneled to the gateway. The gateway applies centralized policy enforcement according to role Finance to limit finance traffic as appropriate. Notice that there are actually two roles used—one for the switch and one for the gateway—to control the user device's traffic.

RADIUS can also indicate if the device traffic should be locally forwarded by the switch, directly toward its ultimate destination, without tunneling to the gateway. You sacrifice some of that centralized PEF enforcement, but you accommodate devices that are very sensitive to delay and jitter, like voice or video traffic. You just don't get all the benefits of dynamic segmentation solution, but you make that sacrifice for certain traffic types based on their needs.

Local switching is typically used when delay-sensitive traffic is involved between access-layer devices, like voice or video communications (VoIP phones, for example) or a third-party firewall already exists in the network and the company wants to continue using the policy function of that firewall instead of the gateway's firewall.

However, local switching lacks the centralized policy and enhanced security of UBT. Traffic


filtering is limited to ACLs, with no stateful or application inspection features. In many cases, companies use only VLANs for security at the access layer, where devices in a VLAN can freely interact with each other. Their traffic is only processed by the firewall if it leaves the subnet. Allowing this type of promiscuous behavior can lead to many security issues: user devices can attack other user devices in the same subnet.

Confidential - For Training Purposes Only

PAPI Communication

Control plane

- Used for AOS-CX switch-to-Aruba gateway communication
- Unlike APs, PAPI connection not secured by IPsec




MOD 10-48

The switch uses two protocols to connect to an Aruba gateway. PAPI is the protocol that runs the control plane and GRE is used to encapsulate the user's traffic (data plane). PAPI runs on UDP 8211 and this protocol is the same used by AP-to-Aruba Gateway communications. However, where APs use IPsec to protect the PAPI connection between the AP and Aruba Gateway, AOS-CX switches do not currently support IPsec.

Enhanced PAPI Security

Enhanced security: drop PAPI messages without valid hash
Prevent rogue switches from connecting to gateway



Enable on both gateway and switches and match the key



MOD 10-49

You can optionally implement an MD5 HMAC function to protect PAPI between the AOS-CX switches and the gateways. If PAPI is not protected, this communication creates a potential security risk. Rogue devices could start sending PAPI messages to the APs and switches and disrupt functionality. They could also try to insert spoofed PAPI messages to the gateway to disrupt communications or obtain information about the network. PAPI enhanced security minimizes this risk by using a MD5 and user-configurable key to hash the message. If a rogue device tries to spoof or tamper with a message, the hash will not check out because the rogue does not know the correct key to create the hash.


If administrators have enabled PAPI enhanced security on the gateway, you must set a matching PAPI key on the switch. The key must be between 10 and 64 characters. The switch then no longer accepts PAPI messages unless they have a valid hash created with this key. The switch will also be able to send its PAPI messages with the hashes created by the configured key. Enhanced security mode is disabled, by default, on Aruba gateways.

UBT Zones

```
port-access role Employee  
gateway-zone zone Building1
```

```
port-access role Contractor  
gateway-zone zone Building2
```

```
ubt zone Building1  
  primary-controller ip 10.1.10.100  
ubt zone Building2  
  primary-controller ip 10.1.20.100
```




UBT clients terminate on different gateways

MOD 10- 50

UBT zones allow users to be tunneled to different gateways or gateway clusters. This is similar to what was discussed in module 9, where a dedicated gateway cluster is used for different type of users (clients). This obviously enhances security and provide better user isolation.

UBT Reserved VLAN Mode: Default

Gateway applies firewall policies, fingerprinting, deep inspection



With UBT there is one GRE tunnel per port established

MOD 10- 51

The gateway divides the arriving tunneled traffic into one tunnel per tunneled-node port. It decapsulates the traffic and applies various controls. The switch communicates the gateway role to the gateway which applies that role to the user traffic. Based on that role, the gateway controls traffic with firewall policies and other policy actions. Finally, the gateway forwards the packet towards its destination. The gateway can apply different firewall rules to different devices in the same VLAN, again, based on dynamic settings applied during authentication.

Keep in mind that the switch is simply using the untagged VLAN on the tunneled-node port. It is not switching the tunneled-nodes' traffic in this VLAN, and therefore you don't need to tag it on the switches' uplinks as you normally would. Instead, the effective VLAN exists on the gateway and at the distribution or core switch. In keeping with this design, the switch is restricted from implementing many normal features—such as an IP address—on any untagged VLAN on a tunneled-node port. Aruba recommends reserving the tunneled-node VLANs for tunneled-node ports only. Client traffic in reserved VLAN mode must arrive untagged!

You should also understand that the gateway can override the default VLAN assignment after it authenticates the device. For example, it might apply 802.1X authentication to tunneled-nodes and receive a dynamic VLAN assignment from Aruba ClearPass. The gateway would then forward the tunneled-node device traffic in this VLAN, rather than the ubt-client-vlan VLAN configured on the switch. If you don't override the user VLAN on the gateway, then the traffic will be forwarded upstream with ubt-client-vlan configured on the switch. The gateway


replicates the broadcast or multicast packet for each user in the user VLAN.

Note that the switch will establish one GRE tunnel per port, even if multiple clients connect on the same port. The maximum number of clients per port in reserved mode is 256. The reserved VLAN is also known as UBT 2.0 and is supported from AOS-CX 10.04.0020 and later versions.

Confidential - For Training Purposes Only

UBT Extend VLAN Mode

- User VLAN tag is carried inside GRE tunnel
- Ingress user traffic can be tagged or untagged
- User VLANs **MUST** exist on the switch and the gateway!



MOD 10- 52

Extend VLAN mode requires user VLANs to exist on the switch on the gateway. In this mode traffic from the client may arrive untagged or tagged. Notice that in this mode tunneled client traffic will be 802.1Q VLAN tagged. In this mode, the switch creates a multicast tunnel between switch and the device designated gateway (DDG) that is used to maintain multicast and broadcast traffic sent from gateway to switch. The switch forwards the multicast and broadcast traffic to the UBT clients that are part of the same VLAN. The unicast, broadcast, or multicast traffic from switch to controller is always sent through UDG tunnel. This mode is also known as UBT 1.0 and is supported from AOS-CX 10.06 and later versions.

Role-Based Auto VLANs

Automates VLAN creation with user roles for:

- VLAN extended mode UBT
- Locally switched user roles
- VSA assigned roles

Simplifies switch configuration and deployment

Enable with only one command

```
Switch(config)# port-access auto-vlan
```



MOD 10- 53

Auto-VLAN enables automatic VLAN creation for the port access clients if the VLAN is not configured statically on the switch. In the case of user-based tunneling (UBT), by default, administrators are required to create VLANs explicitly to onboard the clients for every user. With this feature enabled, clients can be onboarded with the VLAN derived from local user role, downloadable user role, and RADIUS VSA. Use the command `port-access auto-vlan` to enable the VLAN creation automatically for the user-based tunneling clients. This feature is supported only on UBT 1.0.

In the case of non-UBT clients, access VLANs must be manually configured on uplink aggregation or core switches through the gateway on the respective ports. To make a VLAN as part of uplink, execute `trunk-dynamic-vlan-include globally`. This indicates whether dynamically learned VLANs from MVRP and port-access should be included or excluded on ports configured with `vlan trunk allowed all only`. By default, dynamic VLANs are not included in the trunk allowed list.

Tunneled Traffic Behavior: Multiple clients


Switch communicates gateway roles

```
ubt-client-vlan 400
port-access role Employee
gateway-zone zone zone1 gateway-role Corporate
port-access role Voice
gateway-zone zone zone1 gateway-role VoIP
```

Gateway applies policy and VLAN

```
User-role Corporate
Policy 1
Policy 2
vlan 4000
```

```
User-role VoIP
Policy 1
Policy 2
vlan 3000
```



When you configure dynamic segmentation, AOS-CX switches allow you to set a single UBT-client reserved VLAN. The Aruba gateway places all wired client traffic on this VLAN. In real scenarios this is not the expected behavior. To solve this problem, the AOS-CX switch sends to the Aruba gateway the user-role that the client has been assigned. The Aruba gateway can then use this information to assign a new VLAN, which overrides the 802.1Q tag included in the inner frame header.

As an example, the administrator wants to place the traffic from a wired PC which has MAC1 in VLAN 4000 and the traffic from the Phone MAC2 in VLAN 3000. The switch sends to the Aruba gateway the corresponding user role for each MAC. The Aruba gateway then must be configured with the proper VLAN for each user role: in the case of the VoIP role, VLAN 3000 and for the Corporate role VLAN 4000.

Wired Captive Portal

Let's talk about the third method of port access authentication that AOS-CX switches support: captive portal.

Confidential - For Training

Guest Access Considerations



Guest options

- Limited access without passing through a portal
- Authenticate in a web portal
- Accept an agreement in a web portal



Switch options

- Leverage web or portal-based authentication
- Place users in guest VLAN



MOD 10- 56


Different companies have different expectations and needs for their guest access solution. Some companies might allow guests ^limited access without any authentication. However, most companies want guests to log in through a Web portal of some kind. Even if they do not care whether the guests authenticate: they just want the guests to accept an agreement for legal purposes.

The company's expectations will dictate whether you need to configure AOS-CX switches at the access layer to enforce or participate in a web-based authentication solution, or whether you can simply place guests in a guest VLAN. You will focus on web-based authentication—also called captive portal authentication—solutions for most of this module. At the end, however, you will learn how you can enforce authentication for employees, but allow limited access to guests without forcing the guests to log in.

Advantages and Concerns

No 802.1X supplicant
No special config required
Custom web page / UAP

Guest access
Staff access for small organizations



Endpoints must have a UI



MOD 10- 57

Web-based authentication (Web-auth) eases the user login experience—no special device requirements are necessary and the user needs no experience since they’ve probably logged into sites on the Internet hundreds or thousands of times! Users simply open a web browser and are directed to a captive portal login page. They might enter credentials or simply accept an acceptable use policy (AUP) agreement to log into the network.

The simplicity and lower security make this perfect for guest access. However, some small organizations may use web-based authentication for employees or any organization might use this for onboarding Ethernet-based devices.

Web-auth advantages

Web-auth also lets companies control what users see before they receive network access. The company can present a customizable webpage that shows the guest information about the company providing the complimentary access. The company can also force the user to consent to terms and conditions, which can be a legal requirement in many regions.

Environments with external users who may benefit from Web-auth include:

Hospitals

Universities

Libraries, hotels, and other businesses that provide courtesy Ethernet access


Web-auth disadvantages

Web-auth does not provide as rigorous security as 802.1X. Users have a degree of network access before they authenticate. The credentials and authentication protocol might not be as secure as with 802.1X; however, you can improve the security by using HTTPS rather than HTTP.

Because web-auth requires interaction with the user, you cannot use it to authenticate devices without a user interface (the user device minimally needs a supported web browser). You should also note that the local web-Auth features are quite limited. However, the switches can integrate with Aruba ClearPass, which is a sophisticated guest and BYOD solution.

Confidential - For Training Purposes Only

AOS-CX Web-Auth Solutions



MOD 10- 58

In the past, web-based authentication was as simple as a login page and perhaps a backend authentication server. However, these solutions have greatly evolved over the years. A solution such as ClearPass provides a web portal as part of its complete user access management solution. Such a solution helps users log in more easily with features such as self-registration, for example. It also helps companies track and monitor users, and it controls their access in sophisticated ways.

AOS-CX switches currently support one very capable web-based authentication method: captive portal. You get a full guest and BYOD access management solution, where AOS-CX switches can use captive portal authentication to integrate with Aruba ClearPass. Aruba ClearPass provides the RADIUS and portal server functionality; the ClearPass Policy Manager and Guest components are required for this solution. ClearPass OnGuard is required if companies want endpoint integrity features. Note that AOS-CX switches can also integrate with other third-party AAA guest solutions.


Important: Currently AOS-CX switches do not support a local web-auth, where the switches provide the web server and authentication component locally on the switch: you need to integrate the AOS-CX switches with a backend guest solution, like ClearPass.

Note: AOS-CX switches can also integrate with HPE IMC's BYOD solution using the BYOD redirect feature. Or it can use MAC-auth with Failure Redirect (MAFR) to redirect users who fail MAC authentication to register with a server that then adds their MAC address to the

authorized list. These features are intended for similar use cases as the captive portal feature, and the use of each feature is mutually exclusive. BYOD redirect and MAFR are not covered in this training.

Confidential - For Training Purposes Only

Captive Portal User Experience



MOD 10-59

The captive portal authentication feature requires close teamwork between the AOS-CX switch and Aruba ClearPass. First let's look at the user experience.


A user connects a new device an AOS-CX switch port, or associates with an Aruba AP. The user powers on the device and opens their web browser. The user is redirected to a portal page, with options for logging in as an employee or a guest. At this point, the user device is locked down: they cannot browse to other pages or receive access to other network services.

Employees must enter valid employee credentials, perhaps stored in on Microsoft Active Directory. Guests use guest credentials to sign in, or they might be allowed to create their own account and log in with that. The options are controlled by the ClearPass Guest solution, which you can explore during the lab.

After logging in or registering for an account, the user is permitted network access, customized based on their identity and possibly other factors such as the health of their device, location, or type of access (switch versus AP). For example, an employee might receive different access rights than a guest. And a device that complies with security policies will receive more access than a potentially insecure device. If the user connects with the same device later or in a different location, the device is provided the same access rights transparently.

AOS-CX Captive Portal Process

- 1 Client authenticates with MAC-auth or 802.1X
- 2 HTTP/HTTPS request is redirected to captive portal server, username/password auth
- 3 Successful CP server authentication, then re-auth on the switch



- 4 Client gets internet access
- 5 RADIUS VSA or user role applies redirect parameters per client



MOD 10-60

The captive portal solution provides user access based on HTTP/HTTPS redirection. First, the client performs MAC or 802.1X authentication; for guests it would definitely be MAC-auth. This information is sent, via RADIUS, to the AAA server, like ClearPass. Obviously, if this is the first time the user connects, the MAC address won't exist on the server and the server will have to add it. Then the HTTP/HTTPS connection request is redirected to the captive portal (CP) server, which typically performs username and password authentication. After this succeeds, the AAA server sends a RADIUS CoA to the switch to change the access policy for the user, like granting internet access.


This feature is supported only for the clients who are authenticated through a RADIUS server. This feature is client based. The redirect parameters per client can be applied with user roles (local or downloadable) or RADIUS VSAs.

AOS-CX Captive Portal Options

Role: switch config

- Captive portal profile
- Profile URL and hash key
- User-role, associate the profile

RADIUS VSA: switch config



MOD 10-61



If you are implementing LUR, the configuration steps on the switch includes:

Configure the captive portal profile.

Configure the URL and hash key (optional) in the captive portal profile.


Configure user-role and associate the earlier configured captive portal profile.

When using RADIUS VSA, the Aruba-User-Role attribute is sent from RADIUS server in the authentication response packet. This role must be pre-configured on the switch with captive portal profile name. If the role is not in the switch, authentication fails. The following is at least one of the Aruba VSA's in the authentication response for a client from RADIUS server:

Aruba-Captive-Portal-URL: This will have the captive portal profile URL information. A default role will be considered on the switch and gets applied to the client when this VSA is sent.

- Length: 1024 characters
- Type: string
- Aruba vendor ID: 14823
- Aruba attribute type: 43

Authentication Enforced on the AOS-CX Switch




MOD 10-62

You will now look at how the captive portal solution delivers this experience. First, you must implement RADIUS-based MAC-auth on any AOS-CX switch port to which users might connect their devices. The switch can then submit access requests for all devices to ClearPass so that ClearPass can decide how to control the devices' access.

You can optionally add 802.1X in port-based mode if you want registered devices to eventually authenticate with 802.1X. In this case, ClearPass Onboard would deploy the correct configuration settings and certificates for supporting 802.1X to devices connected to the captive portal. Companies usually use this option for employee BYOD devices, which require more security than guest devices. This module focuses more on the guest access option, though.

High-level Authentication Process



MOD 10-63

The slide shows the high-level connection process for an unregistered guest device, organized in three phases:


PHASE ONE: As soon as the device sends any traffic, the switch sends a MAC-auth Access-Request with the device's MAC address. Because the device is unknown, ClearPass sends an Access-Accept with various dynamic settings to limit the device's access and redirect the device to the portal page on ClearPass Guest.

PHASE TWO: The guest uses this page to login, or to register and then login. At the end of this process, ClearPass Guest logs the user into ClearPass Policy Manager (CPPM). Policy Manager sends a RADIUS CoA to the switch.

PHASE THREE: Now when the device sends any traffic again, the switch sends a new MAC-auth Access-Request to CPPM. Policy Manager now knows the guest device MAC address and sends an Access-Accept, with settings for the appropriate guest access.

Let's explore each of these phases in more detail.

Phase One Details: Redirection to Captive Portal



To improve your understanding, let's examine the captive portal redirect process in more detail. Reference the figure while reviewing the process below. In the figure, any reference to HTTP also includes HTTPS.

First an unregistered guest device sends any traffic, typically a DHCP discover message.

The switch port enforces RADIUS MAC-auth, so it sends an Access-Request to ClearPass Policy Manager (CPPM), which must be defined as its RADIUS server.

CPPM does not yet know the guest device MAC address. Typically, when a device cannot authenticate, the RADIUS server sends a RADIUS Access-Reject message, and the switch port denies access. But for a captive portal solution, the ClearPass RADIUS server behaves differently, using a special Allow All MAC-Auth service. When a device fails authentication, this service matches the device to a deny any profile. The profile defines various RADIUS attributes and VSAs to apply to this device session.

- ClearPass sends a RADIUS Access-Accept to the switch, with attributes to apply to the session. These attributes include:
- A captive portal URL, which indicates the URL of the ClearPass portal in which users will log in or register. The URL must use standard format such as: <https://clearpass.example.com>. If you plan to redirect clients to an HTTPS URL, you must ensure that ClearPass has a certificate that is signed by a well-known trusted certificate authority.

- A dynamic ACL with IP rules that permit only DHCP/DNS/ICMP traffic, and HTTP/HTTPS traffic to ClearPass. It also permits HTTP/ HTTPS to the redirect URL.
- A rate limit to reduce the risk of untrusted users or devices launching a DoS attack (this VSA is not required, but it is recommended).
 - Note that you will typically use the Aruba role feature to assign these policies

Although the device is blocked from most network access, it can send DHCP messages and receive an IP address.


And when the user opens a browser and attempts to browse to a site, the device will use DNS to resolve hostnames.

The device then sends an HTTP or HTTPS request to access some website, the default web page of the browser or a URL entered as desired by the user.

The AOS-CX switch intercepts that request and sends a response, posing as that address and redirecting the browser to the captive portal URL. The switch appends some information to the URL in a query string. This information includes the device's MAC address and IP address, a timestamp, and the originally requested URL. Hackers could try to obtain inappropriate access by altering the ACL. To prevent this from occurring, the ClearPass portal can require a hash appended to any requests, created by hashing the URL with a secret key. It only responds to requests that have the correct hash and drops other requests. If this security feature is enabled on ClearPass, you must configure a key on your AOS-CX switches that matches the key on ClearPass.

The endpoint is thus redirected to the captive portal (CP) URL, and the process continues on the next slide.

Phase Two/Three Details: Login Process



MOD 10-65

The unregistered device now accesses the ClearPass guest portal page. As you see, traffic flows between the client and the ClearPass Guest portal without any proxy help by the switch.

The guest might be able to register for an account based on how the portal page is set up. The ClearPass portal's appearance and functionality is highly customizable. You can configure the portal such that guests log in with an account that has been provided to them, or perhaps you want them to self-register for an account. The portal might grant all self-registering users immediate access, force them to activate the account through a text message, or force them to receive sponsor approval. You can even link the account authentication to many different social media and business site logins, like Facebook, for example. ClearPass transparently adds a newly registered account to the guest repository.

When the user clicks a login button, ClearPass Guest transparently authenticates the guest using ClearPass Policy Manager. Because the wired solution uses MAC-auth to authenticate registered devices, ClearPass Guest also makes sure that the guest device MAC address (submitted in the URL) becomes authorized. The portal page begins a login count down—recommended at about 20 seconds for a solution that uses AOS-CX switches.

How the MAC address is authorized depends on the solution. ClearPass might add the guest device to special guest endpoint repository and pair the device with the guest account. Or ClearPass Guest might tell ClearPass Policy Manager to change the guest device, automatically discovered in ClearPass's global Endpoint Repository, to a known device. 1647

any case, ClearPass Policy Manager now recognizes the device MAC address as a known address associated with the known guest.

The successful guest login also triggers ClearPass to send a bounce port bounce RADIUS CoA message to the switch. This message causes the switch to de-authenticate the device and hold the port closed for a certain time period: 12 seconds is recommended. Note that ClearPass can be configured to use either a Disconnect Message (DM) or CoA for this purpose, but you should tell ClearPass administrators to use a port bounce CoA for AOS-CX switches.

It is important that the switch is set up to accept this message. You must enable dynamic authorization when you specify ClearPass as the RADIUS server, and you must make sure that the port is correct. The default dynamic authorization port on both AOS-CX switches and ClearPass is 3799. (If the switch does not accept the CoA message, the client will not reauthenticate, as explained in steps 13 and 14, and the client will continue to be redirected to the portal.)

When the port opens again, the device starts to send traffic (probably a DHCP discover message).

The switch port enforces MAC-auth for the device again.

Now ClearPass finds the MAC address as a known endpoint in its endpoint repository database and therefore authorizes the device for access, assigning it to a profile based on policy settings. For example, it might apply a guest profile that assigns the device to a guest VLAN using the same RADIUS VSA.

Because the login counter on the portal page is several seconds longer than the port bounce time, the authentication process is complete before the countdown ends. The guest is now connected to the network and can receive the appropriate services, based on the settings applied by ClearPass.

The next time that the guest connects the device, the MAC-auth process occurs in the background, and the guest once again receives guest-level access without having to pass through the portal. (You can set up the solution such that ClearPass forces the guest to log in through the portal again after a certain time period.)

BYOD onboard

If you are using the captive portal for BYOD, you must set up the portal a bit differently. Employees typically log in to a device provisioning page on the

portal with their existing domain or LDAP credentials. They then click a few buttons to accept a configuration that prepares their devices for 802.1X by configuring EAP settings and deploying certificates.

After this process, ClearPass sends a port bounce message to the AOS-CX switch port, forcing the port to close for several seconds. The device then automatically authenticates through 802.1X using the settings deployed during the provisioning process. ClearPass uses the 802.1X service to authenticate the user. The port is now open to the user with optional customization based on the ClearPass policies and profiles.

Because the switch port still enforces MAC-auth, the switch might also submit a MAC-auth request for the device, and the device could fail that check. However, the 802.1X authentication settings take precedence for the employees, so the device will receive the correct access and will not be redirected to the portal.

You could also choose to combine guest and BYOD solutions by adding a link to the device provisioning page on the guest login or registration page. Therefore, employees see the same login page as the guests, but by clicking a link on that page will redirect employees to an onboarding page, for example. If you choose this approach, you should authenticate people that are trying to onboard, ensuring that only employees can do this.

Confidential - For Internal Use Only

Basic Web-Auth Setup with User Roles

- 1 Configure RADIUS servers and server group
- 2 Define traffic classes
- 3 Define captive portal redirection policy
Associate classes with policy
- 4 Configure captive portal profile including URL
- 5 Associate captive portal profile to user role
- 6 Enable MAC globally and on the port
User gets role by authentication



MOD 10-66

Configuring an AOS-CX switch to participate in the solution is simple because ClearPass handles most of the functionality. There are six basic steps in setting up captive portal with user roles on AOS-CX switches. Let's explore the six steps.

Steps 1 and 2

1

Configure RADIUS servers and server group

```
radius-server host 192.168.13.31 key ciphertext *****
aaa group server radius dot1x
    server 192.168.13.31 vrf mgmt
```

2

Define traffic classes

```
class ip clearpass-web
    10 match tcp any <clearpass-ip-address> eq 80
    20 match tcp any <clearpass-ip-address> eq 443
class ip dhcp
    10 match udp any any eq 67
    20 match udp any any eq 68
class ip dns
    10 match udp any any eq 53
class ip web-traffic
    10 match tcp any eq 80
    20 match tcp any any eq 443
```



MOD 10-67

Configure RADIUS servers (ClearPass) and a server group.

- As part of the authentication setup, make sure to specify ClearPass as the RADIUS server. Configure the correct shared secret and enable dynamic authorization. Also configure NTP as you learned how to do in previous modules.
- Also make sure that you understand whether ClearPass is deployed as a high availability cluster. The cluster provides shared configuration and databases, but not a virtual IP address. Therefore, you must make sure to specify all ClearPass servers as RADIUS servers on the switch. As you learned in a previous module, AOS-CX switches do not load balance requests between RADIUS servers. Instead, they send all requests to the first defined server unless the request to that server times out and all retries expire. If you want, you can configure one RADIUS server IP address first on half of the switches, and a different server IP address first on the other half of the switches. (Note that you cannot define a hostname for the RADIUS server, so a DNS-based load balancer will not work.)

Define traffic classes.

- Traffic classes define what to match on for the policy. Traffic classes are covered in more depth in Module 13, “Traffic Optimization and QoS”.
- Note that the clearpass-web class map must specify the destination address of the ClearPass (or RADIUS) server that will handle web authentication.

Steps 3 and 4

3

Define captive portal redirection policy Associate classes with policy

```
port-access policy CLEARPASS-REDIRECT
10 class ip dns
20 class ip dhcp
30 class ip clearpass-web action cir kbps 1024 cbs
    2048 exceed drop
40 class ip web-traffic action redirect
    captive-portal
```

4

Configure captive portal profile including URL

```
aaa authentication port-access captive-portal-profile test
url http://192.168.16.11/guest/portal.php
```



MOD 10-68

Define captive portal redirection policy and associate the classes with the policy.

- Here you reference the traffic classes from Step 2. You also specify the redirect action for web traffic. Optionally you can rate limit traffic. Note that minimally you must allow DHCP and DNS, or the captive portal process will break. The redirect action in the example will redirect all web traffic.

Configure captive portal profile including the URL.

- In this example, the redirect URL is a ClearPass guest portal page. Optionally, you can define a hash key for protection with the url-hash-key plaintext command in the captive portal configuration context.

Steps 5 and 6

5

Associate captive portal profile to user role

```
port-access role EMPLOYEE
  associate captive-portal-profile test
  associate policy CLEARPASS-REDIRECT
```

6

Enable MAC globally and on the port User gets role by authentication

```
aaa authentication port-access dot1x authenticator
radius server-group dot1x enable
interface 1/1/3
  aaa authentication port-access dot1x authenticator
  enable
```



MOD 10-69

Associate the captive portal profile to the user role.

- The role ties together the captive portal profile (the redirect URL) and the previously configured policy.

Enable MAC or 802.1X both globally and on the port to let the user get the role by authentication.

- Since captive portal is implemented with MAC authentication and RADIUS CoA on the AOS-CX switches, only MAC authentication needs to be enabled on the port. Optionally, you can enable both 802.1X and MAC-Auth if you are not sure what type of authentication will be necessary on a switch port.

Captive Portal Troubleshooting

diag-dump captive-portal basic

Diagnostic info related to session, client, and port

show port-access captive-portal-profile

Config info for all captive portal profiles:
user role, URL, and the URL hash key

show tech captive-portal

See output for two commands:

- diag-dump captive-portal basic
- show port-access captive-portal-profile



MOD 10-70

The `diag-dump captive-portal basic` command displays the detailed diagnostic information related to the session, client, and port to help troubleshoot captive portal issues, including redirects.

The `show port-access captive-portal profile` command displays the configuration information for all captive portal profiles. Here's an example:

```
switch# show port-access captive-portal-profile
```

Captive Portal Profile Configuration

Name : employee

Type : local

URL : http://1.1.1.1/employee/captiveportal.php

URL Hash Key : SWNGWyMeYubHPDgVlirpEUwNK5Uf+r1vmhBIncQPw1Y=

The `show tech captive-portal` command captures the output from the two previous commands.

Additional Wired Authentication Features




Let's examine some additional wired authentication features.

Confidential - For Training

Using the AP1X Feature


- AP acts as 802.1X client and authenticates on the network
- Supports EAP-PEAP and EAP-TLS
- For TLS TPM, or use a user certificate



MOD 10-72

If your network requires all wired devices to authenticate using PEAP or EAP-TLS protocol, you must enable the 802.1X authentication type on uplink ports of an AP, so that the APs are granted access only after completing the authentication as a valid client: i.e., the AP acts as a supplicant.

AP1X Setup: EAP Method



To enable 802.1X authentication on uplink ports using PEAP or EAP-TLS protocol, complete the following steps in Aruba Central:

In the Aruba Central app, set the filter to a group containing at least one AP. The dashboard context for the group is displayed.

Under Manage, navigate to Devices > Access Points .A list of APs is displayed in the List view.

Click the Config icon. The tabs to configure the APs are displayed.

Click Show Advanced.

Click the Interfaces tab. The Interfaces page is displayed.

Click the Uplink accordion.

Expand the AP1X section.

- To set PEAP based authentication, select PEAP in the AP1X Type drop-down list.
- To set EAP-TLS based authentication:
 - Select TLS in the AP1X Type drop-down list.
 - Select User in the Certificate Type drop-down list.

Select the Validate Server check-box to validate the server credentials using server certificate. Ensure that the server certificates for validating server credentials are available in the IAP database.

Click Save Settings.

Confidential - For Training Purposes Only

AP1X Setup: PEAP Credentials

The screenshot displays the configuration interface for an AP1X device. The top navigation bar includes 'Access Points' and 'Gateways' tabs, with 'Access Points' selected. Below this, a series of sub-tabs are visible: 'WLANs', 'Access Points', 'Radios', 'Interfaces', 'Security', 'Services', 'System', 'IoT', and 'Configuration Audit'. The 'Access Points' sub-tab is highlighted. The main content area shows the configuration for 'Campus-AP2'. Within this section, there are tabs for 'System', 'WLANs', 'Radio', 'Uplink', 'Installation Type', and 'Mesh'. The 'Uplink' tab is selected. Under the 'Uplink' tab, there are three sections: 'USB Port' with a checked checkbox, 'USB Power Override' with an unchecked checkbox, and 'PEAP User'. The 'PEAP User' section contains three input fields for 'Username', 'Password', and 'Retype Password', which are highlighted with an orange box. A large diagonal watermark reading 'Confidential - For Training Purposes Only' is overlaid on the interface. In the bottom right corner, the text 'MOD 10-74' is visible.

Access Points / Campus-AP2

System WLANs Radio **Uplink** Installation Type Mesh

USB Port: ☒

USB Power Override: ☐

PEAP User

Username:

Password:

Retype Password:

MOD 10-74

If you select the PEAP protocol, ensure that the PEAP User and password are configured on the uplink port by selecting an AP group and navigating to Uplink section under the Access Points tab.

Client IP Address Tracking

Enables client IP address tracking on the switch

The default is disabled on global and VLAN levels

Enable client IP address tracking at the VLAN level

```
switch(config)# client track ip  
switch(config)# vlan 100  
switch(config-vlan-100)# client track ip
```


Tracking takes effect only if IP address tracking is enabled at system and VLAN level

MOD 10- 75

The `client track ip` command enables client IP address tracking on the switch. With this feature enabled you will see wired client IP addresses in Aruba Central. The default is disabled on global and VLAN levels. You can enable client IP address tracking globally and at the VLAN level.

MACsec Overview

- L2 hop-by-hop encryption on p2p Ethernet links
- Enables a bi-directional secure link
- Provides layer 2 security
- Mitigates DoS, MITM, eavesdropping



MOD 10- 76

Media access control security (MACsec) provides layer 2 security for wired LANs, protecting network communications against a range of attacks including: denial of service, intrusion, man-in-the-middle (MITM) attacks, and eavesdropping. These attacks exploit layer 2 vulnerabilities and often cannot be detected. MACsec appends a header and trailer to all Ethernet frames, and encrypts data payload within the frame. Receiving device checks header and tail for integrity. If the check fails, traffic is dropped. If the check is successful, the frame is encrypted.

The Media Access Control security (MACsec) protocol:

Provides a layer 2 hop-by-hop encryption on point-to-point Ethernet links, enabling a bi-directional secure link after an exchange and verification of security keys between two connected devices.

Secures switch-to-switch infrastructure using the MKA (MACsec Key Agreement) protocol and Static CAK (Connectivity Association Key).

MACsec encrypts all fields behind the source/destination MAC addresses except for MACsec SecTAG.

MACsec secures switch to switch infrastructure using the MACsec Key Agreement (MKA) protocol and Static Connectivity Association Key (CAK).

The pre-shared key (PSK) includes a connectivity association name (CKN) and a connectivity association key (CAK). The CKN and CAK are configured by the administrator and must match on both ends of the link.


Confidential - For Training Purposes Only

MACsec Encapsulation

MACsec appends a header and tail to all Ethernet frames

Encrypts data payload within the frame

Receiving device checks header and tail for integrity



MOD 10-77

The MACsec frame format includes an additional 32-byte MACsec header, which includes a well-known EtherType field (0x88E5), while allowing the Ethernet source/destination MAC addresses to be left in the clear for Ethernet frame forwarding. All remaining frame payload is encrypted. The receiving side checks information in the header and trailer to ensure integrity.

Configuration of MAC security is beyond the scope of this course.



Knowledge Check

Self-check on key learning points



It's time for a knowledge check to review what you have learned.

Confidential - For Training Purposes Only

Question #1

Which of the options below are advantages of using *colorless ports*?

- a. They eliminate the need to manually configure VLANs, ACLs, and QoS on switch ports connected to end devices.
- b. They eliminate the need to manually configure VLANs, ACLs, and QoS on switch ports connected to other switches.
- c. They provide a more automated solution.
- d. They minimize the need for 802.1X authentication for employee access.

Q

Knowledge Check

MOD 10-79

+

a, c

Confidential - For Training Purposes Only

Question #2

During an endpoint's 802.1X authentication, why might the RADIUS server reject all authentication requests from a Network Access Device (NAD) authenticator?

- a. The user entered an incorrect password.
- b. The RADIUS server was not configured with that Authenticator's IP address.
- c. The authenticator is using the wrong encryption algorithm.
- d. The authenticator was not configured with the RADIUS server's IP address.
- e. The authenticator is configured with the wrong shared secret.

a

Knowledge Check

MOD 10-80

+

b, e

Confidential - For Training Purposes Only

Question #3

RADIUS servers will not become a member of any RADIUS server group until you define a group.

- True
- False



Knowledge Check

MOD 10-81



false

Confidential - For Training Purposes Only

Question #4

You can configure your network systems to attempt to authenticate clients using 802.1X authentication and also to use MAC authentication as a backup, in case the client does not support 802.1X.

- True
- False



Knowledge Check

MOD 10-82



true

Confidential - For Training Purposes Only

Question #5

Downloadable user roles are a great option for when you need a more robust solution involving third-party RADIUS servers.

- True
- False



Knowledge Check

MOD 10-83



false

Question #6

Which of the following are motivations for integrating ClearPass services with your wired web-auth solutions?

- a. Local web-auth features are limited.
- b. The web service provided by AOS-CX switches is very basic.
- c. You can more closely track and monitor users.
- d. AOS-CX switches lack native support for a captive portal solution.
- e. You can leverage ClearPass OnGuard for endpoint integrity features.

a

Knowledge Check

MOD 10-84

+

a, c

Question #7

Which of the following are motivations for using Aruba's AP1X feature?

- a. You need the AP to act as a local authentication server for wired clients.
- b. You need the AP to act as a local authentication server for wireless clients.
- c. You need the AP to act as a local authenticator for wired clients.
- d. Your AP must authenticate before being granted access to the wired network.



Knowledge Check

MOD 10-85



d

Confidential - For Training Purposes Only



Lab Activity

Let's review the lab activity that is associated with this module.

Confidential - For Training

Lab Tasks

Lab 10.02 - Wired Access with Aruba Gateway

Task 1

- Prepare the Gateway

Task 2

- Configure Switch to Gateway Cluster connection

Module 11 -
VXLAN and
GBP



MOD 10-87


The slide provides a brief review of tasks performed in this lab. Please see your Lab guide for details of this lab.

VXLAN and GBP

Implementing Aruba Campus Access

Welcome back! Module 11 focuses on Virtual eXtensible LANs (VXLANs) and group-based policies (GBP). Here you'll learn about what VXLANs are and how to implement them. You'll also learn about how group-based policies with VXLANs provides a better way of doing policies than with IP-based policies.

Note: An introduction to VXLANs is provided in this module. To scale it, you need to complement VXLANs with EVPN, which is beyond the scope of this course.



Objectives

- Describe VXLAN characteristics and operation
- Leverage group based policies (GBP) to ease policy management

After completing this module, you will be able to:

- Describe VXLAN characteristics and operation.
- Explain VXLAN traffic flow.
- Configure VXLAN.
- Leverage group-based policies (GBP) to ease policy management.

Distributed Overlay & VXLAN Introduction

Let's first begin by exploring VXLAN solutions.

Confidential - For Training

Distributed Overlay Architecture


Uniform bridging & routing across a campus topology

End-to-end segmentation using VXLAN-GBP

Efficient layer 2 extension across layer 3 boundaries

Transported across any IP network

Stable, predictable IP based backbone (no STP)



MOD 11- 4

Distributed overlays are an evolution of the traditional campus network design. Distributed overlays are built using EVPN-VXLAN on highly available underlays and are tied to a full policy-based micro-segmentation, based on global roles, across the entire network infrastructure. Role-based policies abstract policies from the underlying network and enable flexible and simplified policy definition and enforcement. This is provisioned by fully automating the overlay to provide a single-pane-of-glass management overview.

EVPN-VXLAN benefits include the following:

Uniform bridging and routing across a diverse campus topology:

- Efficient layer 2 extension across layer 3 boundaries.
- Anycast gateways to ensure consistent first-hop routing services across the campus.

End-to-end segmentation using VXLAN group-based policies (VXLAN-GBP) provides the ability to propagate policy anywhere in the campus.

Transported across any IP network supporting jumbo frames, VXLANs must be deployed only on the edge devices of the fabric.

Note: VXLAN-GBP is defined in an IETF RFC draft: <https://datatracker.ietf.org/doc/html/draft-smith-vxlan-group-policy>.

VXLAN Characteristics

Usage

- Focused on data centers & Aruba campus
- IP-based overlay technology

Overlay

- MAC in UDP (IP) format
- Runs on any IP routed infrastructure

Scaling

- Can leverage IP ECMP for load sharing
- 24-bit VXLAN ID: 16M IDs



MOD 11- 5

The VXLAN technology is most advantageous in a data center environment, since traditional campus networks do not commonly experience the issues that VXLAN resolves. VXLAN is an IP-based layer 2 overlay technology. It encapsulates layer 2 traffic inside an IP packet for improved function and scalability. Basically, a logical layer 2 topology is built on top of a physical layer 3 topology. Therefore, a VXLAN in the logical layer 2 can easily span across a campus network.

A format referred to as MAC-in-UDP is used for this encapsulation. Once the network encapsulates this layer 2 traffic, that traffic can take advantage of any existing IP routed infrastructure inside the data center, or between data centers. Therefore, all of the redundancy, resiliency, and load-sharing capabilities typical of routed infrastructures are available to VXLAN services.

The VXLAN ID is a 24-bit value – double the number of bits available with 802.1Q. This means that over 16 million VXLAN IDs are available, as compared to 4,000 VLAN IDs for traditional VLANs. VXLAN provides new functionalities that would be impossible with classic campus networking such as GBP discussed latter.

VXLAN Terminology


VXLAN

- Like classic VLAN without L3 IP interface
- Totally isolated

VNI

- VXLAN network identifier
- 24-bit value

VXLAN overlay between switches



The slide shows an example topology, which depicts a transport IP network composed of several switches, interconnected in a routed environment. You have several client hosts connected to this infrastructure.


In this example, you have not yet fully deployed the VXLAN functionality, with no VXLAN configuration on the switches. However, the configuration of the network includes VLAN IDs, and you have the switches configured to map VLAN IDs to VNIs (virtual network identifier) acting as VTEPs (entry points into the VXLANs). (You'll learn more about this in the upcoming slides.)

The S1 switch was configured with VXLAN 1001, which was bound to the VLAN 1001 (where PC1 connects). The printer was assigned a VLAN 2001 that maps to VXLAN 2001. Similar configurations exist on the other switches as well for their edge devices.

Switches will encapsulate the edge devices' layer 2 frames into an IP packet, including the appropriate VNI. The result is that PC 1 through 3 are functionally on the same layer 2 multipoint subnet. Similarly, virtual network 2001 contains hosts Printer 1, Printer 2 and the security camera. VXLANs 1001 and 2001 operate like a traditional classical VLAN without any layer 3 IP interface in them. They are totally isolated; and these VMs can communicate with each other, but they cannot communicate with any device outside of their own VXLAN.

VXLAN VTEP

- VXLAN tunnel end point (VTEP)
- Entry or “on-ramp” into the tunnel
- Device that interacts with other VTEPs



The VXLAN tunnel end Point (VTEP) provides an entry point into the VXLAN, while interacting with the other VTEPs to provide connectivity. In topology, each physical switch receives a role assignment. They are responsible for encapsulating traffic from the VXLANs into the tunnel and forwarding the resulting packets to their destination.

The receiving destination VTEP must decapsulate the packet and deliver it toward the intended target. Essentially, the VTEP provides an on-ramp to the VXLAN by performing this frame encapsulation and decapsulation service. VXLANs only need to be deployed at the edge switches of the transport network. Devices within the transport network perform typical layer 3 forwarding.

VXLAN Packet Structure

Header

- IP: 20 bytes
- UDP: 8 bytes
- **VNI: 3 Bytes (24 bits)**

50 bytes total

- 8 bytes VXLAN header
- 8 bytes UDP
- 20 bytes IP
- 14 bytes Ethernet

MTU greater than 1550 bytes

VXLAN Flags

Group Policy ID (User Role)

Virtual Network Identifier (VNI)

Reserved

Outer MAC

14 bytes

Outer IP

20 bytes

Outer UDP

8 bytes

VXLAN header

8 bytes

Original Ethernet frame

VXLAN encapsulated frame

FCS

Aruba recommends enabling jumbo frames with VXLAN

MOD 11- 8

VXLAN encapsulates layer 2 Ethernet frames in UDP packets. These VXLAN tunnels provide layer 2 virtualized network services to connected endpoints. A VTEP is the function within a switch that handles the origination or termination of VXLAN tunnels. Similar to a traditional VLAN ID, a VXLAN network identifier (VNI) identifies an isolated layer 2 segment in a VXLAN overlay topology. Symmetric Integrated Routing and Bridging (IRB) capability enables the overlay networks to support contiguous layer-2 forwarding.

An overlay network is implemented using VXLAN tunnels that provide both layer 2 and layer 3 virtualized network services to endpoints connected to the campus. The VNI associates tunneled traffic with the correct corresponding layer 2 VLAN or layer 3 route table so the receiving VTEP can forward the encapsulated frame appropriately. The symmetric IRB capability allows the overlay networks to support contiguous layer 2 forwarding and layer 3 routing across leaf nodes.

A VTEP encapsulates a frame in the following headers:

IP header: IP addresses in the header can be VTEPs or VXLAN multicast groups in the transport network. Intermediate devices between the source and destination forward VXLAN packets based on this outer IP header.

UDP header for VXLAN: The default VXLAN destination UDP port number is 4789.

VXLAN header: VXLAN information for the encapsulated frame.

- 8-bit VXLAN flags: The first bit signals if a GBP ID has been set on the packet and the fifth bit signals if the VNI is valid. All other bits are reserved and set to “0”.
- 16-bit VXLAN group policy ID: The group ID identifies the policy enforced on tunneled traffic. (This will be discussed later in the module.)
- 24-bit VXLAN Network Identifier: Specifies the virtual network identifier (VNI) of the encapsulated frame.
- 24-bit reserved field


Note: Configure jumbo frames on all underlay links in the fabric to ensure the accurate transport of additional encapsulation to ensure that packet fragmentation does not occur.

Confidential - For Training Purposes Only

VXLAN Tunnel

VXLAN tunnel

- Tunnel between 2 VTEPs




MOD 11-9

One VXLAN VTEP or tunnel end point communicating with another establishes a VXLAN tunnel, which is a requirement for each destination VTEP in the same VXLAN. In this example, switches S1, S2, and S3 are the physical VTEPs assigned to VNI 1001, so each of these switches maintains a tunnel with the other two. Clients assigned to VNI 2001 connect to switches S1, S3, and S4, with a similar set of interconnected tunnels.

The tunnel is merely a mechanism of transport through the IP network. Multiple VNIs can use one tunnel because each VTEP encapsulated packet contains the associated VNI. The network then uses this VNI to send packets to the appropriate tunnel endpoints. This is why traffic from multiple VXLANs can use the same tunnel. The underlying IP transport network uses IP unicasts to deliver broadcast, multicast, and unknown unicast frames for each VXLAN.

VNI “Virtual Switch”



MOD 11- 10

The VNI creates a virtual switched domain on top of an IP-based underlay network. A VXLAN segment differentiates individual logical networks so that several isolated layer 2 VXLAN networks can coexist on a common layer 3 infrastructure. Hosts in different VXLAN segments cannot communicate with each other directly. When multiple VXLANs are configured, each with a unique VNI, it is like having multiple virtual switches each making up a separate broadcast domain. In the example in the slide, there are 2 VXLANs configured: VNI 1001 and VNI 2001. Each VXLAN has VLANs mapped on VTEP devices to connect your endpoints to these virtual switches.


Static VXLAN vs EVPN

Static

- Tunnels manually created
- No scalability
- Difficult to maintain

Dynamic (EVPN)

- Tunnels dynamically created
- Scalable and flexible
- MP-BGP



There are two methods you can use to deploy a VXLAN infrastructure:

- Static VXLANs
- Ethernet VPN (EVPN) over a VXLAN

Let's introduce the two methods.

Static VXLANs

Static VXLAN use flood and learn (or ingress replication) to learn the MAC addresses of the remote hosts, which involves manual configurations of remote VTEPs in the flood list. Tunnels must be manually created, which is challenging in large networks. It is also difficult to maintain. Static VXLANs, therefore, are typically used in small deployments.

Ethernet VPN (EVPN) over VXLAN overview

Ethernet VPN (EVPN) is a standards-based BGP control plane that advertises MAC addresses, MAC and IP bindings, IP prefixes, and etcetera. The initial EVPN standard (RFC 7432) defined the BGP EVPN control plane and specifies an MPLS data plane. The control plane, containing an MPLS data plane, extends to consider additional data plane encapsulation models, including VXLAN, NVGRE, and MPLS over GRE. These are covered in detail in IETF RFC 8365. This section focuses on EVPN and its operation with a VXLAN

data plane for building overlay networks.

You can use MP-BGP EVPN to dynamically discover remote VTEPs and advertise MAC address and MAC/IP bindings in the VXLAN overlay, thus eliminating the flood and learn and the need for creating static tunnels. MP-BGP supports a new EVPN Network Layer Reachability Information (NLRI) carried in BGP, using multiprotocol BGP extensions with a newly defined address family identifier (AFI) and subsequent address family identifier (SAFI). A route distinguisher (RD) is a unique number prepended to the advertised address within the VRF that ensures support for overlapping IP addresses and MACs across different tenants. You can selectively import and export routes across VLANs/VRFs using a BGP-extended community called route target (RT) advertised along with the EVPN routes.

Auto-discovery via EVPN

VTEPs use BGP EVPN routes to discover VTEP neighbors, establish VXLAN tunnels, and assign the tunnels to VXLANs.

IMET route – VTEPs advertise their VXLAN IDs through IMET routes. If two VTEPs have the same VXLAN ID, they automatically establish a VXLAN tunnel and assign the tunnel to the VXLAN. You can use this method to create the replication list for the forwarding of broadcast, unknown, unicast, and multicast (BUM) traffic.

MAC/IP advertisement route and IP prefix advertisement route – In the EVPN deployment, VTEPs advertise MAC/IP advertisement routes or IP prefix advertisement routes, which carry export targets. When a VTEP receives a route, it compares the export targets of the route with the local import targets.

The full coverage and configuration of EVPNs with VXLANs is beyond the scope of this course.

VXLAN GBP



Let's now explore VXLAN group -based policies (GBPs).

Confidential - For Tra


GBP Introduction

Enhances Campus VXLAN

Virtual network based tunneling solution

Enables role-based policies

Role based policies not tied to IP addresses




MOD 11- 13

An overlay network is implemented using VXLAN tunnels that provide both layer 2 and layer 3 virtualized network services to endpoints connected to the campus. The VXLAN Network Identifier (VNI) associates tunneled traffic with the correct corresponding layer 2 VLAN or layer 3 route table so the receiving VTEP can forward the encapsulated frame appropriately. A VTEP encapsulates a frame in a header (see slide 8 and its notes for an explanation of these fields).


To enable GBP VTEPs, you must be able to map user roles to a group-based policy ID. In order to do that, VXLAN and GBP enabled VTEPs use the GBP ID field in the VXLAN packet header, which will contain the policy ID-to-role mapping value. When the G bit is set to 1 in the VXLAN flag field, this indicates that there is a value in the subsequent GBP ID field in the VXLAN header.

Problem with IP-based Policies


Role Contractor



Filtering done at the ingress switch



Destination subnet
10.1.200.0/24



Policy

Contractor	Printer	Allowed
Contractor	Medical gear	Denied
Employee	Printer	Allowed
Employee	Medical gear	Allowed

Impossible to implement with IP based policies


MOD 11-14

When using classic IP-based policies to control user traffic, classification and filtering is done at the ingress switch. Typically, IP-based policies rely on destination IP addressing and TCP/UDP ports to decide if the packet should be dropped or permitted. What if you want to control access to devices that use dynamic IP addresses and are connected to some remote switch? In the example shown in the slide, you want to allow contractors to access printers, but deny access to the medical gear. At the same time, employees should be able to access both printers and medical gear. Using standard IP-based traffic classification is not an option to implement such policies because both the printers and medical gear IoT devices use dynamic IP addressing. Notice, however, that all devices are assigned a user role after the authentication process completes.


The problem is that the remote switch is not aware of user role that is assigned to employee and contractor users—only the switch or switches connected to those users. GBP solves these problems by allowing you to build policies that use the user role as source or destination of the packet instead of the IP address. In order for this to work, filtering must be done at the remote switch and the source user role must be communicated in the data packet between the switches. Let's look at the details of how this works.

Micro-Segmentation Use Case

Clients assigned to the same VLAN/subnet on a single switch



Clients assigned to the same VLAN/subnet on different switches



Role based policy to prevent network connectivity between Contractor and Camera roles




MOD 11-15


With GBP you can permit or deny traffic between devices connected to the same switch or different switches. The VLAN or subnet of the device doesn't matter, because in a GBP policy you can use user role to specify source and destination of the packet. Source and destination devices could connect to the same VLAN as in the example shown in the slide, but might or might not be able to communicate with each other based on the roles assigned to these devices.

Macro-Segmentation Use Case

Clients assigned to different VLANs/subnets/VRFs on a single switch



Clients assigned to different VLANs/subnets/VRFs on different switches



Role based policy to prevent network connectivity between Contractor and Camera roles



MOD 11- 16

In this example in this slide, the source and destination devices connect to different VLANs on the same or different switches. Again, the VLAN or subnet of the device doesn't matter, as GBP policies use the user role to specify source and/or destination of the packet. In this scenario you can also permit/deny traffic between the contractor and camera device.

Role-Based Policy Details

GBP policy ID/role mapping is required on ingress & egress VTEPs

Works for traffic that is locally switched

Works for traffic exchanged across multiple VTEPs

Group policy ID is applied on ingress VTEP after device authentication

Policy is enforced on **egress** VTEP

Role names (e.g. "Contractor" or "Camera") are not sent between VTEPs

Role names are created locally on every required VTEP and mapped to group policy ID




MOD 11- 17

In the slide you see the details of the GBP implementation on AOS-CX switches. The role-to-policy mapping must exist on every edge switch. GBP policies are network-wide so that they can filter local traffic but also traffic exchanged across multiple switches (VTEPs). The group policy ID is applied on ingress switch after the authentication process completes. The GBP policy is enforced on the egress switch. Role names are not sent between switches; instead, each role has a corresponding policy ID that is carried with the user traffic encapsulated in its corresponding VXLAN packet header.

Role-to-Policy ID Mapping


```
gbp enable
gbp role-tag mapping Employee 100
gbp role-tag mapping Contractor 200
gbp role-tag mapping Camera 300
gbp role-tag mapping Printer 400
```

Role Contractor



Role Employee

Roles applied at ingress switch



Roles communicated to egress switch




MOD 11-18

Role names are created locally on every required VTEP and mapped to Group Policy ID

You enable GBP globally using the `gbp enable` command. Next you must map user roles names to group policy ID values. This must be done on every edge switch (VTEP). The policy ID value is carried with the user traffic encapsulated in the VXLAN. The remote switch can now read group policy ID indicated in VXLAN header and compare that with the local role-to-GBP mappings. It then knows what the role name is that is assigned to the source device on source VTEP and is able to enforce the GBP policy.

Policy Enforcement: IP Independent



MOD 11-19

Clients with the same role – default **allowed**, clients with different role – default **denied** +

Now let's go back to the problem defined at the beginning of this section. With GBP, local user roles are mapped to GBP ID values in the local switches' configurations. These GBP IDs are then communicated to the egress switch as part of the VXLAN header in the encapsulated user packets. When the employee user sends traffic to the medical gear device connected locally, local switch will permit this packet as per GBP policy.

When employee user sends a packet to the medical gear connected to the remote switch, the employee user role name will be translated to a GBP ID (as per mapping) and this ID is included in the VXLAN packet header sent to the remote VTEP. The remote switch reads the GBP ID and translates that to a role name: it now knows what the role name is of the employee user. The remote switch permits the packet as per the GBP policy.

Now, let's assume that the contractor user sends traffic to the medical gear connected to the remote switch. The same process takes place, where the Contractor user role name is mapped to the BGP ID, and the remote switch translates the GBP ID back to role name. It then matches that against GBP policies. This packet will be denied as per GBP policy configuration of the medical gear. By default, clients that have the same role will be allowed to exchange traffic and users with different roles will be denied any communication.

GBP Configuration

```
gbp enable
gbp role-tag mapping Employee 100
gbp role-tag mapping Contractor 200
gbp role-tag mapping Medical-gear 300
```

Enable GBP
Map roles to policy IDs

```
class gbp-ip IP-Class-Contractor
  10 match any any Contractor
class gbp-mac MAC-Class-Contractor
  10 match any Contractor arp
class gbp-ip IP-Class-Employee
  10 match any any Employee
class gbp-mac MAC-Class-Employee
  10 match any Employee arp
class gbp-ip IP-Class-Medical-gear
  10 match any any Medical-gear
class gbp-mac MAC-Class-Medical-gear
  10 match any Medical-gear arp
class gbp-ip Allow-Employee-Medical
  10 match udp Employee Medical-gear eq 8644
class gbp-ip Deny-Contractor-Medical
  10 match any Contractor Medical-gear
```

Setup IP and MAC classes

- gbp-ip : matches IP traffic
- gbp-mac: matches L2 traffic

MOD 11- 20

After creating the mappings (see the top part of the slide), you need to define GBP traffic classes. There are 2 types: IP and MAC. IP GBP traffic classes will match L3 traffic while MAC traffic classes are designed to match L2 traffic. Notice that with GBP you can specify the source and destination of the traffic using role names instead of IP addresses. In addition to that you can still match specific TCP/UDP ports or protocol types. See the bottom part of the slide for an example GBP traffic class configuration.

GBP Configuration

```
port-access gbp employee
  10 class gbp-mac MAC-Class-Employee
  20 class gbp-ip IP-Class-Employee
port-access gbp contractor
  10 class gbp-mac MAC-Class-Contractor
  20 class gbp-ip IP-Class-Contractor
port-access gbp medical-gear
  10 class gbp-ip Deny-Contractor-Medical action drop
  20 class gbp-ip Allow-Employee-Medical
  30 class gbp-mac MAC-Class-Medical-gear
  40 class gbp-ip IP-Class-Medical-gear
```

Create policies

```
port-access role Employee
  associate gbp employee
port-access role Contractor
  associate gbp contractor
port-access role Medical-gear
  associate gbp medical-gear
```

Assign policies to roles



MOD 11- 21

Next, you need to create the BP policies that would permit or deny traffic (top part of the slide). No action explicitly entered means that this traffic should be allowed (no need to define a permit action in the syntax). When you add the action drop syntax, that means you want to deny the specified traffic. After configuring GBP policies, you then assigned them to user roles (bottom part of the slide).

NetConductor

Managing policies in large network can be challenging



Aruba NetConductor

Policy manager

- Defines user and device groups
- Creates the associated traffic routing and access enforcement rules for the physical network

Group policy identifier

- Carries configuration and client policy information
- Reduces configuration and security overhead
- Increases mobility and scalability.

Fabric wizard

- Simplifies the creation of the overlays
- Intuitive GUI
- Eases configuration of switches and gateways



MOD 11- 22

As you have learned in the previous examples, GBP management in large environments can be challenging. Building a distributed overlay network, configuring the VXLAN-EVPN fabric, and setting up policies manually would be time consuming and hard to maintain.

Aruba NetConductor is cloud-based solution that can automate and simplify many processes and tasks related to VXLAN-EVPN fabric. Its main components include:

Policy manager: defines the user and device groups and creates the associated traffic routing and access enforcement rules for the physical network.

Group policy identifier (GPID): carries the configuration and client policy information in traffic for in-line policy enforcement, which reduces configuration and security overhead and increases mobility and scalability.

Fabric wizard: simplifies the creation of the overlays using an intuitive, graphical user interface, greatly easing the way virtual components are defined and network configuration instructions are generated and pushed to switches and gateways.

To learn more about Aruba NetConductor, visit <https://www.arubanetworks.com/products/network-management-operations/central/netconductor/>.

Note: Aruba Fabric Composer (AFC) is another stand-alone option that can be used to

manage and deploy this type of solution. AFC is beyond the scope of this course.

Confidential - For Training Purposes Only



Knowledge Check

Self-check on key learning points



It's time for a knowledge check to review what you have learned.

Confidential - For Training Purposes Only

Question #1

VXLAN relies on a VXLAN network identifier (VNI). This is a 24-bit value that can support over 16 million VXLANs.

- True
- False

Q

Knowledge Check

MOD 11-24

+

true

Confidential - For Training Purposes Only

Question #2

With VXLAN technology, a GRE tunnel is created between two VXLAN tunnel end points (VTEP).

- True
- False

a

Knowledge Check

MOD 11-25

+

false

Confidential - For Training Purposes Only

Question #3

Which of the options below accurately describe Ethernet virtual private network (EVPN)?

- a. It improves VXLAN manageability in large networks by dynamically creating tunnels
- b. It is a standards-based BGP control plane that advertises MAC-to-IP bindings, and more
- c. VTEPs use BGP EVPN routes to discover VTEP neighbors and establish VXLAN tunnels
- d. It has a data plane component that secures VTEP-to-VTEP traffic

a

Knowledge Check

MOD 11-26

+

a, b, c

Confidential - For Training Purposes Only

Question #4

One advantage of VXLAN group-based policy (GBP) is that it eliminates the need to focus on IP addressing to control access.

- True
- False

Q

Knowledge Check

MOD 11- 27

+

true

Confidential - For Training Purposes Only

Question #5

One caveat related to VXLAN group-based policy (GBP) is that you cannot control traffic between devices connected to the same switch.

- True
- False

Q

Knowledge Check

MOD 11-28

+

false

Confidential - For Training Purposes Only



Lab Activity

Let's review the lab activity that is associated with this module.

Confidential - For Training

Lab Tasks

Lab 11.01 - Group Based Policies with EVPN

Task 1

- Prepare your lab environment

Task 2

- Verify the Group Based Policy Configuration

Task 3

- Configure Access Control between Roles

**Module 12 -
Security and
Availability
Features**




MOD 11- 30

The slide provides a brief review of tasks performed in this lab. Please see your Lab guide for details of this lab.

Security and Availability Features

Implementing Aruba Campus Access

Module 12 covers additional security and availability features. You'll learn about the intrusion detection system (IDS) and intrusion prevention system (IPS) capabilities of Aruba wireless products, how to configure authentication caching for authentication survivability in a campus network, how to deal with connection issues from your campus to Aruba Central, and managing administrative accounts on Aruba switches.



Objectives

- Describe IDS/IPS capabilities and use cases
- Configure authentication survivability
- Explain Aruba Central survivability
- Describe management user authentication and management roles

After completing this module, you will be able to:

- Describe IDS/IPS capabilities and use cases.
- Configure authentication survivability.
- Explain Aruba Central survivability.
- Describe management user authentication and management roles.

IDS/IPS

Let's explore the IDS and IPS capabilities first.

Confidential - For Training

Devices That Discover Threats

Access mode - default

Provides WLAN service to clients

IDS scan: All regulatory channels

Containment: Best effort, own channel

Spectrum analysis: Own channel

Air Monitor (AM)

IDS scan: All regulatory channels, Total Watch™


Containment: All channels

Spectrum analysis: None

Spectrum Monitor

IDS scan & containment: None

Spectrum analysis: All regulatory channels



MOD 12- 4

Take a moment to review the three modes in which Aruba APs can function to support discovery of wireless threats and issues.

Hybrid mode APs

APs are primarily responsible for serving WLAN clients. However, they can operate in hybrid mode and devote some of their time to detecting and mitigating threats as well. Hybrid mode APs go off channel every 10 seconds to scan one channel for about 100 milliseconds (ms). If they are serving clients, the APs back off the scans. Hybrid mode APs can typically sweep the regulatory channels and detect all wireless devices in the area within five to six minutes. Hybrid mode APs are capable of containing wireless devices that IDS/IPS has determined pose a risk, but only on their own channel and in a best effort manner. The hybrid mode APs also support spectrum analysis to help maintain RF health. Hybrid mode operation is the default.

Air monitors (AMs)

Also called wireless sensors, AMs are dedicated to scanning for rogue APs or other IDS events. They passively scan layer 2 802.11 frames. As passive devices, they listen for frames, but do not transmit RF and do not service clients. AMs move from channel to channel, listening for threats. They listen on all channels, not just channels that are permitted in the region, including rare ones. Because they do not transmit RF, they can do so without

violating the law. With ArubaOS Total Watch, the AMs spend a longer scan time on active channels with detected wireless devices. AMs can contain wireless devices that pose a risk on any channel. If the solution requires intrusion prevention rather than just detection, Aruba recommends deploying AMs.


Spectrum monitors (SMs)

Like AMs, spectrum monitors (SMs) do not service clients. However, they examine layer 1 RF data, rather than layer 2 like AMs. By sending captured RF data to Aruba Central for spectrum analysis, SMs help companies detect RF noise and interference issues. Such interference could be introduced maliciously, but is usually done so accidentally. Examples of devices that cause RF interference include non-802.11 devices like microwave ovens, baby monitors, video cameras, and Bluetooth devices; many other such devices exist as well.

Note: You would typically deploy AMs to support your security efforts and SMs to monitor and maintain the environment's RF health.

Confidential - For Training Purposes Only


IDS/WID Profile: Detection



MOD 12- 5

To configure threat detection, select a Group context and then navigate to Devices > Config > Access Points > Security > Wireless IDS/IPS > Detection. You can use predefined detection levels or set up a custom policy. There are two sections to define policies: one for APs and one for clients. In the slide, you can see the infrastructure (AP) section.

IDS/IDS/IPS Profile: Protection



MOD 12- 6

This is a continuation of the last slide, where you can also see the client section.

Note: The IDS/IPS capability is rule-based: you enable the rules that you need and disable the ones that you don't want. This can be done using the slider and setting the value to high, low, or off, which turns on or off specific rules; or you can select custom and choose the rules that you do or do not want to use individually.

Viewing AP Classifications

Customer: HPE hqlow

Campus1

Manage

Overview

Devices

Clients

Guests

Applications

Security

Analyze

Alerts & Events

Audit Trail

Tools

Reports

RAPIDS Gateway IDS/IPS

IDS Rogues

TOTAL	ROGUES	SUSPECTED ROGUES	INTERFERING	NEIGHBORS	MANUALLY CONTAINED
86	4	3	79	0	0


Detected Access Points

Name	Classification	SSID	Last Seen
Netcore Tech-5F:32:40	Interfering	tomaszek_5G	22 Nov 2022, 14:11:26
ASUSTek COMP-3A:5F:40	Interfering	WutekNET	22 Nov 2022, 15:21:54
08:A7:C0:BE:A7:C0	Interfering	CGA2121_x4cjmqc	22 Nov 2022, 15:06:26
1C:9E:CC:69:16:F0	Interfering	2.4G-vnet-819932	22 Nov 2022, 14:15:34
ARRIS Group,-39:17:80	Interfering	UPC249935047	22 Nov 2022, 13:48:44
Compal Broad-3C:1C:B0	Interfering	UPC8155823	22 Nov 2022, 13:59:39
Compal Broad-6B:41:B0	Interfering	UPC8411359	22 Nov 2022, 14:49:44
Compal Broad-E8:0F:10	Interfering	UPC8869904	22 Nov 2022, 13:48:44
Compal Broad-E8:0F:20	Interfering	UPC8869904	22 Nov 2022, 14:01:26
36:23:87:43:30:10	Interfering	DIRECT-bD-BRAVIA	22 Nov 2022, 14:41:54
36:2C:94:6B:41:B0	Interfering	Wi-Free #InternetUPCNa...	22 Nov 2022, 14:46:48

MOD 12- 7

To review the discovered (non-company) APs along with their classification select a particular context and navigate to Devices > Security > RAPIDS > Rogues. Here you will see the list of all detected APs.

Interpreting AP Classifications




MOD 12- 8

As you see from the slide, IDS/IPS classifies all detected AP radios. To interpret what you see, you need to understand what these classifications mean.

IDS/IPS first divides APs into Authorized APs and Unauthorized APs. Authorized APs are APs controlled by your company and have been added to Central. All other APs are unauthorized.

IDS/IPS further classifies unauthorized APs based on whether they connect to the corporate LAN. Unauthorized APs that are connected to the LAN are called rogue APs. Unconnected, unauthorized APs are called either interfering APs or neighbor APs. Interfering APs are APs that are deployed on the company property, but are not connected to the LAN. Neighbor APs are legitimate APs that belong to neighboring organizations. When a corporation shares a building or even is just close to other buildings, some neighbor APs will be detected. By default, IDS/IPS classifies the unconnected APs as interfering. Admins can create rules, though, that define specific SSIDs as neighbor SSIDs. You can also manually re-classify an interfering AP as a neighbor.

Rogue APs



MOD 12- 9

Unauthorized APs that are connected to the LAN are called rogue APs. If a rogue AP or interfering AP have the same SSID as your authorized APs, IDS/IPS also calls them honeypot APs.

IDS/IPS also supports detecting ad-hoc networks, which will be listed among detected radios. The 802.11 standard defines ad hoc networks as a way for stations to connect together as peers without the need of an AP. The client that hosts the ad hoc network could be connected to the LAN, which means that it is a rogue, or not, or it could mean that the ad-hoc client is interfering or possibly a neighbor.



Ways to Mitigate Threats

Contain

APs: Prevent clients from connecting
Clients: Prevent from connecting to an AP

Protect valid clients

Disconnect only authorized clients from non-authorized APs



MOD 12- 10

Before you look at determining how to deal with detected alerts and IDS/IPS events, briefly review what containment of an AP means. When IDS/IPS contains an AP, it prevents all clients from connecting to it. When IDS/IPS contains a client, your valid APs target the client's association and prevents the client from connecting to the rogue AP.

Instead of completely containing a rogue AP or ad hoc network, you can also choose to protect valid clients. This feature prevents any valid clients from connecting to a non-authorized AP. However, the APs or AMs allow non-valid clients to connect to the rogue APs. This approach can offer a good option for companies that want to protect their own users and devices from man-in-the-middle (MitM) attacks but want to avoid potentially interfering with neighbors.

Aruba IDS/IPS provides three containment methods, two of which are illustrated in the slide. Wired AP containment poisons the rogue AP's ARP table so that the AP's traffic does not reach the corporate default gateway. Some companies prefer to use this form of containment because it can only affect devices that are truly connected to your LAN; wired containment cannot accidentally contain a legitimate neighbor's AP that is on a neighboring network.


When using the de-authentication method, APs or, more often, AMs send de-authentication frames to disconnect clients from the contained (rogue) AP. The drawback of this method is that the AP or AM must continue to send these spoofed frames as clients usually continually try to reconnect to the rogue AP.

The third option is covered next.

Confidential - For Training Purposes Only

Tarpit Containment

- Deauth + fake associations to prevent clients from reconnecting
- Lower bandwidth and airtime requirements than de-authentication containment



MOD 12- 11

Tarpit containment is the third option and can be more effective than de-authentication containment. In addition to sending de-authentication messages, the AM spoofs probe responses and association responses on a different channel or BSSID from the rogue AP. The idea is to get the clients away from the rogue and connect to the AM instead. Once a client is associated with the AM, it stops trying to reconnect to the rogue AP. This process is in called a tarpit. When all clients that try to connect to the rogue AP are in the tarpit, the AM no longer needs to send out frames, so this method imposes lower bandwidth requirements. The AM also does not need to spend as much time on the channel, so it can continue to scan for threats more effectively. Note, though, that the AM is spoofing what looks like a valid SSID. When clients connect, they will not have access to internal resources. The idea behind tarpitting is to get the clients off the rogue device and connected to something else in the most efficient way possible.

Administrators can choose to implement tarpit containment against all clients connected to a rogue AP or just non-valid clients. Since valid clients might accidentally connect to a rogue AP, the company might not want to punish them. The latter approach protects valid clients—the company's own clients—from being stuck in the tarpit and lets them connect to an authorized AP instead.


Wired AP containment works against rogue APs that are wireless routers. It poisons the AP's ARP table so that the AP's traffic does not reach the corporate default gateway. Some companies prefer to use this form of containment because it can only affect devices that are

truly connected to your LAN; wired containment cannot accidentally contain a legitimate neighbor's AP.

Instead of completely containing a rogue AP or ad hoc network, you might choose to protect valid clients. This feature prevents any valid clients from connecting to a non-authorized AP by having hybrid APs or AMs send disassociation frames targeted to those specific clients. However, the APs or AMs allow non-valid clients to connect to the rogue APs. This approach can offer a good option for companies that want to protect their own users and devices from MITM attacks but want to avoid potentially interfering with neighbors.

Confidential - For Training Purposes Only

Containment



You can select containment method by first selecting a context and then navigating to Devices > Access Points > Config > Security > Wireless IDS/IPS > Protection > Containment Methods.

Warning: Care should be taken if you enable one of the proactive containment options. If you tarpit or prohibit a valid client from connecting to a neighbor network, you could be in violation to the law depending on the country in which you reside.

Authentication Survivability




Let's move on to authentication survivability with both the AOS-CX switches and Aruba APs and gateways.

Confidential - For Tra

Authentication Caching for Switches

802.1X and MAC re-authentications succeed when RADIUS server unavailable

Users already authenticated retain their currently assigned RADIUS attributes



```
aaa authentication port-access dot1x authenticator cached-reauth
aaa authentication port-access dot1x authenticator cached-reauth-period <PERIOD>
```

```
aaa authentication port-access mac-auth cached-reauth
aaa authentication port-access mac-auth cached-reauth-period <PERIOD>
```


MOD 12-14

The AOS-CX `aaa authentication port-access dot1x authenticator cached-reauth` command configures the period during which an authenticated client, which has failed to reauthenticate because the RADIUS server(s) is unreachable, remains authenticated. The default cache re-authentication period is 30 seconds.

Authentication Survivability for Gateways

Allows user to connect if there is no connectivity with RADIUS server

Maximum connectivity lost duration is 7 days



Currently supported in mixed and tunneled Mode **ONLY** !
EAP-TLS is the *only* supported authentication mode



MOD 12- 15

Authentication survivability is required when a remote link failure occurs between a gateway cluster and an authentication server that is either in the cloud or a data center. If the connectivity between the gateway cluster and the authentication server is lost for a maximum duration of 7 days, authentication survivability ensures that the known users can securely join the network even if the authentication server is unavailable. The gateway will first try to reach the RADIUS server and if it's not responding, the gateway will revert to survivability mode. This feature is currently supported in tunnel and mixed mode deployments. Also, EAP-TLS is currently the only supported authentication mode.

Authentication Survivability Configuration

The screenshot displays the configuration interface for Authentication Survivability. The navigation pane on the left shows the hierarchy: Gateways > Security > Auth Servers. The main content area is titled 'Authentication Servers' and contains a section for 'Authentication Survivability'. This section includes a toggle switch for 'Enable authentication survivability' which is currently turned on. Below this are three dropdown menus: 'Local cache lifetime' set to '7 days', 'CA certificate' set to '-Choose an option-', and 'Authentication server certificate' set to '-NONE-'. A large orange box highlights the entire 'Authentication Survivability' configuration area.

MOD 12- 16

To configure authentication survivability on a gateway group, complete the following steps:

Configure a gateway group: Set the context filter to a group containing at least one gateway. The dashboard context for a group is displayed.

Under Manage, navigate to Devices > Gateways.

Click the Config icon to view the group configuration dashboard.

Click Advanced Mode and then click Security.

Under Security, click the third-level tab of Auth Servers.

In the resultant page, scroll to the bottom and expand the Authentication Survivability drop-down.

Slide the toggle switch next to Enable authentication survivability.

Select a value within the range of 1 to 7 days from the Local cache lifetime drop-down to set the duration after which the authenticated credentials in the cache expires. When the cache expires, the clients are required to authenticate again with the external authentication server.

From the CA Certificate drop-down, select the client's CA certificate to be configured as

Trusted CA cert on the Gateway device. You can add multiple CA certificates in this field.

From the Authentication server certificate drop-down, select a server certificate used by the local survival server to associate EAP-TLS for 802.1X authentication.

Click Save Settings.

Confidential - For Training Purposes Only



Central Survivability

What if the connection between your company's network and Aruba Central fails? Obviously best practice is to ensure you have two different ISP connections for redundancy. But what if there's a catastrophic failure and both go down? Let's investigate what happens.

Confidential for Trade


Key Management Survivability

Existing users

- Can continuously do fast roaming
- PMK or R1 keys exist for 8 hours

New users

- PMK or R1 keys not shared between APs
- Must re-authenticate when roaming



MOD 12- 18


In case Central is not available, existing users can continuously do fast roaming since PMK or R1 keys of existing users are cached in each AP for 8 hours.

PMK or R1 keys cannot be shared between APs without the Key Management Service (KMS), which implies that new users must run the complete authentication process when roaming.

Tunnels Rekeying

Keys

- Valid for 36h / AP-GW
- 12h before expiration, new keys are generated



MOD 12- 19

Aruba Overlay Tunnel Orchestrator (OTO) service is responsible for changing IPsec key on a regular basis for the secure IPsec connection (this is done on both SPIs that make up that connection). Keys pushed to the APs and gateways and remain valid for 36 hours. 12 hours before expiration, new keys are generated by the OTO server and pushed to the APs and the gateways.

Tunnels Survivability

Initial Orchestration


Cloud is mandatory for initial tunnel orchestration

Survivability

Tunnel config from tunnel orchestrator is preserved in both AP and GW

Legacy Mode

When both AP & GW fails rekeying, AP and GW start legacy IKE process



MOD 12- 20

Aruba Central connectivity is required for the initial tunnel orchestration. If the connection to Central fails, APs and gateways revert to legacy mode. The current tunnel configuration is preserved on both APs and gateways; and for authentication, APs and gateways use TPM certificates. If connectivity to Aruba Central is restored, the APs and gateways remain in survivability mode until the next rekeying.


Management User Authentication



You need to manage your network. Let's look at controlling access to AOS-CX switches.

Confidential - For Tra

Switch Manager Roles



Admins

CLI and Web UI: Full access, including user accounts

Operators

CLI and Web UI: Display-only access

Auditors

CLI: Auditor context commands
Web UI: System > Log page



MOD 12- 22

When a manager authenticates, that manager is assigned a role. AOS-CX switches separate management users into three different default roles:

- Administrators
- Operators
- Auditors

All types of users may access the switch through the CLI or Web UI (or REST API, as well); however, their privileges differ. Administrators have full privileges, including full read-write access to the CLI and Web UI. They can configure any settings, including user accounts.

Operators have no configuration privileges. They can only view non-sensitive switch configuration information. In the CLI, operators can only issue show commands, and they can only view settings in the Web UI.

Finally, auditors have access to CLI commands only in the auditor context (auditor>). Auditors can view event log contents and local accounting log contents. They can also copy command output to a remote server or a local USB drive.

Admins can create custom manager groups as well, but that configuration is beyond the

scope of this course.


Confidential - For Training Purposes Only

How AOS-CX Switch Credentials Are Stored

Default Stored encrypted in the config

Options Set encryption key with the **service export password**

Must be set if config applies to multiple switches



```
hostname Switch
user admin group administrators password ciphertext AQBapT8B3l5gKNNhf1ulTru
```

MOD 12- 23

All AOS-CX credentials are stored in the switch configuration with encryption. The service export password acts as the key for this encryption, so you should set it in the config if you want to deploy the same config across multiple switches.

Best Practices for AOS-CX Local Authentication

Default

Admin account has no password
Prompt to define it at first log in

Recommended

At least one local admin account
with a secure password

Optional

Password length and complexity rules



MOD 12- 24

On AOS-CX switches you should define at least one account for local authentication. The switches have a default admin account, but no default password is defined for it. Instead, the first manager to log in is prompted to set this password. You can also create additional local accounts and assign them to groups, if desired.

You can configure password length and complexity requirements to ensure that admins create proper passwords for other accounts. You create a password policy which specifies settings such as minimum length and required characters. You then activate the policy. If, for future passwords, admins try to create a password that does not meet the requirements, the switch registers an error and the admin is prompted to define a different password.



Knowledge Check

Self-check on key learning points




It's time for a knowledge check to review what you have learned.

Confidential - For Training Purposes Only

Question #1

Match each component with its appropriate description.



Knowledge Check ✓



MOD 12- 26



Access mode

Scans all regulatory channels for threats

Performs spectrum analysis on its own channel

Air monitor mode

Spends more time scanning all regulatory channels for threats

Does not service clients

Does not perform spectrum analysis

Spectrum monitor mode

Does not scan any channels for threats

Performs spectrum analysis on all regulatory channels

Question #2

Aruba AP IDS/IPS features classifies all detected AP radios as either rogues, suspected rogues, interfering, or neighbor.

- True
- False

a

Knowledge Check

MOD 12- 27

+

true

Question #3

Which of the following options accurately describe Aruba's authentication resiliency features?

- a. Authentication caching enables 802.1X and MAC re-authentications to succeed when the RADIUS server is not available
- b. You configure authentication caching on Aruba gateways
- c. Authentication survivability allows previously authenticated and new users to connect if connectivity is lost between the gateway cluster and the authentication (RADIUS) server
- d. Authentication survivability is supported for APs operating in mixed, bridged, and tunneled mode

a

Knowledge Check

MOD 12- 28

+

a, c

Confidential - For Training Purposes Only

Question #4

If connectivity to Aruba Central fails, APs and gateways revert to legacy mode for their tunnel authentication.

- True
- False

a

Knowledge Check

MOD 12- 29

+

true



Lab Activity

Let's review the lab activity that is associated with this module.

Confidential - For Training

Lab Tasks

Lab 12.01 - Service Survivability

Task 1

- Tunneled WLAN Central Survivability

Task 2

- Wired Cached Re-Authentication and Critical Role

Lab 12.02

Lab 12.02 - Admin Authentication

Task 1

- Gateway Admin Authentication

Task 2

- Switch Admin Authentication

Module 13 - Traffic optimization and QOS



MOD 12- 31

The slide provides a brief review of tasks performed in this lab. Please see your Lab guide for details of this lab.

Rev. 23.11




Traffic Optimization and QoS

Implementing Aruba Campus Access



Welcome to module 13 on traffic optimization and quality of service (QoS). In this module you'll learn how to correctly identify traffic types that needed to be treated differently in your network and how to handle those traffic types when Aruba wireless and wired solutions are deployed.

Confidential for Training



Objectives

- Leverage AirMatch to tune the WLAN RF environment for optimal performance
- Use other features to improve the WLAN end user experience: ClientMatch and broadcast/multicast optimization
- Describe and apply features like Aruba AirGroup for improved “zero configuration networking”
- Describe and configure key QoS features to ensure the best service for the most important traffic
- Improve real-time traffic performance over WLANs with Aruba’s (UCC) service
- Configure LLDP and LLDP-MED to help discover and validate network connectivity and improve operations for connected VoIP phones and APs

After completing this module, you will be able to:


- Leverage AirMatch to tune the WLAN RF environment for optimal performance.
- Use other features to improve the WLAN end user experience – ClientMatch, broadcast and multicast optimization.
- Describe and apply features like Aruba AirGroup for improved zero configuration networking.
- Describe and configure key Quality of Service features to ensure the best service for the most important traffic.
- Improve real-time traffic performance over WLANs with Aruba’s Unified Communication and Collaboration (UCC) service.
- Configure LLDP and LLDP-MED to help discover and validate network connectivity and improve operations for connected VoIP phones and APs.

WLAN Optimization

Improving the end user WLAN experience is important and this section covers Aruba's WLAN optimization features.

Confidential - For Training

Channel Availability



MOD 13- 4

Channel availability is a challenge to wireless network design and troubleshooting in every changing radio frequency (RF) environment. RF interference from other 802.11 devices and non-802.11 RF transmitters, such as microwave ovens and Bluetooth devices, can cause signal degradation, poor performance, and dropped connections. Interference is both inevitable and unpredictable. The RF interference varies by device (microwaves, cordless phones, and etcetera), usage patterns (time of day), location (local emissions regulations) and construction. You must be able to manage the RF to ensure a reliable Wi-Fi operation.

You will now look at some of the factors that cause interference and poor communication, resulting in low throughput in a WLAN. The 5 and 6 GHz spectrums are usually cleaner air, as there are more available channels, less co-channel interference, and fewer non-802.11 interferers. Most of the devices shown in the example in the slide are 2.4 GHz devices causing interference.

AirMatch Function


Cloud hosted **RF optimization** service

Dynamically adapts to changing RF conditions

Past 24 hours RF info used for calculation

Channel & power deployment once a day – local time

Mitigate co-channel interference



MOD 13- 5

Aruba's AirMatch provides automated dynamic RF optimization to improve the wireless user experience by providing a stable network that dynamically adapts to the changing RF conditions. AirMatch, which runs now in the Aruba Central cloud, is a centralized RF planning and optimization service. It models the network as a whole and devises a channel and power plan for the RF environment. This service receives telemetry data from APs, including radio measurements, channel range, EIRP range, operational conditions, and local RF events (like radar detection or high noise). AirMatch processes the RF information and generates an RF solution that specifies new channel, bandwidth, EIRP, and mode of operation for every radio. AirMatch also minimizes channel coupling, where adjacent radios are assigned to the same channel.

The AP receives the information and applies the change(s) in its radios' configuration. AirMatch calculates an RF channel and power level solution based on the past 24 hours RF information collected for each AP. Channel and power deployment only happens once a day at 5 AM, by default. This time can be changed: deploy channel plan based on a configured deployment time and considering local time zones.

AirMatch provides neighboring information to other Aruba cloud services, such as AirGroup and Key Management services. Also, AirMatch provides AP partitions for the live upgrade process.

AirMatch Service Configuration

Global

Manage

- Overview
- Devices**
- Clients
- Guests
- Applications
- Security
- Network Services

Access Points | Switches | Gateways

Radios
RF management configuration to optimize the wireless coverage for network.

ACTIVATE OPTIMIZATION ☒

Automatically deploy optimization at
05:00

WIRELESS COVERAGE TUNING

Conservative | **Balanced (Recommended)** | Aggressive

Enabled by default

Config



MOD 13- 6

From the context filter, select Global and navigate to Device > Access Point > Configuration. You will see that AirMatch is enabled by default. (Curiously the name displayed is not AirMatch but Radio optimization). By default, channel and power deployment happens at 5 AM every day, but this can be modified.

The system allows to set the wireless coverage: the balanced option is set by default. The aggressive option increases the RF signal overlapping between APs; this is opposite to the conservative option which decreases the signal overlapping.


Client Match

Aruba patented technology

Load balancing

Band steering and balancing

Sticky clients



MOD 13- 7

ClientMatch is an Aruba patented technology that provides client load balancing across radio bands and AP radios. ClientMatch helps to influence sticky clients to connect to the best radio. When clients associate to the WLAN, their wireless NIC's firmware determines which radio, and on which AP, the client will connect. Sometimes the vendor's firmware is intelligent, and looks at criteria like the strength of the signal and the SNR, to name a couple, when determining which radio to associate. Other times, a vendor's firmware might use the last association the NIC had when it was in this area (this is called a sticky client). A lot of things could have changed since then, plus, the client doesn't understand what's actually happening on a radio when it connects: it doesn't know how many other devices are currently associated to that radio.

ClientMatch is a feature that matches a client to the most appropriate AP and frequency in the area. When APs see a probe request from a user, they all report that to Aruba Central. Aruba Central then looks at a lot of criteria to determine which AP and radio would be best and this is shared with the original AP that received the probe. The AP will then redirect the user to the more appropriate radio, assuming that the most appropriate radio is not itself. This is best done if the client supports IEEE 802.11v, which allows for this interaction between the client and the AP, which any client product sold today should support it. ClientMatch thus load balances clients and if a client tries to connect to a less preferred frequency, like 2.4 GHz, the client can be redirected to a 5 GHz or 6 GHz radio.

Broadcast Optimization

The screenshot shows the Aruba configuration interface for 'Access Points' under the 'WLANs' tab. The configuration is for 'AO10-Corp'. The 'General' tab is selected, and 'Advanced Settings' is expanded. The 'Broadcast/Multicast' option is highlighted with an orange box. A dropdown menu is open for 'Broadcast filtering', showing options: 'Disabled' (selected), 'ARP', 'ALL', and 'Unicast ARP Only'. The 'DTIM Interval' and 'Multicast Transmission Optimization' fields are also visible. A watermark 'Aruba Training Purposes Only' is overlaid on the image. The bottom right corner of the interface shows 'MOD 13- 8'.

Broadcast optimization is a unique Aruba feature. When a client generates a broadcast in a broadcast domain, this is typically flooded in most network scenarios. However, the issue with wireless is that these broadcasts use up a lot of bandwidth. Aruba's feature allows the AP to proxy this information and handle it intelligently. For example, if a client ARPs for another client's MAC address, instead of the AP flooding it, the AP will look at its local ARP table first and if it finds the results, sends the results back as a unicast to the requesting client.

This is tuned per WLAN in the WLAN wizard. From the General tab, expand Advanced Settings, expand Broadcast/multicast and select any of the following values:

All—The AP drops all broadcast and multicast frames except DHCP and ARP, IGMP group queries, and IPv6 neighbor discovery protocols.

ARP—The AP drops broadcast and multicast frames except DHCP and ARP, IGMP group queries, and IPv6 neighbor discovery protocols. Additionally, it converts ARP requests to unicast and sends frames directly to the associated clients.

Unicast ARP Only—This option enables AP to convert ARP requests to unicast frames thereby sending them to the associated clients.


Disabled—The AP forwards all the broadcast and multicast traffic is forwarded to the wireless interfaces. Default mode.

For more information on how this works and additional tuning options, see:
https://www.arubanetworks.com/techdocs/central/latest/content/aos10x/cfg/config_wlan_ssid.htm?Highlight=Broadcast%20filtering.

to the wireless interfaces. Default mode.

Confidential - For Training Purposes Only


Multicast Transmission Optimization



The default rate for sending frames for 2.4 GHz is 1 Mbps, and for 5 GHz is 6 Mbps MOD 13-9

To turn on multicast transmission optimization, access the WLAN wizard. From the General tab, expand Advanced Settings, expand Broadcast/multicast. Turn on the toggle switch to allow the AP to select the optimal rate for sending broadcast and multicast frames based on the lowest of unicast rates across all associated clients. When this option is enabled, multicast traffic can be sent up to a rate of 24 Mbps. The default rate for sending frames for 2.4 GHz is 1 Mbps, and that for 5 GHz is 6 Mbps. This option is disabled, by default.

Dynamic Multicast Optimization



Convert multicast to unicast

Clients receive traffic with a highest possible rate
Less airtime required to deliver the traffic

Large number of multicast clients may cause a problem

Set threshold, if exceeded AP reverts to multicast mode

- Channel utilization threshold
- Number of clients threshold

MOD 13- 10

Dynamic multicast optimization (DMO) will convert the downstream multicast traffic to unicasts before forwarding it to the wireless clients. This allows the client to receive the traffic with the highest possible data rate and thus requires less airtime to deliver the traffic.

To enable DMO, access the WLAN wizard. From the General tab, expand Advanced Settings, expand Broadcast/multicast. Turn on the toggle switch to allow the AP to convert multicast streams into unicast streams over the wireless link. Enabling DMO enhances the quality and reliability of streaming video, while preserving the bandwidth available to the non-video clients.

NOTE: When you enable DMO on multicast SSID profiles, ensure that the DMO feature is enabled on all SSIDs configured in the same VLAN.

DMO channel utilization threshold

Specify a value to set a threshold for DMO channel utilization. With DMO, the AP converts multicast streams into unicast streams as long as the channel utilization does not exceed this threshold. The default value is 90% and the maximum threshold value is 100%. When the threshold is reached or exceeds the maximum value, the AP sends multicast traffic over the wireless link.

NOTE: This option will be enabled only when Dynamic Multicast Optimization is enabled.

DMO Client Threshold

Specify a value between 2 and 255 to set the DMO client threshold. The recommended initial DMO client threshold is 40, it should be then adjusted based on the client's experience

Confidential - For Training Purposes Only

AirGroup



In the last section you learned about controlling broadcast mechanisms using the broadcast optimization feature. Let's now discuss how AirGroup can be used to optimize a similar issue with multicast transmissions.

Confidential For Training

Zero Configuration Networking

Enables service discovery: printers, speakers, screen mirroring, etc.

Designed for single subnet IP networks

Apple Bonjour uses mDNS

DLNA uses SSDP

Clients

Servers



MOD 13- 12

Zero-configuration networking (zeroconf) is a set of protocols that enables network service discovery for printers, speakers, screen mirroring, video surveillance camera, and more. It also offers address assignment and name resolution for desktop computers and mobile devices, in flat, single-subnet IP networks: no manual intervention or specific configuration servers required.

Many years ago Apple introduced Bonjour, its zeroconf implementation that uses mDNS (multicast DNS). Most Apple products support this protocol for discovering multiple services within the Apple eco-system, and you can install some of these services in Microsoft Windows using 3rd-party software.

Multicast Domain Name Service (mDNS) resolves host names to IP addresses within small networks that do not include a local name server. mDNS can work in conjunction with DNS Service Discovery (DNS-SD), a companion zero-configuration technique specified separately in RFC 6763. It also supports Digital Living Network Alliance (DLNA) with the Simple Service Discovery Protocol (SSDP) for Windows, Android, and several IoT devices. DLNA/SSDP are commonly found in networks that deploy Android.

Aruba AirGroup Support

Aruba AirGroup provides the following features

- Send unicast responses to mDNS queries and reduces mDNS traffic footprint
- Ensure cross-VLAN visibility and availability of AirGroup devices and services
- Allow or block AirGroup services based on user roles or VLANs



AirPlay



AirPrint



Google Cast



Amazon Fire TV



DLNA

Aruba AirGroup requires the **Foundation License**

MOD 13- 13

AirGroup is a unique enterprise-class capability that leverages zero-configuration networking and allows devices to communicate over complex access network topologies. AirGroup supports Bonjour-supported and DLNA-supported services on Apple and Android devices respectively. Apple devices constantly send mDNS packets to locate Bonjour services. Similarly, Android devices constantly send SSDP packets to locate DLNA services.

In simple networks, like a home network, discovering devices and services is easier because there is just one subnet. If a network includes a large number of client devices sending mDNS or SSDP queries, more bandwidth is consumed and therefore, the network performance is affected. In large universities and enterprise networks, it is common for devices to connect to the network across VLANs. As a result, client devices on a specific VLAN cannot discover a service that resides on another VLAN. The IP addresses in such networks are link-local scope multicast addresses and each query or advertisement can only be forwarded on its respective VLAN, but not across different VLANs. Broadcast and multicast traffic are usually filtered out from a WLAN network to preserve the airtime and battery life. This inhibits the performance of services that rely on multicast traffic.

Aruba addresses this challenge with AirGroup. AirGroup allows administrators to set policy-based discovery and enables client devices to be location-aware. Zero-configuration networking enables service discovery, address assignment, and name resolution for desktop computers, mobile devices, and network services.

AirGroup also supports DLNA, a network standard that is derived from UPnP (universal plug-and-play) for Android devices. AirGroup adds mDNS or SSDP proxy capabilities to campus WLANs, so that Bonjour or DLNA messages can be served across subnets or VLANs. DLNA uses SSDP for discovering services available on the network. DLNA provides the ability to share digital media between multimedia devices running Android or Windows operating systems. SSDP multicasts and advertises the services and queries managed by AirGroup without affecting the advertisement or discovery process of SSDP devices.

AirGroup features

AirGroup provides the following features:

Send unicast responses to mDNS queries and reduces mDNS traffic footprint.

Ensure cross-VLAN visibility and availability of AirGroup devices and services.

Allow or block AirGroup services based on user roles or VLANs.

AirGroup services

AirGroup supports discovery and management of the following services.

AirPlay: Apple AirPlay allows wireless streaming of music, video, and slide shows from your iOS device to Apple TV and other devices that support the AirPlay feature.

AirPrint: Apple AirPrint allows you to print from an iPad, iPhone, or iPod Touch directly to any AirPrint compatible printer.

Amazon Fire TV: The Amazon Fire TV allows you to stream music, video, and games to television.

Google Cast: The Google Cast service allows you to play audio or video content on a high-definition television by streaming content through Wi-Fi from the Internet or local network. Google Cast does not respond to application ID-based wildcard queries from Aruba Central.

DLNA Media: Applications such as Windows Media Player use this service to browse and play content on a remote device.

DLNA Print: This service is used by printers that support DLNA.

AirGroup limitations

AirGroup has the following limitations:

Template groups are not supported.

Wired AirGroup servers or users are not supported.

Custom-based AirGroup services are not supported.

Confidential - For Training Purposes Only

Enable AirGroup Service

Customer: HPE hqlow

Visibility **AirGroup**

Summary List Config

Campus1

Manage

- Overview
- Devices
- Clients
- Guests
- Applications**
- Security

Analyze

- Alerts & Events
- Audit Trail
- Tools
- Reports

ENABLE AIRGROUP SERVICE ☒

Services (7)

Name	Description	Status	Category	Client roles	VLANs
dial	Streaming services to launch apps	Disabled	Pre defined		
amazon_tv	Amazon TV streaming service	Disabled	Pre defined		
airprint	Printing service for Apple devices	Disabled	Pre defined		
googlecast	Streaming service from Google	Disabled	Pre defined		
dlnaprint	DLNA print service	Disabled	Pre defined		
airplay	Streaming services for apple devices	Disabled	Pre defined		
dlnamedia	DLNA media service	Disabled	Pre defined		

MOD 13- 14

To enable AirGroup, complete the following procedure:

1. In the Aruba Central app, set the context filter to a group.
2. Under Manage, navigate to Applications > AirGroup.
3. Click the Config icon.
4. Under Settings, move the AirGroup Service slider to the right.

Edit AirGroup Service

The screenshot shows the 'Edit Service' configuration page in the Aruba Central interface. At the top, there are tabs for 'Visibility', 'AirGroup', 'Summary', 'List', and 'Config'. The 'Services' tab is active. The main content area is titled 'Edit Service' and includes the following elements:

- A toggle switch for 'Override service configuration for this group' which is currently turned on.
- A 'Name' field containing 'airplay'.
- A 'Description' field containing 'Streaming services for appl...'.
- An 'Enable service' toggle switch which is currently turned on.
- A 'Servers expiration period' field set to 'minutes'.
- A section for 'Allow service' with a dropdown menu and a plus sign icon, labeled 'for these client roles:'. Below this is a note: '(Select the client roles from the selection list)'.
- A section for 'Allow service' with a dropdown menu and a plus sign icon, labeled 'on these VLANs:'. Below this is a field for 'VLANs'.

MOD 13- 15

To enable AirGroup services, complete the following procedure:

In the Aruba Central app, set the filter to a group.

Under Manage, click Applications > AirGroup.

Click the Config icon.

On the required AirGroup service, click the Edit icon under Action.

Hover over the required AirGroup service and move the slider to the right.

To disallow the AirGroup service on the required roles:

- Select the Disallow service on selected roles check box.
- Select the check-box next the required roles.

Disallowing an AirGroup service on roles is optional. If configured, the AirGroup service is not available on the specified roles.

To disallow the AirGroup service on the required VLANs, select the Disallow service on selected vlans check box and enter comma separated VLAN IDs. Disallowing an AirGroup service on VLANs is optional. If configured, the AirGroup service is not available on the specified VLANs.

Click Save.

Confidential - For Training Purposes Only




Quality of Service

Make sure your most important traffic is forwarded appropriately by verifying or configuring the Quality of Service (QoS) features of Aruba's wireless and wired products.

Confidential For Training

Quality of Service Overview

Objective Provide special network treatment to specific traffic



MOD 13-17


QoS enables networks to provide better or special service to a set of users and/or applications, perhaps to the detriment of other users and/or applications. With QoS, you are trying to make the best possible use of available bandwidth.

Different traffic types compete for the same network resources such as bandwidth and fast processing. Without QoS, packets are served as they arrive, typically first-in first out (FIFO). Your network management traffic may be more important to you than some users watching videos during their lunch break, but, by default, there is not differentiation between this traffic.

On the highway, an ambulance gets higher priority than other vehicles. Likewise, QoS is not being fair when dividing the network resources, but rather selecting certain traffic types for improved service, perhaps at the expense of other traffic. QoS does not make the road wider, it just decides who goes first and consequently who must wait.

Note: Understand that QoS is not a standalone service or product, but rather a concept that supports a framework of protocols, techniques, and network attributes that span across the infrastructure.

Queuing Operation Overview



MOD 13-18

Now that you understand what QoS is, you should understand how it works at a high-level, using a slightly simplified QoS scenario. Packets ingress at some switch or router interface: some video packets, some FTP packets, and some voice packets.

First the switch must classify packets. It can assess source/destination MAC and IP addresses, the ingress port, the application type, and more, depending on the product and your configuration. Perhaps all packets arriving on port 1/1/1 are low priority, or all VoIP packets are high priority. The purpose of classification is to identify what packets are more or less important and then mark them. Marking them involves adding information either in the layer 2 or layer 3 headers so that other networking devices downstream can use these markings to implement your QoS policies.

Suppose you arrive at a music performance. An attendant classifies you based on your identification and your entry ticket. Then you are marked with a wrist band. Certain folks are on the performer's guest list and are given a gold "all access" band. Others paid a premium for their ticket and are given a silver band: they get a special VIP line to quickly get food and beverages. Others paid a low price, are given a blue band, and must stand in a long line to get warm lemonade.

In the slide, some packets were identified as video. These packets are marked upon ingress by placing a priority value of 5 in a special frame or packet header. Other packets classified as FTP and are marked with a priority of 1. The VoIP traffic is marked with a priority of 6. 665


Notice that the classification process can also use marking. Perhaps an upstream device has already marked packets, or a VoIP phone has pre-marked its packets.

Then packets are queued up for service. Video packets are placed in the medium queue, FTP in the low queue, and VoIP in the high queue. Now the traffic scheduler takes over, often based on a configured service policy. Perhaps the policy states that the high queue shall be serviced exclusively, until it is empty. The router services other queues if and only if the high queue is empty. It goes back and forth between the medium and low queues, spending 70% more time servicing the medium queue.

Let's look at classification and marking in more detail.

Confidential - For Training Purposes Only

L2-L3 Marking



MOD 13-19

To communicate the desired priority for traffic between network infrastructure devices, the devices use QoS markings. These mechanisms define specific values, which request specific priorities for traffic marked with that value. AOS-CX switches can leverage these values to determine how to classify and service traffic. The values reside within frames or packet headers and so are passed on with the traffic from device to device.

At layer 2, Ethernet frames can include 802.1Q information used for VLAN tagging. One of the fields inside the 802.1Q header is used to do QoS tagging, as defined in the 802.1p standard. This is a 3-bit field, and so supports 8 priority values: 0—7. These fields are called the Class of Service (CoS) bits. If you use the 802.1p user priority field that is a part of the 802.1Q standard, you are said to be using CoS marking.

At layer 3, the IP header has a 8-bit type of service (ToS) field. Initially, 5 bits were unused, and so 3 bits were used for ToS with the IP Precedence RFC, which allowed 8 classifications for traffic. Many service providers needed more classifications so a new standard was created (Differentiated Services) for the ToS field, utilizing 6 bits. If using this standard, you are said to be using Differentiated Services Code Point (DSCP) markings.

Best effort service

This is the simplest service type. All traffic is treated equally in a first-come, first-served manner. If the traffic load is low in relation to the capacity of the network links, then there is no

need for the administrative complexity and costs of maintaining a more complex end-to-end policy. If you intend to use this method, you must ensure that your network is over-designed: a so-called over-provisioned network where all link speeds are much higher than the peak loads on the network.

802.1p or CoS

CoS uses a three-bit field in the 802.1Q VLAN tag to indicate priority. It defines eight values between 0 and 7. Note that because the VLAN tag carries the priority value, native VLAN traffic cannot contain a CoS value. Although this example shows an Ethernet frame that encapsulates an IP packet, an Ethernet frame with an 802.1p field can encapsulate any Layer 3 protocol an Ethernet supports.

ToS and IP precedence

ToS is rarely used today, but it is helpful to understand what ToS is if you hear it referenced. The slide shows that every IP packet includes an eight-bit field called the Type of Service (ToS) field. As originally defined, the ToS field contained two sub-field: a three-bit IP precedence field used for relative prioritization and a four-bit subfield for the specific ToS desired by an individual packet. The remaining bit, bit 7, is called the must be zero (MBZ) bit and is unused.

The three IP precedence bits in the most significant position were to be used to provide eight levels of precedence or priority. The levels are labeled 0 to 7, with 7 being the highest priority and 0 the lowest. The values 6 and 7 are reserved for Internet and network uses, respectively. This enables routing updates and other critical traffic to have a higher priority than user-generated traffic. The highest user-defined value is 5, which represents critical priority.

The four ToS bits were designed to allow applications to instruct routers to choose routing paths with one of the following four characteristics: minimized delay, maximized throughput, maximized reliability, and minimized cost. Only one of the four bits could be “turned on,” that is, set to 1.


DiffServ / DSCP

Differentiated Services (DiffServ), defined in RFC 2474, supersedes a legacy IP QoS protocol called Type of Service (ToS). In DiffServ, the first six bits of the eight-bit ToS field define 64 DiffServ code points (DSCPs). Instead of only defining priority relative to each other, the 64 DSCP values are also intended

to define distinct forwarding behaviors, or per-hop behaviors (PHBs).

Confidential - For Training Purposes Only

Layer 2 CoS Classification



802.1p	Traffic Type (802.1Q)	Traffic Type (802.1d)
7	Network control	Network control
6	Internetwork control	Voice (<10ms latency)
5	Voice (<10ms latency)	Video (<100ms latency)
4	Video (<100ms latency)	Controlled Load
3	Critical applications	Excellent Effort
2	Excellent effort	Spare (less than 0)
0	Best effort	Best effort
1	Background	Background



MOD 13-20

CoS values request up to eight different priority levels for different types of traffic. Use the 3-bit Priority Code Point (PCP) 802.1p field within the 16-bit 802.1Q Ethernet VLAN tag to mark the desired CoS. The standard recommendation differs based on whether the traffic type is 802.1Q or 802.1D.

For 802.1Q traffic, the you should use the most up-to-date recommendations in the IEEE 802.1Q-2014 standard. This standard recommends that you reserve the two highest values (7 and 6) for network control and internetwork control traffic. Network control traffic includes frames for layer 2 control protocols such as spanning tree BPDUs while internetwork control includes layer 3 control frames such as OSPF hellos and LSUs.

AOS-CX switches automatically mark their control traffic for high priority. Voice, value 5, and video, value 4, are given the next highest priority. These recommendations are based on their low latency requirements. Mission critical applications receive the next highest priority, value 3. The standard recommends assigning Value 2 to traffic that requires excellent effort. It is up to you to define what is considered critical or excellent effort applications. Next the switch gives its best effort to traffic with the default priority, which is 0. Finally, 802.1p value 1 is assigned to background traffic, which is less important or can tolerate delay. CoS 1 is deliberately set as the lowest CoS: this enables a traffic service level below the default (best effort) traffic level to be specified.


802.1D recommendations are based on the original standard that defines CoS. Some

network infrastructures still use these recommendations, so you should be aware of them. In this scheme, network control traffic has just one value, 7 which shifts voice and video traffic up one priority level. Controlled load (rather than critical) has the next highest level and then excellent effort. But again, companies can choose themselves what applications fit in these categories. One of the most crucial differences in the 802.1D standard is that it assigns value 2 to "spare" traffic and is intended for traffic with a priority less than best effort. It is crucial that you understand whether other devices in the network consider value 2 as higher than 0 or less than 0.

Note: The you see in standards, vendors, recommendations, and what companies use can be different and confusing. What matters is that you use a consistent policy for how you will mark and prioritize traffic to solve your company's traffic needs.

Confidential - For Training Purposes Only


L3 Classification: Usage of ToS Bits



You know that every IP header includes an 8-bit ToS field. Originally, only the high-order 3 bits were used for IP Precedence. With three bits you have 8 possible values: decimal 0—7. Much later, this 8-bit ToS field was redefined in a set of RFC standards to use 6 bits. Of those 6 bits, for nearly all common implementations, the lowest order bit is not used.

Note: Newer specifications can make specialized use of formerly unused bits, but they are not relevant to this discussion of QoS.

L3 Classification: Interpret DSCP as a Decimal Number



Consider the decimal number 9,325. You know that the 5 is in the one's column in a decimal format, the 2 is in the 10's column, 3 is in the hundreds column, and 9 is in the thousands column.

With the 6-bit DSCP binary number, the columns are 1, 2, 4, 8, 16, and 32, as shown in the slide. So, the binary number 001010 is 10 in decimal: there is a 1 in the eight's column, and a 1 in the 2's column, and $8 + 2 = 10$. Similarly, 001100 in binary is 12 in decimal, and 001110 is 14 in decimal.

So, when communicating with others you can say, "I have created a policy to mark certain packets with a DSCP decimal value of 10 (or 12, or 14, etc.); but there is a second way to interpret these bits.

Using L2 and L3 Classification

VLAN Protocol ID	802.1p priority	1	VLAN-ID

Version	ToS Field	Rest of IPv4 header

802.1p	Traffic Type (802.1Q)
7	Network control
6	Internetwork control
5	Voice (<10ms latency)
4	Video (<100ms latency)
3	Critical applications
2	Excellent effort
0	Best effort
1	Background

DSCP	Name	Examples
56	CS6	Network control
46	EF	Telephony (VoIP)
40	CS5	Signaling
34,36,38	AF41 / AF42 / AF43	Multimedia conferencing
32	CS4	Real-time interactive
26,28,30	AF31 / AF32 / AF33	Multimedia streaming
24	CS3	Broadcast video
18,20,22	AF21 / AF22 / AF23	Low-latency data
16	CS2	OAM
00	CS0 / BE / DF	Best effort
10,12,14	AF11 / AF12 / AF13	Bulk data
08	CS1	Low-priority data




MOD 13- 23

The tables in the slide show default markings or classifications for the CoS (left side of the slide) and DSCP values (right side of the slide).

The right table in the slide recommends the type of applications that you might assign to each AF class. For example, you could assign multimedia conferencing applications to AF class 4. These applications are interactive, so they need very low latency. You could assign multimedia streaming, which can use buffering and has a bit less strict latency requirements to AF class 3. You can assign AF class 2 to applications that require low latency, but less than multimedia. For example, interactive network management applications such as SSH could use this class. Finally, AF class 1 is for bulk data that can tolerate delays: it has priority less than 0.

If needed, then you can use the CS values for legacy IP precedence values. Also, RFC 3246 defines Expedited Forwarding (EF) class 46, to be used for applications more sensitive to jitter and delay. Aruba often recommends this DSCP for voice over IP (VoIP) traffic.

QoS Basics



MOD 13- 24

This slide illustrates the basic steps that traffic experiences when you implement a QoS solution.

A device sends a frame, with or without classification markings.

The switch receives the frame and places it in the ingress queue.


The switch determines if the ingress packet was pre-marked by the host application. And if so, does the switch trust these markings. If there are not trusted markings, the switch can mark or re-mark the frame based on some policy you define. If ingress marking is trusted, the switch can use those markings for any policy you may have defined. If you expect both L3 DSCP and L2 802.1p markings, you should define a policy to control which marking to trust. You might also define an ingress rate limiting policy, applied certain ports.

The switch then uses the classification information and virtual output queuing (VOQ) to determine when to move the frame to the egress queue.

The switch then implements the egress queuing policy, based on the traffic classification, to determine when the frame is sent out the port, along with applying any defined outbound rate limiter.

Let's look at this five-step process in a bit more detail.

QoS Processing on AOS-CX Switches



MOD 13- 25

ArubaOS-CX switches use a special process to locally apply QoS features. This process has five stages:

Ingress. Packets that arrive at a switch interface could be filtered by the rate limit QoS feature. This QoS feature can limit broadcast, multicast and unknown unicast (BUM) traffic, to protect the switch from unnecessary processing.

Prioritization. The switch applies the QoS classifier tool to determine the correct priority for a specific traffic type. ArubaOS-CX defines two values: Local Priority and Color. The Local Priority helps to determine the correct egress priority queue for the traffic. The color determines the traffic's drop eligibility. If queues become congested, more drop-eligible traffic will be dropped first.

Destination Determination. The egress interface for each packet is determined based on the switching and routing logic. Internally the switch moves the traffic to the correct egress interface. QoS tools in this stage are not used and therefore do not affect this stage.

Queuing. The switch assigns each egress packet to an interface queue, based on the packet's Local Priority value. ArubaOS-CX supports up to 8 queues. Packets wait in the queue for their turn to be transmitted.

Transmission. In this stage a scheduler QoS tool defines the order in which packets from

different queues are transmitted out the egress interface.

ArubaOS-CX switches use an intra-switch queuing method called Virtual Output Queuing (VoQ). If the ingress buffer used a single queue, head of line (HOL) blocking could delay the traffic. If the packet at the front of the queue is destined out a congested port, it delays all packets behind it, even though those that are destined to non-congested ports.

VoQ prevents this problem by providing deep ingress buffers with separate queues for each egress port. Physically the ingress buffer consists of internal and external DRAM banks with most traffic typically being buffered in the internal DRAM. This means that packets will stay in the inbound port queue when the outbound queue is under load. Using this technique packets are never drop in the outbound interface but in the inbound. This feature is embedded in the Operating System and is not configurable.


Let's dive into each one of these processes.

Confidential - For Training Purposes Only

Ingress Stage: Rate Limiting

Objective Control inbound traffic, prevent flooding and traffic storms

Apply to BUM and ICMP traffic



Configure per physical or LAG interface



MOD 13- 26

Port rate limiting helps control undesirable traffic. The goal is to allow only enough BUM traffic for the network to function properly, while preventing flooding and traffic storms. A certain amount of BUM traffic is required for normal network operation. Broadcast packets may include ARP and DHCP traffic, for instance. Video streams, and certain types of network protocol packets, are multicasts. Unknown-unicast packets may be intended for devices whose addresses have temporarily aged out of network-forwarding caches. Configuring rate limits can help provide the balance between necessary traffic and flooded traffic.

Rate limiting is a traffic policer QoS tool. It ensures that traffic conforms to a defined rate: a bandwidth limit. An important characteristic of rate limiting is that it does not add delay to traffic to be forwarded: that which is under the threshold and is accepted by the switch. You configure it per physical or LAG interface.


Configure Rate Limiting

Agg-1 configuration

```
interface 1/1/7
description TO_CLOUD_SERVICE
rate-limit broadcast 500 pps
rate-limit multicast 500 pps
exit
interface 1/1/2
description TO_SERVER
rate-limit broadcast 100 pps
rate-limit unknown-unicast 50 pps
```

Acc-1 configuration

```
interface 1/1/1
description TO_COMPUTER
rate-limit broadcast 50 pps
rate-limit multicast 50 pps
rate-limit icmp ipv4 43 pps
exit
interface 1/1/25
description TO_AGG-1
rate-limit broadcast 500 pps
```



Different platforms may have different rate limit sets - refer to the latest user guides

You can implement the rate limiting QoS feature on all ArubaOS-CX platforms; however, not all configurations are supported on all AOS-CX switches:

Broadcast: Supported on all ArubaOS-CX platforms, limit only the broadcast traffic.


Multicast: Supported on all ArubaOS-CX platforms, limit only the multicast traffic.

Unknown Unicast: Supported only on 8400, 8325 and 8320 switch platforms, limit only the unknown unicast traffic.

ICMP: Supported only on 6400 and 6300 switch platforms, the ICMP rate limit can be configured to apply to IPv4, IPv6, or all IP traffic. Only one ICMP rate-limit can be configured at a time. Applying a new ICMP rate-limit replaces any previous ICMP rate-limit.

Note: When multicast and broadcast rate limits are applied at the same time on the same interface, broadcast packets are limited to the lower of the two rate values.

Classification and Marking



MOD 13- 29

Recall that classifiers perform a kind of identification process. They inspect incoming traffic, identify it, and then assign it to a class. You define a classifier-based policy, analyzes packet parameters to apply a specific class of service.

Notice that the classification process can also use marking, as shown in the slide. Perhaps an upstream device has already marked packets, or a VoIP phone has pre-marked its packets. You can configure certain switch ports to trust these markings, and take some action based on them: provide priority service levels perhaps. You can configure other ports to ignore ingress frame markings, and re-mark them according to a local policy that that you define. You're in charge of how you want to handle this process!

Prioritization Stage

Objective

- Determine **internal Local Priority** and **color (drop eligibility)** for each ingress packet
- Based on incoming marks (DSCP or CoS) or Classifier-based Policy

Trust (Ingress Interface)	Custom Mapping	Classifier-based Policy	Local Priority assignment
Yes	No	No	Default mapping
Yes	Yes	No	Custom mapping
Yes	Yes	Yes	Classifier-based policy mapping
None	No	No	Default (LP = 1 – best effort)
None	Yes	No	Default (LP = 1 – best effort)
None	Yes	Yes	Classifier-based policy mapping




MOD 13- 30

The objective of the prioritization stage is to determine the Internal Local priority value and the Color (drop eligibility) for a specific traffic. The AOS-CX switch uses these parameters to specify the behavior, or the class of service used.

The decision to assign a Local Priority and color value could be based on ingress packet L3 DSCP or L2 CoS marks. Each switch maintains a mapping table that helps with the decision. You can display the default mapping table with the `show qos dscp-map default` command or `show qos cos-map default` command. AOS-CX supports the creation of custom mapping as well. The other possibility is to use a classifier-based policy, where the interesting traffic is selected based on specific characteristics (like source, destinations, and/or applications used). In this case the marking may or may not be present.

Trust is an important setting for QoS in ArubaOS-CX. When trust is set on the incoming interface, the switch can use the incoming marks to determine the Local Priority. However, if the trust is not set (or set to none) then incoming marking cannot be used and all traffic is placed in the default best effort behavior, represented by the Local Priority 1. Notice that classify policy-based assignment is not dependent on the trust status and overrides default and custom mapping. The table in the slide summarizes this discussion.

Trust Overview



MOD 13- 31

When traffic arrives at the switch, the switch can trust existing markings or not. The term trust refers to whether the packet markings are considered valid input for the classifier, which decides each packet's service class.

You know that trust configuration plays an important role in assigning the Local Priority and color values. In AOS-CX, you can configure trust at the global or interface level. When the trust level is set at both levels, the interface level configuration takes precedence.

AOS-CX supports three possible configurations: trust CoS, trust DSCP and trust none. The first two options act as the name suggests: the switch accepts the incoming marking and uses it to determine class of service. In AOS-CX 10.5 an interface can only support one trust setting at the time, meaning that you can set DSCP or CoS but not both at the same time, on the same interface.

It is not always recommended to trust ingress markings, especially if traffic is received from non-trusted devices. In these situations, you can protect the network by ignoring the incoming marking setting and setting the trust to none; and then define your own classification (you'll learn about how to do this later in the module).


Global and Interface Trust

Trust global configuration

```
qos trust <cos | dscp | none>
```

Trust interface configuration

```
interface <interface-id>  
  qos trust <cos | dscp | none>
```



MOD 13- 32

If a port has no setting, the port takes its setting from the global setting. Otherwise, the port setting, including none, overrides the global setting. In this way, you could globally not trust incoming marks, but trust DSCP on certain ports. You could globally trust DSCP, but use CoS instead on certain ports, and so on.

By default, the global setting is none, and all ports have no trust setting. This means that all ports use the global setting of none: the switch ignores any markings in ingress frames and assigns all traffic the LP and color associated with CoS 0 in the CoS map; in other words, all traffic would be treated the same. By default, this is LP=0 and color=green. The switch preserves any incoming DSCP when it forwards the traffic, as well as an incoming CoS value if the forwarded traffic is tagged. However, these values have no effect on how the switch queues the traffic.

RFC 2697 defines the use of colors with DSCP codes. A single rate three color marker (srTCM) meters an IP packet stream and marks its packets either green, yellow, or red. Marking is based on a committed information rate (CIR) and two associated burst sizes, a committed burst size (CBS) and an excess burst size (EBS). A packet is marked green if it doesn't exceed the CBS, yellow if it does exceed the CBS, but not the EBS, and red otherwise. The srTCM is useful, for example, for ingress policing of a service, where only the length, not the peak rate, of the burst determines service eligibility. See RFC 2697 for more information: <https://www.ietf.org/rfc/rfc2697.txt>. The use of colors is optional.

Note. It is important to clarify that a port not having a trust configuration is different from the port having a configuration of none. If a port has no setting, the port takes its setting from the global setting. Otherwise, the port setting, including none, overrides the global setting.

You can configure the prioritization stage as shown in the slide. The custom map has two options:

Use the `cos-map` option to define the local priority assigned to incoming packets for a specific 802.1 VLAN priority code point (CoS) value. The CoS map values are used to mark incoming packets when QoS trust mode is set to `cos`.

Use the `dscp-map` option to defines the local priority assigned to incoming packets for a specific IP differentiated services code point (DSCP) value. The DSCP map values are used to prioritize incoming packets when QoS trust mode is set to `dscp`.


Configuring custom map involves these switch commands:

```
qos cos-map <CoS> local-priority <local-priority> color <green | yellow | red>
```

```
qos dscp-map <DSCP> local-priority <local-priority> color <green | yellow | red>
```

Confidential - For Training Purposes Only


Trust Set to DSCP



MOD 13- 33

Many networks classify traffic at the edge, so it can be useful for the AOS-CX switch to use the already applied classifications to determine the correct priority for the traffic. If other network infrastructure devices are using DSCP to classify traffic, set the global or port trust mode to dscp. When a port has a DSCP trust setting, whether that is set on the port or taken from the global setting, the switch looks at the incoming DSCP in traffic's IP headers and uses its DSCP map to assign the traffic an LP and color.

Default DSCP Map



Default DSCP Map


DSCP	LP	Color		DSCP	LP	Color		DSCP	LP	Color	Name
Name				Name							
0	1	green	CS0	19	2	green		34	4	green	AF41
1-7	1	green		20	2	yellow	AF22	35	4	green	
8	0	green	CS1	21	2	green		36	4	yellow	AF42
9	0	green		22	2	red	AF23	37	4	green	
10	0	green	AF11	23	2	green		38	4	red	AF43
11	0	green		24	3	green	CS3	39	4	green	
12	0	yellow	AF12	25	3	green		40	5	green	CS5
13	0	green		26	3	green	AF31	41-45	5	green	
14	0	red	AF13	27	3	green		46	5	green	EF
15	0	green		30	3	red	AF33	47	5	green	
16	2	green	CS2	31	3	green		48	6	green	CS6
17	2	green		32	4	green	CS4	49-55	6	green	
18	2	green	AF21	33	4	green		57	7	green	CS7
								58-63	7	green	

MOD 13- 34

The slide shows the default DSCP map for AOS-CX switches, which follows the standards and works in most circumstances. In this example, the map has an entry for DSCP 46, which assigns the incoming traffic LP 5 and color green. However, you can customize this map, setting your own local priority, color, and name for each DSCP. Your goal in customizing the map should be to match the system the devices use from where the switch is receiving the incoming traffic. For example, you may want to change DSCP 46, used for VoIP traffic, to LP 7 if the company wants to give this traffic highest priority.

Note that, if the incoming traffic has a CoS value as well as a DSCP, the switch preserves the CoS value in the outgoing frame, if tagged. However, the CoS value and map have no effect on the traffic's LP and color. Those derive from the DSCP only and it is this value that affects the queuing and scheduling of this traffic.

Classifier-Based Configuration



MOD 13- 35

You can manually override or define what should be used for the marking. This can be done globally or per-interface. For an interface, you can remark all traffic with a specific marking, like this:

```
Switch(config-if)# qos trust none
```

```
Switch(config-if)# qos {dscp | cos} <value>
```

This work re-mark all incoming traffic with the specified DSCP or CoS value.

The problem with remarking all traffic the same is that you might have different devices and/or applications that need to be treated differently. If that is the case, you would implement a classifier-based policy. The example in the slide shows how you can use a classifier-based policy configuration to assign the Local Priority. This policy can be applied to an interface or globally. Typically you would use the global application, since you probably want a consistent method that is always used across the entire company.

Validate Classifier-Based Configuration

```
Switch# show policy interface 1/1/1 in  
Direction
```

Sequence	Name	Additional Policy Parameters	Comment	Class Type	action

Inbound

	POLICY-VOICE				
10					
	VOICE	ipv4			
		local-priority	5		

Policy applied

```
ACC-1# show policy hitcounts POLICY-VOICE
```

```
Statistics for Policy POLICY-VOICE:
```

```
Interface 1/1/1* (in):
```

Hit Count	Configuration

10 class ip VOICE action local-priority 5	

402	10 match udp any any eq 5060 count

402 Hits



MOD 13- 36

The example in the slide shows how to validate the classifier-based policy configuration from the previous slide. Use the show policy interface <interface-id> in command to verify if the policy was correctly set on the interface. The command show policy hitcounts <Policy-name> displays the number of hits counts that the policy received.

What is Wi-Fi Multimedia

WMM

- Covers QoS over Wi-Fi networks
- Four categories: voice, video, best effort, background
- Different treatments according to the categories

WMM prerequisites

- AP and controller are Wi-Fi CERTIFIED™ for WMM, which is enabled
- Client must be Wi-Fi certified™ for WMM and configured for WMM marking




MOD 13- 37



Now that you have the basics for QoS on the AOS-CX switches on the wired side, let's talk about wireless QoS. Wi-Fi Multimedia (WMM) is used to implement QoS in the Wi-Fi networks. It defines four QoS categories: voice, video, best effort, and background. Wireless networks and products use WMM to prioritize traffic, so these category applications have access to the necessary network resources. All Aruba Wi-Fi wireless products support WMM.

How WMM works



Higher priority access category has a lower transmission time interval: transmission opportunity (TXOP)

MOD 13- 38

DSCP is most typically responsible for classifying packets in the wired network, while WMM is responsible for WiFi QoS classification and implementation. WMM maps traffic to four access categories:

AC_BK: background traffic

AC_BE: best effort traffic

AC_VI: video traffic





AC_VO: voice traffic

These access categories map to the four queues a WMM device requires.

A TXOP is a bounded time interval during which a station can send as many frames as possible. The higher priority access category has the lower transmission opportunity (TXOP) assigned, and it can transfer more data frames in the same amount of time.

Default WMM – DSCP Mapping

DSCP	WMM
0 – 7	Best Effort (0)
8 – 15	Background (1)
16 – 23	Background (2)
24 – 31	Best Effort (3)
32 – 39	Video (4)
40 – 47	Video (5)
48 – 55	Voice (6)
56 – 63	Voice (7)



Problem scenario

Voice traffic

MOD 13- 39

In the slide, you see default the DSCP-to-WMM mappings used by Aruba wireless devices. The problem with this default mapping is that DSCP value 46, which is often used to mark VoIP traffic, is by default mapped to the WMM video traffic class.

WMM-DSCP Mapping

The screenshot shows a configuration page for 'WiFi Multimedia' under the 'WLANs' tab. It features a table with two columns: 'Share' and 'DSCP Mapping'. The rows represent different WMM categories: Background, Best Effort, Video, and Voice. The 'Voice' row is highlighted, showing a 'DSCP Mapping' value of 46.

	Share	DSCP Mapping
Background WiFi Multimedia Share:	0 %	
Best Effort WiFi Multimedia Share:	0 %	
Video WiFi Multimedia Share:	0 %	
Voice WiFi Multimedia Share:	0 %	46



MOD 13-40

Effective QoS must be established end-to-end, which means that both Wi-Fi and wired networks must prioritize traffic for QoS. Therefore, you could fix this in one of two ways:

- Fix the mapping issue on the ingress switches
- Fix the mapping issue on the AP

You already know how to do it on the switches. Optionally you could define a mapping mechanism on the Aps that would translate the WMM value to the correct DSCP value and vice versa. The IEEE 802.11e standard defines the mappings between the WMM access category (AC) and DSCP tags, which is why DSCP 46 was designated as “video”, even though most companies use it as voice.

The DSCP WMM AC mapping allows you to customize the mapping between WMM ACs and DSCP tags to prioritize various traffic types. You can apply and configure WMM AC mappings to a WMM-enabled SSID profile. The system will use the default DSCP WMM AC mapping if you haven't configured the DSCP mapping.


You can fix the problem stated earlier by modifying WMM to DSCP mapping. Access a group context. In our example, define the new mapping for WMM voice: DSCP 46. This will guarantee that VoIP traffic marked with DSCP 46 will be forwarded as voice traffic by the AP.

End-to-End QoS

Trust client VOIP QoS: ☒

Firewall setting on gateways

QoS is effective only when implemented *end-to-end*



MOD 13-41

End-to-end QoS requires that all infrastructure devices honor priority values. To make sure wired switches honor DSCP simply enable DSCP trust on every port that connects to an AP, switch, and gateway. This way priorities will be respected by all infrastructure devices in the network.

UCC



You can accommodate VoIP solutions with Aruba's Unified Communication and Collaboration (UCC) service. You'll learn more about Aruba's UCC in this section.

Confidential / For Training

UCC Service


Heuristics based solution

VoIP protocol support

- Skype for Business
- Wi-Fi calling
- SIP

Call visibility, monitoring, troubleshooting

Can modify DSCP values for voice, video



Requires Advanced AP license

MOD 13- 43

When a wireless client connects to a WLAN, it competes with other clients, and even with the AP, to gain access to the media and send its data. This competition is not always fair, and in some situations, clients wait a longer time to gain access to the media or/and the amount of access is not often equal to other devices. This unfair behavior could introduce communication problems, especially to real time applications. And this could lead to a poor user experience.

The Aruba UCC service addresses these problems by L3-marking the real time traffic, including voice and video. This service also provides call visibility, monitoring, and troubleshooting of UCC traffic.

UCC in AOS 10 supports the following VoIP protocols:


Skype for Business

Wi-Fi - Calling

SIP (generic call)

Note: UCC requires of the Advanced AP license in Central.


UCC Workflow for Bridge Mode



The following steps describe the workflow of UCC for an AP-only deployment:

1. The user starts a Skype for Business call.
2. The AP, which runs DPI, identifies a new UCC session.
3. The AP sends the call metadata using the OpenFlow protocol (the AP runs an OpenFlow agent) to an OpenFlow controller. In AOS 10, this is Aruba Central.
4. The OpenFlow controller service in Central sends the metadata to the internal UCC service. This runs the proper Application Layer Gateway (ALG) to decide if the traffic should be prioritized; and the result is sent back to the OpenFlow controller service.
5. The OpenFlow controller in Central instructs the AP to mark the traffic using the proper DSCP value, assuming it should be prioritized.

UCC Workflow for Tunnel Mode



MOD 13-45

The following steps describe the workflow of UCC for a gateway deployment:

1. The user starts a Skype for Business call.
2. The gateway, which runs DPI, identifies a new UCC session.
3. The gateway sends the call metadata using OpenFlow protocol (the Gateway runs an OpenFlow agent) to Aruba Central.
4. The OpenFlow controller service in Central sends the metadata to the UCC service; this runs the proper Application Layer Gateway (ALG) to decide if traffic is prioritized. The result is sent back to the OpenFlow controller service.
5. The OpenFlow controller in Central instructs the Gateway to mark the traffic using the proper DSCP value, if necessary.

Key Considerations

General

- Live call status is not visible only 2 mins after call ends
- Heuristics based solution

Support

- UCC does not currently support UBT wired users
- Skype for Business SDN API is not supported
- No WLAN & end-to-end quality metrics for Wi-Fi calling and video calls

Cloud survivability

- Loss of dashboard data
- AP and gateway heuristics method mark VoIP traffic with DSCP 48



MOD 13-46

There are some considerations that you need to be aware of concerning the UCC methods discussed on the last two slides:

1. A live call status is not visible in the UCC dashboard: you need to wait for a couple minutes after the call has ended.
2. Heuristic detection method is the only one available in the Aruba cloud UCC service.
3. UCC does not support UBT wired clients.
4. Skype for Business SDN API is not supported in the Aruba cloud UCC.
5. WLAN and end-to-end quality metrics are not available for Wi-Fi calling and any video call: only Wi-Fi metrics are displayed.

In case Aruba central is not available, the AP and gateway heuristic detection method will still work (using the processing on the AP and/or gateway) and this mechanism will mark VoIP traffic with DSCP 48.

UCC Configuration

Global Visibility Saas Express **UCC** List Summary Config

Manage

- Overview
- Devices
- Clients
- Guests
- Applications**
- Security
- Network Services

Analyze

- Alerts & Events

SETTINGS

ACTIVATE UCC ☒

RETAIN CLIENT QOS ☐

PROTOCOLS (4)

NAME	VOICE	VIDEO	DESKTOP S...	SERVER	ACTION
Facetime	-	34	-	-	
SIP	46	34	-	-	
Skype for Busi...	46	34	34	-	
Wi-Fi Calling	46	-	-	-	



MOD 13- 47

To activate UCC, set the context filter to Global and then navigate to Applications > Config> UCC. Enable UCC service. Notice that there are some pre-defined DSCP values for different protocols and type of traffic: if you want to change the value click on the pencil under the Action column.

When Retain Client QoS is enabled, all configured ALGs are disabled, all traffic is treated as real-time transport protocol (RTP) traffic, default prioritization defined by the client is honored, and call visibility is provided.




LLDP MED

You can use LLDP to validate network device connectivity, document your network, and accommodate media endpoint devices. This section introduces the use of LLDP and LLDP-MED for use in QoS.

Note: CDP could also be used, but the cover of this topic is beyond the scope of this course.

LLDP and LLDP-MED



```
sw-edge1 (config) # vlan 12  
sw-edge1 (config-vlan-12) # voice
```

MOD 13-49

The Link Layer Discovery Protocol (LLDP) is an open-standard link layer protocol used by network devices for advertising their identity, capabilities, and neighbors on a local area network based on IEEE 802 technology, primarily Ethernet. It can be used to verify layer-2 connectivity between devices as well as learning basic information about directly-connected devices.

The Media Endpoint Discovery (MED) is an enhancement of LLDP, known as LLDP-MED. It is commonly used to provision devices, like determine the correct power settings for a connected VoIP phone or AP, or automatically assigning the correct VLAN and QoS settings for a connected device. It can enable plug-and-play networking. Features include the following:

Set LAN policies based on the connected device: VLAN, layer 2 802.1p CoS values, layer 3 DSCP settings, and more.

Device location discovery allows the creation of location databases: VoIP and enhanced 911 services.

Extended and automated power management of power-over-Ethernet (PoE) end devices.

Inventory management: Track your network devices, and determine device-specific info - manufacturer, software and hardware versions, serial or asset number.

Support for the fast-start capability

You know best practice is to place VoIP phones in their own VLAN and configure a high priority for that VLAN (typically a tagged VLAN). You connect a PC via an Ethernet cable to the phone. The PC sends its frames on a native, untagged VLAN. But setting up the switch port to receive the tagged traffic is not enough. The phone must also know what VLAN ID to use for its packets. You can manually configure that on each phone, but this would be a tedious process.

With LLDP-MED, the switch automatically sends the correct VLAN ID to the IP phone, along with 802.1p and DSCP settings. With LLDP-MED, you get automatic, plug-and-play provisioning for LLDP-MED capable VoIP phones. The switch can also communicate other information to the phone, including its physical location. This information can help to support emergency call services (ECS).

When phones receive PoE from the switch, LLDP-MED can help the switch allocate and deliver exactly the power that the phone needs.

LLDP-MED can also help in tracking and troubleshooting VoIP phones through a central management platform. SNMP servers can read detailed VoIP endpoint data inventory, PoE status, and information to help troubleshoot the IP telephony network and call quality issues.

Review interoperation with authentication

You manually configure switch ports with the tagged VLAN for VoIP, or rely on dynamic assignment using 802.1X and RADIUS attributes.

You manually configure each IP phone with a matching tagged VLAN, or rely on LLDP-MED to automatically configure it for you.

You can use 802.1X with LLDP-MED. Simply set up 802.1X on the port. Make sure that IP phones support 802.1X and that the RADIUS server has policies and accounts for authenticating them. If you want the server to send the tagged VLAN assignment dynamically, make sure that it has the proper policies for doing so.

For this scenario, you probably want the computer and IP phone to authenticate separately so that an unauthorized user cannot piggyback on the IP phone's session. Make sure to set the 802.1X client-limit to 2 so that the port operates in user-mode and authenticates each device separately. If you

want to enforce 802.1X but the IP phones do not support it, combine 802.1X for users and MAC-auth for IP phones.

Voice VLANs

To create a voice VLAN, configure the voice command in the VLAN context, like this:

```
Switch(config)# vlan <VLAN-ID>
```

```
Switch(config-vlan)# voice
```

LLDP-MED extensions support voice VLANs.

Confidential - For Training Purposes Only

LLDP-MED Configuration

Advertise MED TLVs

```
interface 1/1/1
  lldp receive
  lldp transmit
  lldp med capability
  lldp med network-policy
  lldp med poe [priority-override]
```

Advertise MED location TLV

```
interface 1/1/1
  lldp med-location civic-addr <country-code>
  <CA type><CA value>[<CA type><CA value>]
```

Specify an ELIN

```
interface 1/1/1
  lldp med-location <octet>
```



MOD 13- 50

LLDP is enabled by default, both globally and on interfaces. Configure LLDP if needed using the syntax shown in the figure. Use the `show lldp tv` command to determine the attributes included and/or supported via LLDP:

Switch# show lldp tv

TLVs Advertised

=====

Management Address

Port Description

Port VLAN-ID

System Capabilities

System Description

System Name

OUI

You should look for the following TLVs:

MED Capability: The capabilities TLV enables the switch to detect LLDP-MED capable endpoints, their class, and their capabilities.

MED Network Policy: The switch communicates the proper VLAN ID and QoS settings in this TLV. You will learn how to configure those on the next page.

MED Power (PoE): The switch advertises its capabilities as a power sourcing equipment (PSE). The VoIP phone can also use this TLV to request a specific amount of power after it first powers up or dynamically throughout the connection. Note that PoE+ defines its own LLDP TLV for negotiating the power (a dot3 TLV called poeplus_config), and the phone might use this TLV instead. In either case, the switch will allocate the requested amount of power and reserve it for this port even if the phone is not drawing the full power at the moment.

If you need to enable any of these TLVs, use these commands:

```
Switch(config)# interface <interface-ID>
```

```
Switch(config-if)# lldp med capability
```

```
Switch(config-if)# lldp med network-policy
```

```
Switch(config-if)# lldp med poe [priority-override]
```

Here's an explanation of these commands:

capability specifies advertisement of supported LLDP MED TLVs. The capability TLV is always sent with other MED TLVs, therefore it cannot be disabled when other MED TLVs are enabled. Default: enabled.

network-policy lets endpoints and network devices advertise their VLAN IDs, and IEEE 802.1p (PCP and DSCP) values for voice applications. This TLV is only sent when a voice VLAN policy is present. It is enabled, by default.

The **priority-override** option overrides user-configured port priority for Power over Ethernet. When both **lldp dot3 poe** and **lldp med poe** are enabled, the **lldp dot3 poe3** setting takes precedence.

Use the **lldp select-tlv** command to change what TLVs the switch will send and receive. See the Fundamentals Guide for more information on this command. The **show lldp tlv** command shows the LLDP TLVs that are configured for send

and receive:

MED location TLV

It is best practice to populate the location ID TLV for each switch port that might connect to a VoIP phone with the correct location information. For example, the command for specifying the civic address of a media endpoint such as a phone is:

```
Switch(config)# interface <interface-ID>
```

```
Switch(config-if)# lldp med-location civic-addr <country-code> <CA type> <CA value> [<CA type> <CA value>]
```

A CA type specifies the type of location being indicated, such as 3 for a city. You can specify many CA types and values in order to program the city, street, address, and building, for example. For a list of CA types, see your switch Management and Configuration Guide.

You can alternatively specify an Emergency Line Identification Number (ELIN), which is used in North America, if such numbers have been assigned to your phones:

```
Switch(config)# interface <interface-ID>
```

```
Switch(config-if)# lldp med-location <octet>
```



Knowledge Check

Self-check on key learning points



It's time for a knowledge check to review what you have learned.

Confidential - For Training Purposes Only

Question #1

You need to optimize AP channel and power settings in your RF environment. Which solution should you choose?

- a. Aruba ClientMatch
- b. Aruba AirMatch
- c. Mobility gateway clustering
- d. Aruba broadcast and multicast optimization
- e. Aruba AirGroup



Knowledge Check

MOD 13- 52



b

Confidential - For Training Purposes Only

Question #2

You need to balance the load between 2.4 and 5 GHz radio bands and between APs. Which solution should you choose?

- a. Aruba ClientMatch
- b. Aruba AirMatch
- c. Mobility gateway clustering
- d. Aruba broadcast and multicast optimization
- e. Aruba AirGroup



Knowledge Check

MOD 13- 53



a

Confidential - For Training Purposes Only

Question #3

You need to optimize the operation of zero configuration networking services in a large campus environment. Which solution should you choose?

- a. Aruba ClientMatch
- b. Aruba AirMatch
- c. Mobility gateway clustering
- d. Aruba broadcast and multicast optimization
- e. Aruba AirGroup



Knowledge Check

MOD 13- 54




e

Confidential - For Training Purposes Only

Question #4

Match each VSX component with its appropriate description.



Knowledge Check ✓



MOD 13- 55



Classify	identify traffic based on some criteria
Mark	tag packets with a priority value
Queue	organize packets into queues
Service	treat packets based on policy

Question #5

You can use ingress rate limiting to control broadcast, unknown, and multicast (BUM) traffic, but it does add some delay to the forwarded traffic.

- True
- False



Knowledge Check

MOD 13- 56



false

Confidential - For Training Purposes Only

Question #6

To get the most out of your QoS deployment, it is best to classify and mark traffic as soon as possible.

- True
- False



Knowledge Check

MOD 13- 57



true

Confidential - For Training Purposes Only

Question #7

You are concerned about tech-savvy end users that might send packets with a modified DSCP marking, in order to have their packets treated with a higher priority. What should you do to mitigate this issue?

- a. Configure an access list on the access switch
- b. Modify the default DSCP map on the access switch
- c. Use a classifier-based configuration on the access switch
- d. Change the DSCP trust setting on the access switch



Knowledge Check

MOD 13-58



d

Confidential - For Training Purposes Only

Question #8

One advantage of Aruba's Unified Communication and Collaboration (UCC) service is that you can monitor call status live, as the call is happening.

- True
- False



Knowledge Check

MOD 13- 59



false

Confidential - For Training Purposes Only

Question #9

One advantage of LLDP is that it supports LAN policies based on connected devices without any enhancements or extensions.

- True
- False



Knowledge Check

MOD 13-60



false

Confidential - For Training Purposes Only



Lab Activity

Let's review the lab activity that is associated with this module.

Confidential - For Training

Lab Tasks

Lab 13.01 - Traffic Optimization

Task 1

- Airmatch Configuration

Task 2

- WLAN Optimization

Task 3

- Wired QOS

Task 4

- Wireless QOS Marking

Task 5

- Wireless WMM Voice Class

Module 14 -
Monitoring



MOD 13- 62


The slide provides a brief review of tasks performed in this lab. Please see your Lab guide for details of this lab.

Monitoring

Implementing Aruba Campus Access

Confidential - For Training

This is the final module in the course and is all about monitoring your network.



Objectives

- Use your knowledge of Central alerts to improve your situational awareness
- Leverage Aruba Client Insights to know what is connected to your network
- Describe Aruba's User eXperience Insight (UXI) solution and how it helps to improve customer and end-user satisfaction

This module is really about improving your situational awareness, to better serve your customers and end-users. After completing this module, you will be able to:

- Use your knowledge of Central alerts to improve your situational awareness.
- Leverage Aruba Client Insights to know what is connected to your network.
- Describe Aruba's User eXperience Insight (UXI) solution and how it helps to improve customer and end-user satisfaction.

Central Alerts

Let's begin this module by covering alerts in Aruba Central.

Confidential - For Training

Alerts Page

Alerts

Events

Summary

List

ALERTS

43

CRITICAL 24

MAJOR 18

MINOR 0

WARNING 1

>

CLICK HERE FOR ADVANCED FILTERING

Show Acknowledged Alerts

☐

Open Alerts (43)

ACKNOWLEDGE ALL

Occurred On	Category	Site	Group	Severity	Description
Nov 17, 2022, 01:43:05	Radio Noise Floor	AnyCorp BO Sydney (SYN)	AnyCorp Branch Offices	major	Noise floor on AP BO-SYD-AP-01 operating on channel
Nov 17, 2022, 01:30:05	Radio Noise Floor	AnyCorp BO Tokyo (NRT)	AnyCorp Branch Offices	major	Noise floor on AP BO-NRT-AP-01 operating on channel
Nov 16, 2022, 23:22:05	Radio Noise Floor	AnyCorp BO Johannesburg (JNB)	AnyCorp Branch Offices	critical	Noise floor on AP BO-JNB-AP-01 operating on channel
Nov 11, 2022, 02:10:04	Radio Noise Floor	AnyCorp BO Tokyo (NRT)	AnyCorp Branch Offices	major	Noise floor on AP BO-NRT-AP-02 operating on channel
Nov 9, 2022, 01:10:01	Uplink speed flapping		AnyCorp Branch Offices	critical	Uplink GE0/0/1 speed flapped 60% on device BO-BLR-C
Nov 9, 2022, 01:09:44	WAN Uplink Autoneg...		AnyCorp Branch Offices	critical	WAN ports autonegotiation speed changed from 1000 I
Nov 8, 2022, 22:29:05	Radio Noise Floor	AnyCorp BO São Paulo (GRU)	AnyCorp Branch Offices	critical	Noise floor on AP BO-GRU-AP-02 operating on channel



MOD 14- 4

The Alerts and Events panes display all types of alerts and events generated for events pertaining to device provisioning, configuration, and user management. The Alerts and Events dashboards display a list of alerts and events generated for events pertaining to device pre-provisioning and configuration. You can view the alerts and events in the List view and Summary view.

You can acknowledge alerts as well as view their details. Alerts are acknowledged automatically when the event count drops below the lowest severity threshold configured for the alert. Users with admin access can acknowledge alerts irrespective of the severity configuration. As manually acknowledging an alert does not reset the count data, the alert service continues to aggregate events. When the number of new events meets the configured threshold, an alert is triggered again.

Advanced Alert Filtering

Aruba Central allows you to filter the alerts based on the alert categories. To filter alerts based on alert categories, complete the following steps:

From the Alerts page, click Click here for advanced filtering to filter the alerts based on alert categories.

Select the alert category and click Filter. You can select multiple categories from the

advanced filtering option.

The Open Alerts table displays the list of alerts generated in each alert category. The filter summary bar displays the total number of alerts in the selected categories.

Optionally, to clear advanced filtering option, from the Alerts summary bar, click Clear All. The advanced filtering gets cleared.

Once the issue is resolved alerts will be removed

Confidential - For Training Purposes Only

Alert Configuration

Alerts Events Summary List Config

ALERT SEVERITIES & NOTIFICATIONS Search...

USER	<u>ACCESS POINT</u>	SWITCH	GATEWAY	CONNECTIVITY	AUDIT	SITE
						Default Recipients All Enabled
	New Virtual Controller Detected ✓		Virtual Controller Disconnected ✓			New AP Detected +
	AP Disconnected ✓		Rogue AP Detected ✓			Infrastructure Attack Detected ✓
	Client Attack Detected ✓		Uplink Changed ✓			Modem Plugged ✓

MOD 14- 5

Aruba Central allows network administrators and users with administrative permissions to configure alerts. The configuration view is not visible to users who do not have edit permission to the Alerts and Events module.

1. In the Aruba Central app, set the filter to Global. The dashboard context for the selected filter is displayed.
2. Under Analyze, click Alerts & Events. The Alerts & Events page is displayed in the List view.
3. In the Alerts & Events page, click the Config icon. The Alert Severities & Notifications is displayed.

Creating Custom Alerts

The screenshot shows the 'Alerts' configuration page for a switch port transmit rate. The breadcrumb is 'SWITCH > Switch Port Tx Rate'. There are tabs for 'Alerts' and 'Events'. On the right, there are links for 'Summary', 'List', and a 'Config' button. Below the breadcrumb, there are links for 'Expand All' and 'Collapse All'. A trash icon is visible. The 'Rule 1' configuration section includes a 'Transform Function' dropdown set to 'percentage'. The 'Severity' section has four radio buttons: 'Critical', 'Major' (selected), 'Minor', and 'Warning'. Below the severity section, there are four input fields for 'exceeds:' values, with the second field set to '80'. The 'Duration' is set to '5 mins'. The 'Devices Filter Options' section has a 'Group' dropdown set to 'Select the Group'.

MOD 14- 6

Use the tabs to navigate between the alert categories. Select an alert and click + to enable the alert with default settings. To configure alert parameters, click on the alert tile and do the following:

Severity—Set the severity. The available options are Critical, Major, Minor, and Warning.

Duration—Enter the duration in minutes.

Device Filter Options—(optional) You can restrict the scope of an alert by setting one or more of the following parameters:

- **Group**—Select a group to limit the alert to a specific group.
- **Label**—Select a label to limit the alert to a specific label.
- **Device**—Select a device to limit the alert to a specific device.
- **Site**—Select a site to limit the alert to a specific site.

Other Filter Options—(optional) In addition to the device filter options few alerts also has further filtering options in order to restrict the alert scope by setting one or more of the following parameters:

- **SSID**—For fewer Connectivity alerts, you can select an SSID to limit the alert to a specific SSID.
- **Interface**—For few switch port alerts, you can mention an interface value to limit the alert criteria to a specific port.

- **Band**—For few Access Point alerts, you can select the band, 2.4 GHz or 5 GHz or 6 GHz to limit the alert to a specific band. 6 GHz Band is only supported for devices with 6 GHz capability.
- **Subnet**— For few gateway alerts, you can limit the alert for a specific DHCP pool identified by its subnet.
- **Provider Tag**—For few gateway alerts, you can limit the alert to a particular uplink identified by its link tag (Provider Tag).

Notification Options

- **Email**—Select the Email check box and enter an email address to receive notifications when an alert is generated. You can enter multiple email addresses, each separated with a comma. The Default Recipient check box is selected by default. If you want to disable specific email addresses from the default list to avoid sending alert notification, click the number displayed in parenthesis and click against each email address. To add or delete default recipient, see Adding Default Recipients. Uncheck the Default Recipient check box in order to disable alert notifications to all the default email addresses. The number displayed in the parenthesis denotes the total number of email addresses that have been already configured as default recipients to receive notifications when an alert is generated.
- **Streaming**—Select the Streaming check box to receive the streaming notifications when an alert is generated.
- **Webhook**—Select the Webhook check box and select the webhook as required. You can also Create Webhook for alert notification.

6. Click Save.

Default Recipients

The screenshot shows the Aruba Central interface. At the top, there are tabs for 'Alerts' and 'Events'. Below these, there are icons for 'Summary', 'List', and 'Config'. The main section is titled 'ALERT SEVERITIES & NOTIFICATIONS' and includes a search bar. Below the search bar, there are tabs for different device types: 'USER', 'ACCESS POINT', 'SWITCH', 'GATEWAY', 'CONNECTIVITY', 'AUDIT', and 'SITE'. The 'ACCESS POINT' tab is selected. In the top right corner of this section, there is a button labeled 'Default Recipients' which is highlighted with an orange box. Below the tabs, there is a grid of alert categories, each with a name and a checkmark indicating if it is enabled. The categories include: 'New Virtual Controller Detected', 'Virtual Controller Disconnected', 'New AP Detected', 'AP Disconnected', 'Rogue AP Detected', 'Infrastructure Attack Detected', 'Client Attack Detected', 'Uplink Changed', and 'Modem Plugged'.

MOD 14- 7

To set default recipients for alert notification, complete the following procedure:

1. In the Aruba Central app, set the context filter to Global. The dashboard context for the selected filter is displayed.
2. Under Analyze, click Alerts & Events. The Alerts & Events page is displayed in the List view.
3. In the Alerts & Events page, click the Config icon. The Alert Severities & Notifications is displayed.
4. In the Alert Severities & Notifications page, click Default Recipients. The Default Recipients dialog box is displayed.
5. Click the + icon to add the email address that you want to add as a default recipient to receive notifications when an alert is generated. You can add multiple email addresses as required.
6. Click Save.

Configuring Alert Notifications at Site Level

Alerts **Events** **Summary** **List** **Config**

ALERT NOTIFICATIONS

Suppress Alerts ☐

Email Configuration ☒ **Override** ☐ Append

Site level email destinations, if configured, will override email destinations from global alert configuration pages.

Email Configuration

MOD 14- 8

Aruba Central enables you to configure site-specific email addresses for notifying alerts. When alerts are generated for a specific site, the email notification is automatically sent to the email addresses configured for that site. The email addresses configured in the site dashboard overrides the email addresses configured in the global dashboard. To add an email address in the site dashboard, complete the following steps:

1. In the Aruba Central app, set the context filter to one of the options under Sites. The dashboard context for the selected filter is displayed.
2. Under Analyze, click Alerts & Events. The Alerts & Events page is displayed in the List view.
3. Click the Config icon. The Alert Notifications page is displayed.
4. In the Email Configuration Override window, click + to add an email address.
5. In the text-box, enter a valid email address.
6. Click Save.

Device Insights



Maximize your awareness of device connectivity with device insights. This section provides an introduction the to Device Insights feature in Aruba Central.

Confidential For Training

Central Device Insights Overview

Cloud-based profiling engine

Recognizes wired and wireless endpoint clients

Does not rely on physical collectors

Telemetry collected directly from Aruba devices

Available to **ANY** Aruba Central customer

Requires **Foundation** license



MOD 14- 10

Device Insights is a cloud profiling engine built to recognize endpoint clients (IoT, etc.) that connect to wireless and wired networks. Aruba has extended the device recognition engine built for ClearPass Device Insight (CPDI) for use by customers running Aruba infrastructure managed by Aruba Central cloud. Instead of relying on physical collectors used by CPDI, Aruba now utilizes the telemetry data collected directly from Aruba APs, gateways and switches. Any Central customers can now see profiling information displayed in the GUI's Profile tab within the Clients page. There's also a separate profile tab for each client.


Foundation licensing is needed for AI-powered profiling, along with visibility into static/flow-based attributes. When available, Advanced licenses will offer the ability to utilize tags for granular security features.

Access Device Insights

The screenshot displays the Aruba Central web interface. At the top, the 'aruba Central' logo is on the left, and a search bar with the text 'Search or ask Aruba' is on the right. Below the logo, the customer name 'Customer: HPE hqlow' is shown. The left sidebar contains a menu with the following items: 'Global' (highlighted with an orange box), 'Analyze', 'Alerts & Events', 'Audit Trail', 'Tools', 'Reports', 'Launch', 'App Catalog' (highlighted with an orange box), 'Maintain', and 'Firmware'. The main content area features the 'App Catalog' header with a globe icon. Below this, the 'ClearPass Device Insight' section is visible, including a padlock icon, the title 'ClearPass Device Insight', the description 'Discover and profile devices connected to the network.', and a 'Launch' button (highlighted with an orange box). A large diagonal watermark reading 'Confidential - For Training Purposes Only' is overlaid across the center of the interface. The text 'MOD 14- 11' is located in the bottom right corner of the interface.

Since Aruba Central Cloud 2.5.5 the CPDI standalone solution launch point was moved to App Catalog in the main Central GUI. This is available only on non-MSP accounts.

How Does it Work ?



How does Client Insight collect data without using collectors? Each Aruba AP, switch, and gateway is designed to automatically send telemetry data to the profiling engine hosted in Aruba Central.

Note: The CPDI standalone solution is intended only for those customers that are not Central cloud customers.

For telemetry data, the minimum version for each Aruba device type follows:

Aruba AP: 8.6.0.9

SD-Branch gateway: 8.7.0.0-2.3.0.0

AOS 10 gateway: 10.2.0.0

AOS-CX switch: - 10.08

Telemetry Data Collected

Static attributes	Telemetry data	AOS-CX switches	APs	AOS10 gateways
	MAC Address	Yes	Yes	Yes
	DHCP Options: DHCP Option 55 DHCP Option 60 DHCP DORA Options	No	Yes	Yes
	User Agent	No	Yes	Yes
	OS via SMB	No	No	Yes


Flow attributes	Telemetry Data	AOS-CX switches	APs	AOS10 gateways
	Applications	No	No	Yes
	Protocols	No	No	Yes
	Destination Connections	No	No	Yes
	DNS Domain Names	No	No	Yes



MOD 14- 13

The slide summarizes type of information sent to Aruba Central's Device Insights service by the different types of Aruba devices. Remember that this is information about the devices (endpoints) that the Aruba devices with which they are interacting.

Device Dashboard



Device Insights dashboard provides aggregated information about known clients and their device types.

Client Profile

Devices > 1c4d7007fd14 / 172.16.1.8



Status
Idle

Hostname
CODZW1UZZL

Type
Windows 8/10

 Filter Devices

 Reclassify

 Export

PROFILE | NETWORK | HISTORY | SECURITY

PROFILE INFORMATION

DEVICE CATEGORY Computer	MAC ADDRESS 1c4d7007fd14	FIRST SEEN Sep 29 2022, 08:15:54 pm	IP 172.16.1.8
DEVICE FAMILY Windows	MAC VENDOR Intel Corporate	LAST ACTIVITY Nov 23 2022, 09:31:44 pm	TAGS [Computers & Servers]
DEVICE TYPE Windows 8/10	HOSTNAME CODZW1UZZL	UPDATED AT Nov 23 2022, 09:31:44 pm	
DEVICE SOURCE CP0012014	USERNAME -		

CLASSIFICATION

CLASSIFIED BY
System Rule

CONDITION(S)

FINGERPRINT ATTRIBUTES	OPERATOR	VALUE
DHCP Option55	Equals	1,3,6,15,31,33,43,44,46,47,119,121,249,252

MOD 14- 15

For Client Insights, in addition to the profiling data displayed within the Clients Profile tab on the Global Clients page, there is also the Profile tab view when drilling down to an individual client. The Profile tab for an individual client contains more detailed information, such as the Classification mechanism used to profile the client, static client attributes, flow attributes (only if telemetry is received from SD-Branch Gateway or AOS 10 gateway), and network activity.



User Experience Insight

See things from the user's perspective with Aruba's User eXperience Insight (UXI) technology. UXI was briefly introduced in the Aruba Campus Access Fundamentals course and associated HPE Press book. This part of the module, you become much more familiar with UXI.

Aruba User Experience Insight



Onboards like a user

- Sensors installed near users or on a Zebra device
- Logs in testing applications and services like a user
- No cable run required



Onboard to any network

- Wireless and wired testing
- Measure DHCP, DNS, etc.
- Logs into captive portals



Test from the edge

- Are services and apps working?
- Zebra: Roaming, SIP call and low impact synthetic testing
- Sensor: Test throughput, and through the network



Actionable insights

- Realtime alerts, trend and baseline, and root cause analytics
- Integration to your NOC
- Stream to reporting tools

MOD 14-17

Aruba UXI leverages a cloud-hosted dashboard and physical sensors deployed onsite. Together, they function as your remote technician. The UXI sensors onboard onto your network as an end-user and actively monitor wireless and wired connections, providing a detailed performance analysis from an edge perspective via a highly intuitive cloud-hosted UXI dashboard. Similar to the sensor, Aruba now has sensor software which can be installed on a Zebra Android-based device. Aruba has integrated UXI with Zebra's Wireless Insights engine to provide real roaming statistics—did it roam, why did it roam or why did it fail to roam?—and true SIP call measurement to be able to provide call metrics like if a call dropped who dropped the call. Aruba's UXI can measure end-to-end VoIP.

Note: An overview of UXI can be found here:

https://www.arubanetworks.com/assets/ds/DS_User-Experience-Insight.pdf.

Here is the process of setting up UXI:

Step 1: Sensors onboard via cellular or standard PoE. You just plug in the sensor and it looks for the network it's supposed to test. As soon as it finds that network, the cellular function (if being used) turns off. (Note: the agent just onboards to a wireless network, not a wired network.)

Step 2: Just like a normal user the sensor associates to a network, goes through all of the same DHCP, DNS, and gateway (DORA) process. If there is a captive portal involved, the

sensor will log in and log out of that captive portal. UXI measures load, and redirect timings (and success) of that captive portal.

Step 3: Configure the sensor/agent to test services and applications which are critical to that location. First, the sensor will test internal applications: things on the inside of the firewall. Maybe printers, authentication servers, the default gateway, and etcetera. In the case of Zebra, the sensor may perform a test using a basic ping/webserver of a work app or voice analysis of calls. (Note: the agent does not do web application tests, or throughput tests because it would impact the performance of the device itself)

Step 4: All test results, incidents, and alerts are sent to the UXI cloud dashboard real-time. These views give you the ability to trend, baseline, and troubleshoot in a very organized way. AIOps enables easy network operation center (NOC) notifications. People who tend to use this dashboard are, at a high level, people who are responsible for a particular service: they just want to know “is my application working at the xx location?” This dashboard also provides a simple way to see an external service, like Facebook for example, has an outage and/or seeing if everything is working on the inside of the network, and perhaps only a specific application is having an outage. Alternately we stream all data/alerts/incidents to a container of your choice.

Along the way, the device assesses performance of wired and/or wireless networks, the steps to get onto a network, execute transactions, exchange data, and etcetera. Every step of the way, every metric that’s possible to capture is captured.

The sensor onboards just as a client would and therefore:

Can measure end-to-end application performance from client’s perspective.

Is not subject to any of the same points of failure as the network .

Is vendor neutral.

UXI provides a single pane of glass into the metrics that are gathered. UXI and Aruba Central/AI Insight provide assurance of connectivity to and through the network with optimal performance and Wi-Fi coverage. As you would expect from the names and descriptions, individually they provide insights into how the network is performing and should perform. Collectively, they provide holistic intelligence from the perspective of the customers/end-users and the network infrastructure. Note that AI Insight requires an Aruba infrastructure;

there are other requirements too, like ideally having over 2,000 APs).

Confidential - For Training Purposes Only

UXI Components

Purpose-built sensors and agents



- Actively test wired and wireless
- Baseline & trend
- Test across any network
- Power outage notifications
- Device roaming, voice analysis
- Testing across any network
- Becomes another sensor
- Global baselining metrics
- Packet capture details

Simple to set up and deploy on any wired and wireless networks

Cloud-based UI




- Device and network analytics
- Prioritize critical issues
- Integrate into your workflow
- Single pane NOC views

Understandable real-time interface

MOD 14-18

UXI consists of front-end sensors and agents, and a back-end cloud-based UI. The sensors act like an end-user, creating traffic typical of an end user. The slide summarizes sensor and agent features and advantages. The Cloud-based systems provide real-time monitoring and analytic intelligence, as summarized in the slide.

Aruba Central Integration



MOD 14- 19


The UXI user experience status can be incorporated into the Aruba Central Network Health page. The green smiley icon is an example shown in the slide of the UXI status.

UXI Tests



Wi-Fi

- Scan
- Association
- Packet capture
- Measure RSSI, util., etc.
- Captive portal



Network

- Authentication
- Gateway
- DNS
- DHCP
- Iperf3
- Ping/Hping
- HTTP Get /Curl
- Telnet



Application

- Test any web application (easy recording/script)
- Office 365, Gmail, AWS (w/ customer login)
- Dropbox (w/ customer login)
- Slack (w customer login)
- Skype for Business, Blue Jeans and other VoIP services
- Netflix
- YouTube

MOD 14- 20

Here's a sample list of what is tested. Note that the list on the right of the slide of the SaaS apps tests is only a sample; there's more and Aruba's team is constantly adding to the list. If there's something specific that you need to test, it's possible to create your own test.

Test to and through your network


- Log into website and applications
- Office 365, Gmail, AWS
- Pings, web server
- Team and other VoIP services
- Netflix, YouTube, Dropbox, and etcetera
- Libraspeed and iPerf
- Voice Analysis and roaming

At-A-Glance Status



You must currently log in to a separate portal to access the UXI information. The status dashboard is a simple way to see, at-a-glance, whether everything is performing well or if something needs attention. You can immediately see if there is an issue and then drill down to pin-point the problem. When you see green, all is good and performing well. When you see red, you know that you have work to do in order to make your users happy.

Installing New Sensors




MOD 14- 22



Before installing sensors, you must assign them to the customer account and configure them in the UXI account on Aruba's UXI portal. Sensors do not properly boot if you do not complete this process: you must add sensors in the UXI GUI. If your sensor has internet access via its Ethernet port or its built-in LTE connection, the sensor downloads its configuration and upgrades software as needed.

Aruba recommends that any sensor you install without an ethernet port should be installed as a sensor with LTE, to facilitate easy of deployment. You could wire (Ethernet) and preconfigure the sensor, and then ship it to its site; but this could be a lot of work if you are deploying many sensors.

UXI App: New Sensors



MOD 14- 23






Some sensors may be connected to an 802.1X wired port or may need a proxy to get to the Internet. For this reason, you can use a smartphone with the UXI app and use Bluetooth to connect to the sensor.

First, you must set up the sensor with a wired 802.1X or proxy configuration. Here's how. The app scans for the sensor then asks the UXI system if there is a configuration for this sensor. The app downloads the configuration and sends it, via Bluetooth, to the sensor. This gives the sensor the wired 802.1X or the proxy configuration. You can use this same method to send the configuration to a sensor that has neither an LTE nor an Ethernet connection.

Let's take a look at that app.

UXI APP





MOD 14- 24



The UXI app is available for both Apple iOS and Android devices, but its only available for the newer G-sensors (the older F-sensors lack Bluetooth capability). Nothing much to do here: just open the app and scan for the sensor. Once the app finds the sensor using Bluetooth, it checks with the UXI system for a configuration and firmware if needed.

Configured via APP




MOD 14- 25



The app receives the configuration from the UXI cloud system and sends the configuration and software updates to the UXI sensor via Bluetooth. Once done, the sensor is up and running and you can move on to any other needed sensor installations.

Groups



MOD 14- 26



If you have several sensors and they are not grouped, then all the sensors can do any test configured. But you can use groups to assign different SSIDs and run different tests. In the example show in the slide, some sensors are in Building-1, some are in Building-2, with different SSIDs and tests to run. So, you create different groups for the two buildings.

You can also use groups to zero in on certain areas of your network. In the dashboard, pick a specific group to view and only sensors and the tests for that group will display.

And the sensors are easy to configure.

Configuring Groups

Dashboard

GROUPS

Wireless

Wired

LOCATIONS

Sensors

Groups

TESTING

Groups

Group Management

NAME	ALIAS	NO. OF SENSORS	
New-York	NY	1	Manage
USA-Group	G1	2	Manage

+ Add Group



MOD 14- 27




To add a group, navigate to Setting > Group and click Add Group. Enter a name and at least two to three characters for an alias. Use the pencil icon next to the group under Group Management to change these attributes or delete the group. Click Manage to see a location map for the group, where you can also add or remove sensors from the group. When you add a sensor or edit a sensor you can also select a group. And these groups can have a hierarchy.

Sensor Testing Networks

4 network tests:
3 wireless / 1 wired

- 3 SSIDs + 1 Wired
- 4 SSIDs
- 1 SSID + 3 Wired

← Minimum 1 SSID



Multiple wired tests Trunk port with multiple VLANs


MOD 14- 28

You can configure four networks to run your tests. This could be three SSIDs and one wired port, as shown in the slide. Or perhaps 4 SSIDs, or one SSID and three wired ports: whatever you need. There is one restriction, though: you must always have at least one SSID network (the software cannot conduct four wired tests). To perform multiple wired tests you must configure the switch port as a trunk with multiple VLANs.

Let's look at how this plays out.

Sensor Testing Cycle: Operation

4 network tests: 3 wireless / 1 wired



MOD 14- 29




You setup four networks of interest. Then run the tests, starting with the first SSID. When all the tests have been completed for a specific network, the sensor will release the IP address. if obtained, from DHCP and de-authenticate from the network. Then it moves on to the next network, and so on and so forth, performing these tests in a round robin fashion.

Sensor Testing Cycle: AP Selection

Choose the best BSSID

- Highest energy
- Best throughput
- And more




MOD 14-30



Typically, a sensor detects multiple APs with the same SSID. It selects a particular AP based on several factors, including energy detection and predicted throughput, rather than just raw Received Signal Strength Indicator (RSSI). Plus, this decision is not strictly up to the client: the network considers several factors, too. Basically, the sensor acts like a client: it associates and makes the same decision a computer or user device typically would. You can, however, select a specific BSSID to be test.

UXI Test Types




MOD 14- 31

There are three areas for testing:

- Services (core test): Will test your connectivity to the networks.
- Internal test: Used for services inside your network.
- External test: Will verify access to services outside of your corporate network.

Services Test



MOD 14- 32

The types of core tests include the following:

- Wi-Fi test: Examines the Wi-Fi environment around the sensor and collects BSSIDs, RSSI values, channel information, and related information.
- SSID test: Checks the availability of the configured SSIDs on the sensor in the Wi-Fi environment.
- Association test: Makes sure the sensor can associate as well as how long it would take to associate to the configured SSID/AP.
- Ethernet test: Gets an IP address and marks the speed and duplex of the line, if you have a wired port configured.
- DHCP test: Checks whether the sensor can obtain an IP address using DHCP. (Note: UXI will skip this test if you have configured a static IP address.)
- Gateway test: Checks whether the sensor can reach the provided DHCP or configured static gateway.
- DNS test: Will use the obtained DNS (via DHCP) or statically configured DNS nameserver to resolve `cdn.capenetworks.io` for both the primary DNS and secondary DNS.

- External connectivity test: Ensures the sensor can connect to <http://cdn.capenetworks.io/auth>.

Once UXI has successfully completed all the core tests, the sensor will run the internal and external tests configured for the networks.

Confidential - For Training Purposes Only

Adding Internal / External Tests

Networks & Groups

The enabled tests below apply to **all sensors** across 2 groups and 3 networks.

Change Selection

Tests

+ Add Test

INTERNAL SERVICES	TARGET	TESTS	TESTING
Clearpass-test	10.254.1.71	Ping	☒ ON
Telnet Server3	67.1.1.2	Telnet	☒ ON
VoIP MOS	10.254.1.102	VoIP MOS	☒ ON
Webserver	10.254.1.21	Port 80, Port 443, Ping, HTTP Status Codes, Val...	☒ ON
iPerf3	10.254.1.102	iPerf3	☒ ON



MOD 14- 33

The installation wizard in the example in the slide added three internal/external tests. To add, change, and/or delete these tests, you will need to navigate to Settings and select Service & App Tests. Here you will find the tests displayed in two sections: internal and external tests. In this window, you can click Add Test to create a new test. You can click the pencil icon beside a test to modify it or delete it, or you can temporarily disable a test by toggling the on/off slider button. Any test created without setting the network or groups will run on all sensors.



Knowledge Check

Self-check on key learning points



It's time for a knowledge check to review what you have learned.

Confidential - For Training Purposes Only

Question #1

Aruba Central allows you to filter alerts based on alert categories.

- True
- False

a

Knowledge Check

MOD 14- 35

+

true

Question #2

Aruba Central can provide deeper device insights if you add physical collector devices in the access and distribution layer of your network.

- True
- False

a

Knowledge Check

MOD 14- 36



false

Confidential - For Training Purposes Only

Question #3

Aruba's User eXperience Insight (UXI) solution allows you actively test wired and wireless networks, so you can see things from an end-user's perspective.

- True
- False

a

Knowledge Check

MOD 14- 37

+

true

Confidential - For Training Purposes Only

Question #4

One consideration with Aruba's User eXperience Insight (UXI) solution is that the sensors require an Ethernet or Wi-Fi connection for initial deployment.

- True
- False

a

Knowledge Check

MOD 14- 38

+

false

Confidential - For Training Purposes Only

Question #5

With Aruba's User eXperience Insight (UXI) solution, one sensor can be used to test both wired and wireless networks.

- True
- False

a

Knowledge Check

MOD 14- 39

+

true

Confidential - For Training Purposes Only



Lab Activity

Consider the lab activity.

Confidential - For Training

Lab Tasks

Lab 14.01 – Monitoring with UXI Sensor

Task 1

- Monitoring Using the Aruba UXI Sensor

Task 2

- Integrate with Aruba Central

Task 3

- Reset the lab Customer Environment



MOD 14- 41

The slide provides a brief review of tasks performed in this lab. Please see your Lab guide for details of this lab.

Next Steps



Certification

How to Study for Aruba Certification Exams – Free!
Sign up, take the exam!

Advanced knowledge

The next courses in this series is:
Security, Central, and SD-WAN courses



You have successfully completed this course. Nice job! Now, what are the next steps?

For certification, you should study diligently for the exam. For more information on scientifically proven, effective study techniques, consider taking the short course *How to Study for Aruba Certification Exams* – its free! At the time of this writing, the URL for this course is here:

https://inter.viewcentral.com/events/cust/search_results.aspx?cid=aruba&pid=1&lid=15&tstamp=1670423864146&event_id=754

When you are confident and prepared, sign up for and take the exam.

To take your knowledge the next level, sign up and attend other courses to broaden your knowledge!

Confidential - For Training Purposes Only

6280 AMERICA CENTER DR. SAN JOSE, CA 95002
TEL: 408.227.4500 | FAX: 408.227.4550
www.ARUBANETWORKS.com

EDU-ICAC-ILT-v23.11