



Aruba Campus Access Fundamentals

STUDENT GUIDE



Aruba Campus Access Fundamentals

22.41

Student Guide

December 2022

Aruba Campus Access Fundamentals

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-ACAF-ILT-v22.41

DECEMBER 2022

Aruba Campus Access Fundamentals

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
Course Objectives	7
Agenda Part 1	9
Agenda Part 2	10
Resources	11
Module 1: Networking Fundamentals	12
Objectives	13
Basic Network Concepts	14
OSI Model & Encapsulation	19
Physical Media	36
TCP/IP Stack	46
Networking Devices	62
Knowledge Check	71
Lab Activity	78
Module 2: AOS-CX Switching Fundamentals	80
Objectives	81
Initial Switch Setup	82
Introduction to VLANs	113
Discovering the Network	134

Link Aggregation	146
Knowledge Check	163
Lab Activity	172
Module 3: Basic IP Setup	174
Objectives	175
IP Addressing	176
Inter-VLAN Routing	194
Static IP Routing.....	204
OSPF	209
Knowledge Check	249
Lab Activity	256
Module 4: Network Redundancy	258
Objectives	259
Spanning Tree.....	260
VRRP	298
Introduction to VSX	308
Knowledge Check	312
Lab Activity	319
Module 5: VSF.....	321
Objectives	322
VSF Introduction.....	323
VSF Configuration	332
Auto VSF	336
VSF MAD	342
Knowledge Check	345
Lab Activity	349
Module 6: Introduction to Aruba Solutions	351
Objectives	352
Aruba ESP	353
Aruba Switch Portfolio	356
Aruba WLAN Portfolio	366
Introduction to Aruba Central.....	384
Introduction to Aruba ClearPass.....	390

Knowledge Check	393
Module 7: Central for Device Management.....	400
Objectives	401
Device Onboarding.....	402
Aruba Central Licensing	411
Central Groups, Sites, and Labels	420
Knowledge Check	437
Lab Activity	443
Module 8: Device Profiles & AP Onboarding.....	445
Objectives	446
Device Profiling	447
Connecting APs to Aruba Central	457
Knowledge Check	463
Lab Activity	465
Module 9: WLAN Fundamentals	467
Objectives	468
RF Bands & Channels	469
RF Basics	480
WLAN Mobility	502
802.11 Frames	507
WLAN Design.....	514
WLAN Setup.....	523
Knowledge Check	532
Lab Activity	537
Module 10: Implementing Secure WLANs	539
Objectives	540
Introduction to AAA	541
802.1X Authentication	549
Secure WLAN Setup	571
Roles and Rules	573
Knowledge Check	582
Lab Activity	589
Module 11: Implementing Guest Access	591

Objectives	592
Guest Access Introduction.....	593
Captive Portal Authentication	596
Guest WLAN Setup	605
Knowledge Check	609
Lab Activity	612
Module 12: WLAN Security	614
Objectives	615
WLAN Security	616
Certificates	634
Cloud Authentication	642
Knowledge Check	646
Lab Activity	653
Module 13: Monitoring and Maintenance.....	653
Objectives	654
Monitoring with Aruba Central	655
LED Status	661
Firmware Upgrades	670
SNMP	675
AI Insights	680
Alerts & Reports	692
UXI	697
Knowledge Check	706
Lab Activity	714
Module 14: Troubleshooting.....	716
Objectives	717
Password Recovery & Factory Reset	718
Troubleshooting Central Connectivity	723
Spectrum Analysis.....	729
Central Troubleshooting Tools.....	731
Knowledge Check	738
Lab Activity	743
Next Steps.....	745

Rev 22.41



Course Overview

Aruba Campus Access
Fundamentals

Hello, and welcome to the course! Aruba Campus Access Fundamentals.

Course Prerequisites: Ensure Your Success

IMPORTANT

For success in this course, you should know the content in
Get the Edge: An Introduction to Aruba Networking

- Numerical conversion
- Network fundamentals
- AOS-CX network OS
- L2 switching & VLANs
- Spanning-Tree Protocol
- IP addressing
- IP routing
- WLAN fundamentals
- RF basics
- Deploying WLANs with Central

Your success is important to us! To do well in this course, you must know the content in the Aruba course *Get the Edge: An Introduction to Aruba Networking*. This includes the specific topics you see here.

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 1 of this course, you will not be charged for this course.

Introductions

Name	Company, Location	Job Responsibilities
		
Experience: LAN/WAN/WLAN		Class Expectations



MOD 0-3

Please introduce yourself

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Logistics



Breaks



Lunch



Restrooms



MOD 0-4

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



MOD 0- 5

If you haven't already done so, I strongly encourage you to sign up for your HPE Learner ID at the URL shown here and included in the resources tab at the top of this course window.

You need this 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.

Welcome To Aruba Central 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



MOD 0-6

Thanks for choosing to learn about Aruba Central! We hope you will take the next step to get certified on Aruba's technologies.

To learn more about our certification & training program visit us online at arubanetworks.com/certification

We encourage you to join the Airheads Community, a vibrant group of networking professionals 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 and self-directed labs.

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

Course Objectives

After completing this course, you will be able to:



- Describe networking components, models, encapsulation and media types
- Access switches and configure VLANs, tagging, IP addressing, LLDP, ICMP, and link aggregation
- Configure routing and DHCP relay functions
- Configure resiliency features like Spanning Tree, VRRP, VSF, and VSX
- Describe Aruba features and product portfolios – Aruba ESP, Central, ClearPass, and more
- Onboard devices to Aruba Central, and configure usability features – groups, sites, templates, and licensing
- Use device profiling, connect APs to Aruba Central
- Describe fundamental Wi-Fi concepts: RF bands, channels, and properties, WLAN mobility and 802.11 frames
- Describe RF design and configure WLANs

MOD 0-7

Here's your course objectives. Please pause the presentation if you like and review the first half of these objectives

Course Objectives

You will be also able to:



- Explore AAA, 802.1X, roles, and access rules
- Describe guest access and authentication and configure guest WLANs
- Describe WLAN security, certificates, and cloud authentication
- Use Aruba Central monitoring capabilities
- Analyze LED status, perform firmware upgrades, and enable SNMP
- Describe AI Insights, Aruba Central Alerts & Reports, and Aruba User eXperience Insight (UXI)
- Perform password recovery and factory reset procedures
- Troubleshoot Central connectivity, explore Aruba Central troubleshooting tools, and enable spectrum analysis

MOD 0- 8

And here's the second half of the objectives. Pause the presentation if you like to review.

Agenda: Part 1

1 – Network Fundamentals

2 – Switch Fundamentals

3 – Basic IP Networking

4 – Network Redundancy

5 – VSF

6 – Intro to Aruba Solutions

7 – Central Device Mgmt



MOD 0- 9

Here's the first set of modules you'll explore.

Agenda: Part 2

8 – AP onboarding, Device Profiles

9 – WLAN Fundamentals

10 – Implementing Secure WLANs

11 – Guest Access

12 – WLAN Security

13 – Monitoring and Maintenance

14 – Troubleshooting



MOD 0-10

And here are the next set of modules.

Resources



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



MOD 0-11

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

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Rev. 22.41

aruba
a Hewlett Packard
Enterprise company


Networking Fundamentals

Aruba Campus Access
Fundamentals



Welcome back. This is Module 1, Networking Fundamentals.

CONFIDENTIAL - FOR TRAINING



Objectives

- Define networking, LAN, WAN and their components
- Explain OSI model & encapsulation
- Discuss different types of physical media
- Compare unicast, multicast, and broadcast
- Explain TCP/IP stack
- Discuss different types of networking devices

Basic network concepts provides a foundation for deeper understanding.

After completing this module, you will be able to define networking, LAN, WAN and networking components.

You will be able to explain the OSI model and encapsulation, and discuss different types of physical media.

You will be able to compare unicast, multicast, and broadcast, explain the TCP/IP protocol stack, and discuss different types of networking devices.


Basic Network Concepts

We begin with basic networking concepts.

CONFIDENTIAL - FOR TRAINING

What is a Computing Network?

Group of computing resources that permit digital data exchange between devices




MOD 1- 4

What is a computing network?

Computing communication is the process of sending or receiving data from other entities. Computing network engineering makes this happen. It is all resources and infrastructure – both hardware and software - working to create bridges, paths and highways used to transport data between devices.

Taking pictures from your smartphone is great, but it's worthless if you cannot share them. Moving data out of your smartphone is an example of computing communication.

Network Classifications



MOD 1- 5

Based on the geographical coverage a computing network can be categorized as a Local Area Network (LAN) or a Wide Area Network (WAN).

A LAN is a group of computer devices that are geographically in the same place. For example, a group of devices within a building can be considered a LAN.

LANs are used in several settings:

- Small Office/Home Office (SOHO)
- Office LANs
- Building LANs
- Campus LANs

A WAN on the other hand is a group of computer resources that can communicate over large geographical distances – typically a few kilometers or miles, and perhaps thousands of miles, such as the Internet. The Internet is considered a WAN since it permits communication across countries and continents.

Typically, WANs are deployed by Internet Service Providers (ISPs) since those companies have the economic resources to interconnect long distances. Examples of WAN technologies include the following:


- Internet
- Multi-Protocol Label Switching (MPLS)
- Asynchronous Transfer Mode (ATM) (obsolete)

- Frame Relay (largely obsolete)
- Dark fiber

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

What is a Protocol?

Set of rules that computer devices follow to establish and maintain communications



MOD 1-6

Communication is the main purpose of a computing network, and this communication is enabled using protocols. A protocol is a group of rules or instructions that computer devices follow in order to properly establish and maintain communications.

Consider this analogy:

Alice meets Bob for the first time, and starts the conversation by saying "Good morning, my name is Alice".

Bob replies, "Good Morning Alice, my name is Bob."

This brief conversation is actually a procedure. Notice that Alice starts the communication with a greeting, and then she identifies herself. Bob's reply is also a procedure. He acknowledges Alice, and then he identifies himself by name. The implicit rules in this conversation help to establish and maintain a conversation. Likewise, computing devices exchange messages in a specific order, following specific rules.

OSI Model and Encapsulation




Let's move on to OSI model and encapsulation.

CONFIDENTIAL - FOR TRAINING

OSI Reference Model

Standard communication model for computing devices created by the ISO
Organizes computing communication in seven Layers



MOD 1- 8

In the mid-1980's during the fast evolution of computing, every vendor wanted to implement their own, proprietary communication protocol. These proprietary protocols created interoperability challenges. The International Organization for Standardization (ISO) solved the problem by presenting a standard communication model for computing devices: The Open Systems Interconnection (OSI) model.

The OSI model organizes computing communication in 7 layers. Each layer defines a phase of message processing.

The OSI layers are shown in the figure, and described below:

Layer 7: Application Layer

Layer 6: Presentation Layer

Layer 5: Session Layer

Layer 4: Transport Layer, which organizes data into segments, as you will soon learn.

Layer 3: Network Layer, which organizes data into packets

Layer 2: Data Link Layer, which organizes data into frames

Layer 1: Physical Layer, which organizes the data into bits, and transmits those bits using²⁰

physical hardware over wires, fiber optic cable, or RF Signals.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Layer 1: Physical Layer

Goal


Define physical aspects of signal transmission and receipt across media

Modulation


Convert logical data bits into physical signals

Demodulation

Convert signals into logical data bits



Media types



MOD 1- 9

The physical layer dictates the physical aspects of how signals are transmitted and received across some media.

Computing devices convert logical data bits into the correct physical principle depending on the media in use, this process is known as modulation.

The inverse process, converting signals into logical data bits is known as demodulation.

This layer also defines the material characteristics and the components to use to achieve the correct transmission of the messages. For example, consider the modem that is used in a home setting; your tablet connects to the modem using Radio Frequency (RF) signals that travel across the air. This modem connects to the Internet Service Provider (ISP) network using fiber optic cables. Thus, the router converts data into optical signals. Also, the printer connects directly to the modem via a twisted pair copper patch cable. This means that print jobs are converted into electrical voltage signals before being sent to the printer.

Layer 2: Data Link Layer

Goal Frame delivery and error detection

Media Access Control


- Controls how to access media
- Older Ethernet uses Carrier Sense

Link Layer Addressing

- Each station has a unique “name”
- **Ethernet** uses a **MAC** address

Error Detection

- During Layer 1 transmission



MOD 1-10

The Data Link provides three main functions:

Media Access Control

In polite human conversation, only one person talks at a time. While Bob talks, Alice politely listens. She detects when Bob stops talking and knows that it is her turn to speak. Similarly, for most media, only one device may successfully transmit at a time, so access to the media must be controlled. Likewise, many media access control techniques leverage a carrier sense mechanism: before a station may transmit, it must sense the state of the carrier or media. If a transmission signal is detected, the station must wait to transmit, until the currently transmitting station is complete.

Link Layer Addressing

In a crowded room, Bob may call out, “Hey, Alice can I buy you an ice cream?” Everyone hears these sounds waves, which travel over a shared media – the air in the room. However, only Alice responds - she was identified as the intended recipient. Likewise, each station on a LAN has a unique “name”. Instead of an alphanumeric name like Alice or Bob, each station is identified by a 6-byte hexadecimal number called a MAC address. All stations on a shared media receive the message, but only the device identified as the intended recipient processes and responds. Like humans, the other stations realize, “This message is not for me” and simply ignore the message. This information is added to the data from the upper layers as so-called header information, about which you will soon learn.

Error Detection

For the receiver side, Layer 2 helps to detect errors that could occur during layer 1 transmission. This avoids unnecessary processing of corrupted or incomplete messages. This is accomplished by adding a checksum of the data to the trailer. This is also called a Cyclic Redundancy Check (CRC)

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY


Layer 3: Network Layer

Goal

Establish communications across multiple LANs and WANs

Logical addressing

Unique source and destination identifiers are maintained across the path



Path discovery and selection

Algorithms and protocols find all possible paths and select the best one



MOD 1- 11

The main goal of the Network Layer is to establish device communications across multiple LANs or WANs, using the best available path. This is achieved using two fundamental techniques:

- Logical addressing. A unique Layer 3 identifier for the source and destination is maintained across the path.
- Path discovery and selection. The Network Layer runs algorithms and protocols to find all possible paths, and then choose the best path. Later in this course, you will learn more about this, and things like the Routing Information Protocol (RIP) and Open Shortest Path First (OSPF).

The communication between two computing devices can take a specific path but not necessary the same one will be used in the future. Protocols and algorithms used in this layer could update the path any time depending on multiple factors. You will learn more about it in this training.

Layer 4: Transport Layer

Goal


Control link reliability using segmentation, and error control

Segmentation

Sender divides data into smaller pieces

De-segmentation

Receiver reassembles the pieces



The Transport Layer controls the reliability of a given link through segmentation, de-segmentation, and error control. In this layer some protocols, like the Transmission Control Protocol (TCP) are connection-oriented. This means that the transport layer can keep track of the messages and retransmit those that fail. Other protocols, like the User Datagram Protocol (UDP), are stateless or connectionless. This means the transport layer does not keep track of the messages. The advantage of this is that processing these connections is relatively fast and easy to compute.

There are three main responsibilities of the transport layer, as described below:

- **Segmentation.** The sender's TCP or UDP process accepts files from the application and divides them into smaller pieces (typically 1500 bytes) called segments. Each piece is passed down to the lower layers and transmitted individually – over an Ethernet link in the example shown in the figure.
- **De-segmentation.** The receiver accepts each segment, puts them back in the correct order if need be, and reassembles the information. This then can be processed by the application.
- **Error Control.** Refers to the verification of the information received to avoid errors that could occur the lower Layers (1-3).

Note. Error detection is a process that happens in different Layers: 2, 4 and sometimes in 7 of


the OSI model.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Layer 5: Session Layer

Goal Setup, maintenance and teardown a session between computing devices

Session Individual conversation between two computers



MOD 1- 13

Layer 5 is responsible for setup, maintenance, and tear down of sessions between two computing devices. A session is a particular conversation between two computer devices.

Suppose that some user opens a browser and connects to a web page like <http://arubanetworks.com>. A session is created. Then the same user opens a different browser to the same destination, since the application is different a new conversation or session is created. Thus, two separate sessions are maintained.


Then the user might establish a connection to a different host for purposes of file transfer, and even more sessions to remotely configure seven Aruba switches. A typical computer could generate thousands of sessions.

Layer 6: Presentation Layer

Goal

Transforms data into the format that the application accepts

- Compression/Decompression
- Encryption/Decryption
- Code Translation (EBCDIC to ASCII)



MOD 1- 14

The Presentation Layer transforms data into the formats that the application accepts. Typical processes include:

- Compression/Decompression
- Encryption/Decryption
- Code Translation (EBCDIC to ASCII)

For example, the figure shows how an application passes the clear-text message "Hello" to the presentation layer process, which encrypts this message before transmission. This provides confidentiality. If any bad actors or hackers intercept this data, they will not be able to read the message. Of course, upon receipt of an encrypted message, only the intended receiver has the correct digital keys to decrypt the data.

Layer 7: Application Layer

Goal

End user interacts with this layer via an application


Functions

- Identify communication partners
- Provide network resources to user applications

Web browser



FTP client



MOD 1- 15

The Application Layer is the closest to the end user, which means that both the OSI Application Layer and the user interact directly with the software application.

Application Layer functions include:


- Identifying communication partners. The application layer determines the identity and availability of communication partners for an application with data to transmit.
- Provide network resources. This layer provides network services to user applications, such as file transfer, email, video conferencing, and many others.

Some examples of application layer include:

- Hypertext Transfer Protocol (HTTP), which relies on TCP for transport
- Hypertext Transfer Protocol Secure (HTTPS) which relies on TLS for the presentation and TCP for transport.
- File Transfer Protocol (FTP), which also relies on TCP
- Domain Name System (DNS), which typically uses UDP, but sometimes also uses TCP
- Trivial File Transfer Protocol (TFTP), which uses UDP as its transport mechanism

Layer Headers and Encapsulation

Each layer has a specific responsibility in the communication



MOD 1- 16

In the OSI model, each layer has a specific responsibility during network communications. In the computing world, devices use headers to establish a communication exchange and to control information on a particular layer, or the layers above it.

Note: A header that is generated on a specific layer by the sender can only be read at the same level on the receiver side.


See the figure above for an introduction to the encapsulation and decapsulation process.

Encapsulation. Is the process where each OSI Layer adds a header. This process is always done by the sender device.

Decapsulation. Is the process to read and interpret the header information. This process is always done by the receiver device.

Protocol Data Units (PDU)

Structure that considers the header and payload for each layer.



The OSI model introduces the concept of a Protocol Data Unit (PDU). This is simply a structure that considers the header and payload or data for each layer.

The following table summarizes the PDUs from Layer 7 to Layer 2.

Note: PDU1 does not exist, remember that Layer 1 refers to signals that cross the media.

There are three key terms related to PDUs that you should know:

Segment. Refers to the encapsulation that is done in layer 4. A segment is equivalent to PDU4. So, you might hear networking people speak about a TCP segment or a UDP segment.

Packet. Refers to the encapsulation that is done in layer 3. A packet is equivalent to PDU3. You might talk about layer 3 IP packets.

Frame. Refers to the encapsulation that is done in layer 2. A frame is equivalent to PDU2. You might talk about Ethernet frames or Wi-Fi frames.

You might notice that PDU2 – a layer 2 Frame - not only includes a header but also a trailer that is appended after data, as shown in the figure.

The trailer is typically used to detect errors during the transmission of the message. You recently learned about this during the discussion of layer 2 of the OSI model. Layer 2

protocols like Ethernet and Wi-Fi add a trailer, often labeled (CRC) or perhaps as a frame check sequence (FCS).


Also, notice that there is a maximum amount of data that can be transmitted over certain media – the so-called maximum transmission unit (MTU) size.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Maximum Transmission Unit (MTU)

Max transmission size of an interface

- ✓ Each segment connected to a router may have differing MTU size
- ✓ Routers might need to fragment packets
- ✓ Path MTU Discovery protocols can help to alleviate this issue
- ✓ Tunneling protocols add header info, may cause MTU issues



MOD 1-18

The MTU is the maximum transmission size an interface can handle, not including the Ethernet header of 18 or 20 bytes. It could be different for different interfaces, and for different media types. Upper layer protocols (ULP) break up data into appropriate size to fit within the interface max size.

Ethernet networks have classically used an MTU of 1500 bytes. However, as data rates increased, this small MTU size limited performance. Higher speed Ethernet interfaces may use MTUs much larger MTUs of around 4KB, 8KB, or larger – and are referred to as jumbo frames.

You just learned that all the ULP headers are considered part of the layer 2 payload. Thus, this detracts from actual raw data that can be sent.

Also, a router might connect several networks with different MTU sizes. When a router receives a frame from a network that uses 8KB MTU size, it may have to break it up into smaller, 1500-byte chunks. However, if the IP header sets the Do Not Fragment (DF) flag, the router must discard the packet and send ICMP message to the originator of the packet.

The need for routers to handle differing MTU sizes has been recognized as a performance inhibitor, and so it might be best to ensure that MTUs are consistent throughout most of your network. Newer Path MTU Discovery protocols might help you by auto-negotiating an end-to-end MTU. This alleviates routers of the need to break up packets, so they can focus on their main job: moving packets along the best path toward their destination.

Some people run into MTU issues when doing tunneling of packets, especially when using IPsec, PPPoE, GRE, etc. This is because tunneling protocols add additional header overhead, thus increasing the MTU size. This may happen without knowledge of the end systems sending packets. Packets get dropped, and you must be aware of this possibility and look for dropped packets as a part of your troubleshooting process.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY


Physical Media



Airwaves for Wi-Fi and copper or fiber optics for Ethernet: let's talk about physical media.

CONFIDENTIAL - FOR TRAINING

Wired Connection Using Copper



MOD 1-20


The most popular cable type is unshielded twisted pair (UTP). It is a cheap solution to interconnect devices in a LAN. The main consideration is the cable's maximum length, 300 feet or 100 meters.

This connection type uses electrical signals to communicate information with data rates up to 40 gigabit per second.

UTP cabling contains 8 color-coded wires that group into 4 pairs. The connector UTP uses is an RJ-45 connector.


UTP cables have categories that represent cable quality and the signal frequency band that can be carried by the cable. For example, UTP category (or cat for short) 6 cable supports 10Gbps speeds for up to 55 meters in distance; in contrast UTP cat 6E cable allows 10Gbps and can span a distance of up to 100 meters.

STP vs UTP



STP Cable

Shielded Twisted Pair



UTP Cable

Unshielded Twisted Pair



SFTP Cable

Shielded Foiled Twisted Pair



MOD 1-21

Copper cables used in networking come in several variants. Three common variants are shown in the figure.

UTP cable has no shielding. This is the most used and most basic type of cable. The cable contains pairs of wires twisted together to help reduce and prevent electromagnetic interference.

STP cables are like UTP cables, where the wires are twisted together and then wrapped with a shielding or screening material which consists of foil wrapping or a copper braid jacket. Shielding the copper cables can have a tremendous effect on the stability, performance, speed, and longevity of the cable and ultimately your network. Because STP cables have shielding they are more resistant to electromagnetic interference, like fluorescent lighting, nearby motors, or any other electro-mechanical equipment. This provides better signal quality and allows for higher transmission speeds.

STP and FTP cables use the same twisted pairs, but the individual pairs are wrapped in foil. This helps to reduce interference between each pair. Plus, the entire bundle is then wrapped in a strong braided screen. This provides additional mechanical strength and better grounding.

Cable Category


Category1	Cable Type	Maximum Transmission Speed	Signaling rate
Cat. 3	UTP	10 Mbps	16 MHz
Cat. 5		100Mbps	100 MHz
Cat. 5e		1Gbps	100 MHz
Cat. 6 (6e)		1Gbps	250 MHz
Cat. 6a	STP	10Gbps	500 MHz
Cat. 7		10Gbps	600 MHz
Cat. 8		40Gbps	2000 MHz



MOD 1-22

The Cat in Cat5e, Cat6, etc. is short for Category. Network cables are divided into categories based mainly on bandwidth (measured in MHz), maximum data rate (measured in megabits per second), and shielding. The table summarizes commonly used categories.

Wired Connection Using Fiber Optics




Use fiber optic cables to connect devices in your LAN that more than 300 feet or 100 meters apart. This technology uses light impulses to transmit data using fiber. Maximum length can vary widely depending on cable and transceiver type, from a few meters, to around 500 meters, perhaps up to 40 kilometers.

Optic fiber achieves faster data rates. A less expensive fiber transfers 10 Gigabits per second while a higher-quality fiber provides 100 Gigabit per second data rates.

We organize fiber optic cables into two main groups: Multimode (MM) is less expensive for shorter distances. Single Mode (SM) is more expensive but supports longer distances.

There are several standard fiber connectors, but you typically use a Lucent Connector (LC) , as shown here. You connect each end of the cable into a transceiver, which you have installed into a network device or server. Fiber optic transceivers are the active components that transmit and receive data, in the form of light pulses.

Direct Attach Copper (DAC) Cable



Twinax copper cable
Factory-terminated to transceiver



MOD 1-24


Direct attached copper (DAC) cables are twinax cables that are factory-terminated to a transceiver. You use them for relatively short connections, often from a switch in one equipment rack to another switch, server, or other device in the same rack. You might even connect from one modular switch blade to another blade in the same switch.

There are DAC cables that connect from one line card to another, as shown in the figure. There are also breakout DACs that connect from one device and up to four other devices.

A common option from Aruba is a DAC cable with SFP+ transceivers.

For more information, search for "ArubaOS-Switch and ArubaOS-CX Transceiver Guide" currently at this URL: <https://www.arubanetworks.com/techdocs/Switches/xcvrs/5200-3362/index.html#GUID-D7147C7F-2016-0901-0A24-0000000004A4.html>

Wireless Connection Using Radio Frequency



MOD 1-25

Most endpoints only use a wireless connection. This connection type has a typical range of 30 to 45 feet, or 10 to 15 meters. The range depends on variables like the Access Point position, the endpoint type, and the furniture in the area. Objects affect the radio frequency (RF) signal by absorbing or reflecting it. Optimal distance between APs (and thus APs and endpoints) is best determined by understanding client need, and then performing a site survey to accommodate that need. These topics are beyond the scope of this course.

The data rate depends on the Wi-Fi technology in use. The latest Wi-Fi technology is IEEE 802.11ax, with data rates up to 4.8 Gigabits per second.


To transmit data, Wi-Fi devices modify aspects of an electromagnetic signal, like amplitude, phase, or frequency. These modulated radio waves transmit through the air. This wireless transmission technology frees users from wired connections and allows them to roam.

Duplex


Duplex

- Communication system with two or more parties
- Bi-directional communication between parties
- 2 types: Half and Full

Half Duplex



Full Duplex



MOD 1-26


A Duplex communication system has two or more connected parties that communicate in both directions. There are two duplex communication systems types: Half duplex (HDX) and full duplex (FDX).

With half duplex, both parties do not communicate simultaneously. It is one direction at the time. An example is a walkie-talkie. Each person presses a push-to-talk button. When the button is pressed, the user does not hear the remote person. To listen, they release the button. A more modern use case for half duplex is the typical Wi-Fi network.

Note: The latest MuMIMO technology for WLANs does allow for multiple transmissions at the same time.

With full duplex, both parties communicate simultaneously. An example is a telephone. Parties on both ends speak and are heard by the other party at the same time. Again, a more data network-specific example is the typical Ethernet network using modern switches and devices: all of which support full-duplex connectivity.

Types of Traffic: Unicast, Multicast & Broadcast



MOD 1- 27

Computing communications have three types:

Unicast communications

Unicast traffic refers to one-to-one communication: a transmitter and a receiver. Imagine several learners in a classroom. Bob is the instructor. He calls out, "Alice, I have a message for you from the front desk." This message came from one source, Bob, and is destined to a single destination. Alice. Similarly, when a PC needs to transfer a file to a server, the two devices use unicast communications.

Multicast communications

Multicast traffic refers to one-to-many communication: a transmitter and multiple receivers. In our classroom analogy, suppose lunch is brought in. Bob calls out, "All vegetarian meals are on the green table". The message came from one source, Bob. It is destined for several people, the vegetarians. A common example is video streaming. A video source sends a video stream to capable devices interested in receiving that information: on their PC, tablet, or smartphone.

Note: Multicast traffic includes many devices, but not all network devices.

Broadcast communications

Broadcast traffic communication is one-to-all. In our classroom, Bob calls out, "Attention,"

everyone! It is break time, and there are free doughnuts in the lobby.” The word “everyone” means all people in the classroom are intended to receive this message.

Message types are more than just concepts. This is real traffic that every network-enabled device generates. Devices choose the type based on data they want to exchange and the protocol generating it. Later, you learn how traffic’s destination addresses define whether it is unicast, multicast or broadcast.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

TCP/IP Stack




By far, the most widely-used protocol stack is TCP/IP.

CONFIDENTIAL - FOR TRAINING

TCP/IP Stack

Conceptual model and set of communication protocols
Specifies how data is packetized, addressed, transmitted, routed, and received



MOD 1-29

The Internet Protocol (IP) suite is a conceptual model and set of communication protocols used in Internet and computing networks. It is commonly known as the TCP/IP stack because the foundational protocols in the suite are Transmission Control Protocol (TCP) and the IP. This model provides end-to-end data communication. It specifies how data should be packetized, addressed, transmitted, routed, and received. This functionality is organized into four abstraction layers.

The TCP/IP model is often compared to the OSI model, as shown in the figure. You see that the OSI model has seven layers, which map to the four layers of the TCP/IP model. As you can see the functionality of the OSI model's application, presentation, and session layers all map to a single TCP/IP Application layer.

This means that TCP/IP-based applications are responsible for interfacing with users and creating the data (application layer), putting it in the proper format (presentation layer), and managing sessions (session layer).

The transport layers of each model directly map. Recall that this layer is responsible for flow and error control. This layer establishes an end-to-end connection of two devices whose logical connection traverses a series of networks. The two choices used in a TCP/IP network are UDP and TCP.

The OSI Network Layer maps to the TCP/IP Internet layer. Both perform the same function - routing data across logical network paths by defining a packet format and an addressing

format.

IP version 4 (IPv4) and IP version 6 (IPv6) are the protocols that work on this layer.


The OSI model's data link and physical functions are all defined in a single layer in the TCP/IP model: the network interface layer. As you learned in the section about the OSI model, these layers contain protocols relating to the physical communication medium (copper wires, fiber optics, and radio waves). Recall that physical media access control also occurs here. Ethernet and Wi-Fi are very common protocols that work at this layer.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Ethernet

Family of computer network technologies used in Layer 1 and Layer 2

Layer 2 Ethernet uses MAC addresses – a 6-byte address assigned to each NIC



MOD 1-30

Ethernet is a family of computer networking technologies that are used at layer 1 and layer 2 of the OSI model. At layer 2, this standard provides for the assignment of 6-byte (48-bit) media access control (MAC) addresses to each device. Each manufacturer assigns a globally unique hexadecimal MAC address to each Network Interface Card (NIC). This address is stored in permanent read-only-memory (ROM) on the NIC. In other words, every Ethernet NIC in the world has a unique MAC address. How is this possible?

The most significant first 3 bytes of each MAC address are referred to as the manufacturer's Organizationally Unique Identifier (OUI). A standards body assigns each manufacturer in the world a unique 3-byte number. All NICs from that vendor therefore have the same number in the first 3 bytes.

The last 3 bytes of the MAC address comprise the NIC serial number. Recall that with hexadecimal a single byte is comprised of two digits. Thus, each MAC address is represented with 12 hexadecimal characters, each Byte or octet is typically separated by either a colon or dash.

For example, for the MAC address 8c:85:90:76:6c:95, the OUI is 8c:85:90, while the NIC serial number is 76:6c:95. As you progress through this course, you will often see fictitious, 2-byte MAC addresses such as 8c::95. This just makes it easier for you to focus on fundamental concepts, as opposed to always trying to read out large, 6-byte hexadecimal numbers.


To verify the vendor of a NIC by its MAC address, access a web site like <https://macvendors.com> , which allows you to perform OUI lookups. If you were to lookup the address in this example, you would determine that the manufacturer is Apple.

Note: Private IP address ranges are defined in RFC 1918. These are only to be used within an organization and never on the Internet. Likewise, there is a set of layer 2 MAC address ranges – any address where the most significant octet begins with is x2, x4, xA, or xE – that are private MAC addresses. Other than these private addresses, MAC addresses are globally unique. But understand that when a router receives a layer 3 IP packet, wrapped in a layer 2 frame, it strips off that frame, adds a new frame, and transmits it toward its destination. Thus, the requirement is that MAC addresses only need to be unique on the same layer 2 network (subnet/VLAN).

Also, newer Apple, Windows, and Linux platforms can pseudo-randomize their MAC address to improve privacy. These systems use the private MAC address ranges for this purpose.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Ethernet Frame



MOD 1- 31

Suppose that Host B receives data from the upper layers of the OSI model. Think of Ethernet as an employee of the upper layers. They have given Ethernet this payload, and asked Ethernet to get this over to host A. To do this, the Ethernet process must wrap this payload with a header and trailer, as shown in the figure. Each field in the header and trailer is described below:

Preamble. Informs the receiving system that a frame is starting and enables synchronization.

SFD (Start Frame Delimiter). Indicates that the destination MAC address field begins with the next byte. There might be several hosts on this Ethernet segment. The preamble and SFD are seen by every host as a kind of signal, "Hey, here comes a frame, and it might be for you!" How does the receiving station know whether the frame is for it, or some other host? By reading the next field: the destination MAC address.

Destination MAC(DMAC). Identifies the NIC of the receiving computing system. In this example Host A reads the MAC address and knows, "Hey, this is my MAC address – I should accept this frame". All other stations read the MAC address, see that the message is not for them, and drop the frame. This is like someone yelling in a room, "Hey Marta, can I buy you an ice cream?" Only Marta will respond.

Source MAC(SMAC). Identifies the NIC of the sending computing system. This is how Host A knows where the message came from, and therefore to whom it should respond.


EtherType. Defines the layer 3 protocol inside the frame, for example IPv4 or IPv6. Host A may be running an IPv4 protocol stack and an IPv6 protocol stack. The type field tells Host A which stack to pass the payload to for processing.

Payload. Contains the data that to be transmitted. FCS (frame check sequence). Allows detection of corrupted data using the Cyclic Redundancy check (CRC) or checksum method. As Host B creates this frame, it performs a mathematical calculation on most of the header and payload bits. It places the result of this calculation in the FCS field. When Host A receives the frame, it performs the same calculation. If it comes up with the same result that is stored in the FCS field, the frame is not corrupted. But if you have bad wiring, or strong electromagnetic interference (EMI) the data may become corrupted. In this case, Host A's calculation will not match the FCS, and so the frame is discarded – no need to accept bad data.

It is important to understand that the Ethernet payload not only carries application data – it also includes the header information from the upper layers. Remember, TCP adds a header, which IP considers part of the payload. Then IP adds a header, which Ethernet considers part of the payload. Let us look now at this IP header information.

Note: As promised, the figure only shows the first and last bytes of each 6-byte MAC address, to ease the learning experience.

IPv4 Header



MOD 1- 32

The figure shows the IPv4 header information added to a packet at layer 3, with a special emphasis on the most important fields in this header, as described below:

- **Time-to-Live (TTL):** This field mitigates Layer 3 route loops. If routers are misconfigured, IP packets could be bounced around between two or more routers forever. To prevent this, the originator places a number in this field from 0 to 255. For example, suppose that TTL:15. Each router that routes this packet decrements this field by 1. So, after this packet “hops” through the first router, TTL = 14. After the next router, TTL = 13, and so on. The router that decrements the TTL field to zero discards the packet.
- **Protocol:** This field indicates whether this IP packet carries TCP traffic or UDP traffic. Recall that application data is passed down to TCP or UDP, which adds a header and passes this on down to IP. From IP perspective, TCP, UDP, and other protocols are at higher layers in the model, and so can be collectively referred to as Upper Layer Protocols (ULP).
- **Source IP Address:** This originator of this IP packet adds its IP address to the source address field. In the figure, host 10.1.1.1 sends a packet to host 10.2.2.2. Thus, the Source IP Address field = 10.1.1.1, and the Destination IP address field = 10.2.2.2.
- **Destination IP Address:** The ultimate destination of the packet, as described above.

The other fields are shown here in gray, because they are not relevant to your learning in this


and several of the modules that follow. You will learn about these other fields as they become relevant. For more detail, please see RFC 791:
<https://www.rfc-editor.org/rfc/rfc791#page-11>

However, you may want to understand the variable IP header length – variable due to the Options field. The IP header length is specified in the Internet Header Length (IHL) field. Four bits specify the number of 32-bit words, with a minimum value of 5. So, $5 \times 32 = 160$ bits = 20 bytes long. If any of the options are used, the size might be larger. The maximum IHL value = 15, and $15 \times 32 = 480$ bits = 60 bytes.

Remember, this is the header that layer 3 routers use to do their job. They receive an inbound packet, analyze the destination IP address field, and forward packets along the best path toward that destination. In the figure, router R1 receives this packet from host 10.1.1.1 and forwards it to R2. R2 receives it and forwards it to destination host 10.2.2.2.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

TCP Three-way Handshake



MOD 1- 33

Recall that TCP acts as a kind of employee for applications, which rely on TCP to get data to some destination host. Before TCP processes any application data, it must first contact the target device and establish a reliable, flow-controlled connection.

In the figure, the Host B's HTTP application has data – sitting in a kind of out box, waiting to be transmitted to Host A.

It is as if HTTP says, “Hey TCP, I need you to get this data to Host A”.

TCP says, “OK boss, but first I must establish a connection with Host A to ensure that your data is sent in a controlled, reliable manner”.

There are four primary fields in the TCP header that are used to establish this connection – the sequence (SEQ) number, the acknowledgement (ACK) number, the synchronize (SYN) flag, and the acknowledgement (ACK) flag.

Host B starts by creating a TCP packet. This packet carries no data: it only signals Host A, “We should establish a connection”.

In this example, Host B sets SEQ = 1, and raises its SYN flag (sets it to a binary 1).

Host A receives this and knows, “This packet came from Host B with the SYN flag raised – Host B wants to connect”. Host A must inform Host B that it received SEQ = 1, and so

expects $SEQ = 2$ next.

To do this, Host A sends a packet with $ACK = 2$, and both the SYN and ACK flags are raised. Understand that Host A is independent of Host B. Host B tagged its packet with $SEQ = 1$. Host A chooses to tag its packet with $SEQ = 8$. Host A is saying, "I'm using my packet number 8 to acknowledge receipt of your packet number 1".


Host B receives and responds to this packet, "You expected $SEQ=2$ next, so here it is...I am using my packet number 2 to acknowledge receipt of your packet number 8. I expect packet number 9 next." As you can see, both the SYN and ACK flags are raised in this packet.

At this point, a reliable, flow-controlled connection is established. Each host has successfully received and acknowledged a sequence of packets from the other. Now TCP is ready to start transferring that data for the application.

HTTP uses a specific port to transmit the information; in this case the destination port number is 80; the source port in this case is randomly selected on Host B; in this case Host B selects the port 36,890 as the source port.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

TCP Sequence Numbers



MOD 1-34


TCP may be asked to transfer very large files – too big to transfer all at once. Therefore, Host B's TCP process breaks up the large file into smaller pieces called segments and assigns a unique sequence (SEQ) number to each byte segment. In this example you want to transmit the word "DATA"; Host B then creates 2 pieces, one for the letters "DA" and other for the letters "TA"

In the figure, TCP assigns Sequence = 3 to the first piece (DA) and transmits it to Host B. Host A acknowledges this by sending a packet with Acknowledgement (ACK) number = 4. Essentially, Host A is saying, "I received byte 3, and so I expect byte 4 next".

Host B then takes the next piece (TA), adds Sequence = 4, and sends it to Host A. Host A acknowledges this by sending ACK = 5. As Host A receives this data, it reassembles the pieces back into a complete file and passes it up to the application.

Note. TCP can send a block of information and Acknowledge receipt of data blocks using a single ACK message.

TCP Port Numbers




MOD 1-35

Recall that TCP acts as a kind of employee for applications. TCP receives requests from HTTP, FTP, and many other applications. It is as if the application says, “Hey TCP – establish a connection with the target host, and make sure my data gets there reliably”. TCP uses the destination port number to keep track of each application it serves.

To accommodate this, each application is assigned a standard port number. FTP uses port 20 for data transfer, HTTP uses port 80, and so on. When Host B’s TCP process accepts data from HTTP, it sets Destination Port = 80. Thus, when Host A receives this data, it knows to pass this data up to its HTTP application, as opposed to DNS or FTP.

In TCP you can use up to 65,535 ports; however, the well-known applications use the first 1,024 ports. From 1,025 to 65,525 ports are known as ephemeral port numbers. Applications use these ports to establish connections by randomly picking a port from that range and use that as source port for TCP session.

TCP Header



The figure shows the entire TCP header information added to a layer 4 segment, with a special emphasis on the most important fields in this header, as described below:

Destination port and source port

The first fields added are the source and destination port numbers. You just learned hosts use the destination port to distinguish the various applications that it serves. But what if you form two or more TCP connections? The source port is used to distinguish multiple TCP sessions so that data carried by there sessions go to the right application.

Sequence and acknowledgement numbers

You also learned about the sequence and acknowledgement numbers. They are used along with the ACK and SYN flags to ensure reliability and flow control. Host B sends SEQ=5, Host A responds with ACK=6, and so on. But what if Host A responded with ACK = 5? Host B says, "Hmm. I guess host A did not receive segment 5, so I will resend it". In this way, TCP ensures reliability, resending any packets that are lost in transmission.

Flags

You learned how the SYN and ACK flags are used to formally synchronize TCP packet (sequence) numbers and acknowledge receipt of packets. There is also a Final (FIN) flag, which tells the receiver, "Hey Host A, this is the last piece of the file."

There is the urgent (URG) and Push (PSH) flags to indicate that certain data is urgent and should be quickly pushed up to the receiving application. The URG flag is used in conjunction with the Urgent Pointer field, which indicates where the urgent data is located. There is also a Reset (RST) flag, used to reset the connection.

Window size


There's a checksum for reliability, and a window size field, used for flow control. If Host B is sending too much data at a time, Host A may not be able to handle the load. Host A may lower the window size, telling Host B, "Slow down, please do not send so much data at a time".

The remaining fields are not either rarely used or not relevant to this initial discussion. You will learn about these fields as you continue your training, as appropriate. For more information, see RFC 793:
<https://www.ietf.org/rfc/rfc793.txt>

Again, you may want to understand the Header Length field, to understand the overall number of bytes a TCP header adds. Like IP there is a minimum value of 5, and $5 \times 32 = 160$ bits = 20 bytes. Max Header Length value is 15, and $15 \times 32 = 480$ bits = 60 bytes.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

UDP Header




Disadvantages

- Connectionless
- No reliability
- No flow control

Advantages

- Small header size
- Low overhead
- Low delay



MOD 1- 37

The figure shows the UDP header information on a layer 4 segment. There are only source and destination ports, which you use exactly like TCP. There is a Length field and an optional checksum. For applications, TCP establishes a connection with the other side, retransmits lost packets, and makes sure not to overwhelm the receiver.

In contrast, UDP offers no connection, no reliability, and no flow control. Does the destination receive the packets? Is the data I'm transmitting overwhelming the network or destination host? UDP does not care. Understand that this does not mean that reliability is sacrificed: the upper-layer application implements the reliability it needs.

UDP is less sophisticated but comes with a small header size meaning a lower overhead. The lack of flags and flow-control services means UDP transfers data with lower delay. It is not reliable but is quick. This is perfect for applications like Voice over IP (VoIP) or video streaming. These applications do not need TCP's reliability. Programmers build the needed reliability into applications.

In UDP, you use up to 65,535 ports. Remember, well-known applications use the first 1,024 ports and ports 1,025 to 65,535 are ephemeral port numbers, same as TCP.



Networking Devices

A functional society is composed of good people.

A functional network is composed of network devices.

CONFIDENTIAL FOR TRAINING

Networking Devices: Switch


L2 device: Forwards Ethernet frames based on destination **MAC** addresses

Multiple ports connect endpoints

- PCs, printers, cameras, and more
- Commonly 8, 24, 48, or more physical ports
- Transparent to endpoints

Special L2 protocols

- Performance, reliability, security
- STP, LLDP, 802.1Q



MOD 1- 39

Switches have multiple ports that are used to connect computing devices into one or more local area networks (LAN). This can include PCs, printers, video cameras, voice-over-IP (VoIP) phones, and more. You may encounter switches that have 8, 24, 48, or more ports. Understand that switches are transparent to endpoints. Connected devices are not aware of the existence of the switch: they perceive themselves as being directly connected to each other.

A switch is a layer 2 network device that forwards Ethernet frames, based on destination MAC addresses. Recall that MAC addresses are part of the data link layer, and so switches are considered layer 2 network devices. The figure shows the Aruba icon used for a switch, and below that a physical switch example. You also see a simplified MAC address table that is maintained in switch memory. The switch has learned that a device with MAC address 90:::01 is connected to port 1 of the switch, and a device with MAC address 90:::03 is connected to port 2.

Switches also have special languages or protocols that they use to increase network performance, reliability, and security. Examples of these protocols include the following:

Spanning Tree Protocol (STP)

Link Layer Discovery Protocol (LLDP)

802.1Q

You will learn more about these protocols later in this course.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Networking Devices: Router


Layer 3 device: Forwards packets based on destination IP address

Functionality and features

- Connect separate networks into an inter-network
- Offer WAN connectivity to networks

L3 routing protocols

- Learn all possible paths, choose a best path
- RIP, OSPF, BGP



MOD 1-40

A router is a layer 3 network device that forwards packets based on the destination IP address. Since IP addresses are part of the network layer, routers are considered layer 3 devices.

You know that switches connect computing devices into one or more networks. It is a router's job to connect those separate networks together, to create an inter-network. The Internet is the largest, most common example of an inter-network. The figure shows four routers, used to interconnect five networks (links between the routers are also considered networks).

While switches might have a few dozen or a couple hundred ports, routers have a relatively small port count – often between 2 and 6 – which are used to connect to WAN networks.

When you take a long trip in your car, there are often several paths you could take to arrive at your destination. You consider all possible paths, and then choose the best path. You judge the best path based on some criteria that is important to you, such as whether the route is the fastest, or the shortest, or whether the route is particularly scenic. Then you drive along that path. Similarly, routers run protocols between them to learn all possible paths between all available networks. They then choose a best path for each destination, and forward packets along those best paths.

Examples of these routing protocols are:

Routing Information Protocol (RIP)

Open Shortest Path First (OSPF)

Border Gateway Protocol (BGP)

You will learn more about these protocols later in this course.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Networking Devices: Multi-Layer Switch

Performs both L2 switching and L3 routing

L3 Routing

- Internal routing functionality
- Runs RIP, OSPF, BGP

L2 Switching

- Multiple ports connect endpoints
- Uses STP, LLDP, VLANs



MOD 1- 41

The OSI model sets clear distinctions between a layer 2 switch and a layer 3 router. However, networking professionals know that it is often good to have a single device that can perform both functions - a hybrid between a router and a switch.

A multi-layer switch has all the functionality of a layer 2 switch.

They support typical layer 2 functionality, such as STP, LLDP, and VLANs. Multiple ports connect various endpoint devices together into one or more networks.

A multi-layer also has all or most of the functionality of a layer 3 router, so it can route between those separate LANs internally - an internetwork in a box! Like any router, it supports routing protocols like RIP, OSPF, and BGP.

Of course, multi-layer switches support layer 1 to actually transmit and receive data, and it can even perform some layer 4 functions – especially as relates to certain security features. These devices help to build a secure and flexible network.

Networking Devices: Access Points

Enables wireless users to access wired resources and roam about

Functionality

- Bridges wireless devices and wired networks.
- Transform Ethernet into Wi-Fi frames and back



Varieties of APs

- Internal or external antennas
- Single or Dual Ethernet Ports
- Indoor or outdoor



MOD 1-42

Access Points (AP) enable wireless users to access wired resources. Without being tethered to Ethernet cables, users are free to roam about.

APs are translational bridges: they accept Wi-Fi frames from endpoints, translate them into Ethernet frames, and forward them on to some wired resource. This could be an internal corporate server, storage, application, or an Internet resource. When that wired resource responds, the AP accepts that Ethernet frame, translates it to a Wi-Fi frame, and then transmits it over the air to the wireless host, as shown in the figure.

There are many varieties of APs to meet your technical and budgetary needs. They may be built with internal or external antennas, with one or dual Ethernet ports. Some may be for indoor use only, while others may can be mounted outside.

Wi-Fi systems may be designed to use Autonomous APs or Controller-based AP. Autonomous APs can perform all necessary functions to create a functional system. They process wireless and ethernet frames and provide a certain amount of manageability and control. As the name implies, they work autonomously, without control from some external device. This type of deployment is relatively simple and easy, but it is not very scalable. It is suitable for smaller environments.

With Controller-based AP solutions, the APs send and receive Wireless and Ethernet frames, much like Autonomous APs. However, much of the processing and management functions are off-loaded to one or more centralized controllers.

With Cloud-based APs, the APs is managed in a similar way to Controller-based solution, but there is no controller. Instead, all management and monitoring functions are offloaded to software that runs as a cloud service like Aruba Central.

These centralized configuration and control features might increase the initial complexity of the deployment, but once in place they provide far superior visibility into network health, enable more proactive network management, and can scale from medium to very large deployments.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Networking Devices: Firewall


L2-7 filtering and security

Functionality

- Block unauthorized/inappropriate access
- Permit authorized/appropriate access

Types

- Stateless vs Stateful
- Functionality built into switches, gateways



MOD 1- 43

A firewall is a security device. It denies unauthorized access and permits appropriate, authorized access.

Firewall functionality may operate at nearly any layer, from layer 2 to layer 7. Firewalls might permit or deny access based on a simple Layer 2 MAC address or Ether Type filter, or based on a layer 3 IP address/mask filter. It might also look at certain fields in the layer 4 headers: TCP, UDP, etc.

Some firewalls have the intelligence to be application-aware. They understand how certain applications are supposed to work and block any nefarious activity that seeks to circumvent this normal operation. Or it might allow a particular user to access a corporate-authorized application on the web – perhaps a vendor catalog, but deny access to inappropriate applications and/or web sites.

Stateless firewalls lack any of this intelligence. They simply compare each layer 2, 3, and/or 4 header information to access control lists of permitted and denied traffic. Stateful firewalls are aware of normal TCP/IP communications. For example, they understand that all TCP connections begin with a three-way handshake, as you just learned. If this procedure is not followed, the firewall intelligently denies traffic.

Aruba builds firewall features into many of their products: switches, multi-layer switches, and especially into their gateway and AP products.



Knowledge Check

Self-check on key learning points



Its time for a knowledge check.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Question #1

Which of the statements below accurately describe basic network concepts? (Select two.)

- a. A computing network is a group of computing resources that permit digital data exchange between routers.
- b. A LAN is a group of computing devices in the same place, all interconnected via a router.
- c. A WAN is a group of computer resources that can communicate over large geographical distances.
- d. A protocol is a set of rules that computer devices follow to establish and maintain communications.

a


Knowledge Check

MOD 1-45

C,D

Question #2

Match the OSI reference model layer to an appropriate description.



Knowledge Check



MOD 1-46

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Question #3

The layer 2 protocol data unit (PDU) includes the headers for layers 3 through 7.

- True
- False



Knowledge Check

MOD 1-47

T

Question #4

Which of the statements below accurately describe physical media? (Select two.)

- a. Copper cabling variants include shielded twisted pair (STP) and unshielded twisted pair (UTP).
- b. Multimode (MM) fiber optic cable can typically handle longer distance transmissions at a lower cost than single mode (SM) fiber.
- c. You typically use direct-attached copper (DAC) cables to directly connect two servers or other end systems to each other.
- d. Wi-Fi devices transmit data by modifying aspects of radio waves, such as amplitude, phase, or frequency.
- e. Multicast traffic is used to transmit data from one host to all other hosts on a LAN.

a


Knowledge Check

MOD 1-48

A,D

Question #5

Match the TCP/IP stack layer to an appropriate description.



Knowledge Check



MOD 1-49

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Question #6

Your colleague is new to networking and asks you about network devices. Which statements below would help your colleague understand this topic? (Select two.)

- a. You use switches to connect many end systems together, at layer 3 of the OSI model.
- b. Routers are often used in WAN environments, and forward packets based on IP address information.
- c. Multi-layer switches perform functions of a firewall and a router.
- d. Stateful firewalls have more intelligence than stateless firewalls, perhaps including application awareness features.



Knowledge Check

MOD 1- 50

B,D

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY



Lab Activity

Consider the lab activity.

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

Lab 01.02: Networking Fundamentals

Task 2.1

- Analyze headers and encapsulation with a packet analyzer

Task 2.2

- Compare UDP vs TCP headers

Modul 2 – AOS-CX Switch Fundamentals



MOD 1- 52

The figure provides a brief review of lab tasks. Please see your lab guide for details. When complete, you will continue with Module 2 – AOS-CX Switch Fundamentals.

Rev. 22.41

aruba
a Hewlett Packard
Enterprise company


AOS-CX Switching Fundamentals

Aruba Campus Access
Fundamentals



Welcome back everyone! This is Module 2 – AOS-CX Switching Fundamentals.

CONFIDENTIAL - FOR TRAINING PURPOSES



Objectives

- Connect to and access a switch
- Perform initial switch setup
- Describe and configure VLANs, tagging, and IP addressing
- Use LLDP and ICMP for network discovery and diagnosis
- Configure link aggregation to improve performance/resiliency

This module lays the foundation on which much of your future studies relies. After you take that new switch out of the box and connect it, then what?

After completing this module, you will be able to connect to and access switches and perform initial switch setup and validation.

Then you configure VLANs and make certain ports members of those VLANs, configure some ports as trunk ports, and configure other ports as layer 3 ports with IP addressing.


Now that the network is functional, you can use LLDP and ICMP for network discovery, documentation, and diagnosis. Then you will learn to configure link aggregation to improve performance and resiliency.

Initial Switch Setup

Let's begin with accessing and performing the initial setup of the AOS-CX switches.

CONFIDENTIAL - FOR TRAINING

Out of Band Management Network




In-band management (left-side of the figure) is when you remotely connect to and manage switches via a normal Ethernet port. This seems convenient; in certain environments like data centers, it could be a bad idea. A switch or router misconfiguration breaks your administrative access. There is also a security concern. Without tight security controls, hackers can access and control your devices.

Out-of-band (right-side of the figure) management (OOBM) uses alternate, dedicated connections to network devices, separate from the end user data network.

You can monitor and configure network devices without relying on the production network. Plus, you establish a trust boundary, with only one entry point for the management interface, it is far more difficult for bad actors to attack you. But how do you make this connection?

Accessing the Serial Console Port: OOBM



You typically perform initial switch configuration using out-of-band management (OOBM) on a console port for the initial switch configuration. This special port exists on all switch models and types, to facilitate configuration, troubleshooting, and management.

Depending on the switch model, the console can be a USB-C port or an RJ-45 port, however, for both port types you must establish a serial connection between the management station (perhaps your laptop) and the switch. To do this, you need software on your laptop or management station that emulates a serial session. PuTTY, TeraTerm, and SecureCRT are common examples. Make sure to configure your terminal emulation software with the parameters shown in the figure.


If the switch has an RJ-45 port, you also need the serial cable that ships with the switch, and a USB-to-serial converter - modern laptops no longer include a serial port. If the switch has a USB-C port, you need a USB-C cable – connected directly from the switch to your laptop. All AOS-CX switches also ship with a Bluetooth USB dongle, to enable easy initial setup with the CX mobile app. For more information, see:

<https://community.arubanetworks.com/community-home/digestviewer/viewthread?MID=35150>. If the URL changes do an Internet search on “Howto: Using the Bluetooth USB Dongle for a Console Connection on a CX Switch”.

Management (mgmt.) Port



SSH • HTTPS • SNMP • NTP



- Exclusively for monitoring and management
- Complete isolation from data network
- Requires an IP address

- Common in data centers, or for critical core devices
- Not typical for campus access layer devices



MOD 2-6

Another way to make an OOBM connection is via the special Ethernet Management. All OS-CX switches include an out-of-band Management (mgmt) port to completely isolate management traffic from all other traffic. This isolation is achieved using VRF (Virtual Routing and Forwarding), which is like a dedicated IP network stack separate from the default stack. So, management traffic does not use bandwidth in the data links, and data traffic cannot contact this interface. You must assign an IP address to this interface.

Also note that many access switches have a so-called switch virtual interface (SVI) enabled as a DHCP client, to support automatic connections to Aruba Central for Cloud management. You will soon learn more about SVIs. Aruba tools like NetEdit, AirWave, and Central are great for centralized management of an existing network and complement CLI access. The common use case is in a datacenter, or for critical core devices.

Interface Numbering for Fixed-Port Switches

- Used with VSF
- Indicates cluster member ID
- Defaults to 1

- Used in modular switches
- Indicates line card number
- Fixed port switches use 1

- Port number

Member / Slot / Port

Interface 1/1/48

Member ID = 1 Slot = 1 Port = 48



VSF member 1

Interface 2/1/44

Member ID = 2 Slot = 1 Port = 44



VSF member 2

VSF

MOD 2- 7

AOS-CX references interfaces following the member/slot/port notation.

Member.

When the Aruba Virtual Switching Framework (VSF) feature is running multiple switches can be seen as one virtual or logical switch; this number indicates the member ID in the cluster. By default, an AOS-CX switch will be member 1 indicating that these features are not running.

Slot.

In modular switches like the 8400 and 6400 models, this number represents the slot being used by a particular line card. For fixed switches (8320, 8325, 6300 and 6200 for example) this number will be always 1.

Port.

This number refers to the individual interface in the line card (modular switches) or in the chassis (fixed port switches). Port numbers always start at 1 on AOS-CX switches.

The image shows an example with a fixed port switch with VSF enabled. If VSF is disabled all switches will have default member ID 1.

Note. VSF is available for 6300 and 6200 switch series.



Port Numbering for Chassis-Based Switches

Member/Slot/Port

- Used with VSF only
- Indicates cluster member ID
- Defaults to 1

- Used in modular switches
- Indicates line card number
- Fixed switches use 1

- Line card/chassis Interface number



MOD 2-8

Here's a couple of typical Aruba OS-CX switches – lots of Ethernet ports right? How do you identify them? You use the typical member/slot/port notation.

The member number is only used with a special feature - the Aruba Virtual Switching Framework (VSF), which makes multiple switches appear as one virtual switch. The member number indicates a physical switch's unique member ID. The default member ID is 1, indicating that the feature is not running.

In modular switches like 8400 and 6400, the slot number represents the slot a particular line card uses. For fixed switches like 8320, 8325 and 6300, this number is always 1. And the port number refers to the individual interface in the line card for modular switches, or in the chassis for fixed switches.

Look at this fixed-port switch in the top-right corner of the figure. It does not run VSF, so the Member ID is always 1. It doesn't have slots because it is a fixed-port switch, so the slot number is always 1. Then the individual ports are numbered, from port 1 at the upper left, to port 48 at the lower right. So, you refer to this last port as interface 1/1/48.

This modular switch in bottom-right of the figure does not run VSF, so its member ID is always 1. But modular switches have slots, into which you install special line cards. The third slot in from the left would be identified as slot three. The individual ports are numbered on each card; therefore, the port highlighted for the switch at the bottom right is port 24. Thus, you refer to it as interface 1/3/24.

Now that you know how to refer to switch ports, let's move on to switch configuration.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

AOS-CX Configuration Prompt

Default hostname is Switch
You soon learn to change it

Symbol indicates "context"
What you can view/configure



MOD 2-9

This section covers the command line prompts for the AOS-CX operating system.

Default prompt.

When you access the ArubaOS-CX switch Command Line Interface (CLI) for the first time, the prompt in the terminal application looks like that shown in the figure. The word Switch is the device's default hostname, which you soon learn how to change.

The CLI is organized into various contexts or levels, as indicated by the symbol following the hostname. Each context determines which parts of the switch you manage, and which commands are available to users. Let's explore!

AOS-CX Prompt Modes: Overview

Operator

Switch>

Some view commands, no config

Manager

Switch#

Execute, view commands, no config

Global

Switch(config)#

Global config commands

Interface

Switch(config-if)#

Sub-context config commands



MOD 2- 10

Prompt modes: Overview


The operating system defines the contexts shown in the figure. The operator context level is indicated by a greater-than (>) symbol. When you see this, know that you can run a sub-set of show or view commands, but cannot change configuration.

The manager context has the hashtag symbol, where you can run nearly all show commands to view status and configuration, but you cannot modify that configuration.

The global configuration context ends with the “(config)#” in the prompt. This is where you run commands that change the switch’s configuration. These commands have a global effect on the device itself. But to configure specific components or features, you must access a sub-config mode.

What you see in the prompt depends on what you are configuring. This example is config-if – short for configure interface. You will see different prompts when you configure other aspects of the device. You’ll see many sub-contexts throughout this course. Mainly, you must understand how to navigate through these contexts. As you’ll soon see, its easy!

AOS-CX Prompt Modes: Move to Other Modes



When you log in with manager privileges, automatically in manager mode
No need to use **enable** command.

MOD 2-11

Prompt modes: Move to other modes


You begin as a lowly operator. To access manager mode, you enter the enable command. Your prompt changes from > to # (see figure).

To move from manager to global config context, enter configure terminal. You see the “config” prompt displayed. To move from global to the interface context, use interface followed by the port number.

You cannot move directly from operator to global or Interface using a single command. You must progress in the order shown in the figure.

Note: When you log in with manager privileges, the AOS-CX switches automatically places you in manager mode. Thus, there is no need to use the enable command.

AOS-CX Prompt Modes: Return to Previous Modes



MOD 2- 12

Prompt modes: Return to previous modes

To move from the interface level to the global level, use `exit`. You just learned that you cannot skip a level moving down from operator to manager to global and so on. But you can move directly from any sub-config level to the manager context with the `end` command.

So, you can move from the Global to Manager level with either `exit` or `end`. You can also use `Ctrl-z` instead of `end`. To move from the manager to the operator level, use `disable`. To move from the interface to the manager level, use `end`. And from the operator level, if you want to log out of the switch, use the `exit` command.

AOS-CX: Context Sensitive Help

Show available commands in the current context

```
Switch# ?  
boot                Reboot all or part of the system  
checkpoint          Checkpoint information  
<<<Omitted Output>>>  
Switch#
```

Show the available parameters for a command

```
Switch# show access-list ?  
all                All access-lists  
ip                 Internet Protocol v4 (IPv4)  
ipv6               Internet Protocol v6 (IPv6)  
log-timer          Set ACL log timer length (frequency)  
mac                Ethernet MAC Protocol  
Switch(config)# access-list
```



MOD 2-13

The AOS-CX CLI provides help features. To see the available commands that you can run in the current context, enter the question mark (?) symbol. The question mark does not display on the screen when you enter it.

Or perhaps you know you want to use some variation of the show access-list command, but you are not sure what options you can specify. Just enter the command followed by a space and a question mark. So, if you wanted to see all access lists configured on the device, you would enter show access-list all, and so on. Pretty cool!

AOS-CX: Command Abbreviation and Completion

Command Abbreviation

Enter enough letters to uniquely specify a command

```
Switch# conf  
Switch(config)#
```

Command Completion

Enter partial command, press TAB, CLI completes it

Press TAB twice to see command options

```
Switch(config)# int + Tab  
Switch(config)# interface  
  
Switch(config)# interface + Tab + Tab  
lag      loopback  mgmt      tunnel
```



MOD 2- 14

AOS-CX supports command abbreviation – just enter enough letters to uniquely specify any command, and the CLI accepts it. For example, you can simply enter `conf` instead of typing `configure terminal` to access global config mode.

The CLI also supports command completion. If you enter enough letters of a command to uniquely identify it, then press `<Tab>`, the CLI completes the command for you. If you have not entered enough letters to match a command, the CLI does not complete it.

In this example you have entered `int`, then pressed the `<Tab>` key. The CLI completed the command for you, so you see `interface` spelled out.

If you press `<Tab>` twice after a word, the CLI displays command options. For example, enter `interface`, then `<Tab>` `<Tab>`. You see options that you can enter after the `interface` command.

Show version

Version info for network OS, service OS, and BIOS

Switch# **show version**

ArubaOS-CX

(c) Copyright 2017-2022 Hewlett Packard Enterprise Development LP

Version : **FL.10.09.1031**
Build Date : 2022-06-14 04:21:23 UTC
Build ID : ArubaOS-CX:FL.10.09.1031:dea0ef74d35b:202206140343
Build SHA : dea0ef74d35b187ab79b41ffba4d68ec38cb5a63
Active Image : primary

Service OS Version : FL.01.09.0002
BIOS Version : FL.01.0002



MOD 2- 15

The show version command displays version information about the network operating system software, service operating system software and BIOS. In the figure, you see that this system is running the AOS-CX, version GL.10.04.0003. Near the bottom you see the Service OS and BIOS versions. This is useful for your general knowledge, and to help to determine if upgrades are needed.

Command context: operator (>) or manager (#)

Authority: operators or administrators.

Show images

Version info for network OS, service OS, and BIOS

```
Switch# show images
-----
ArubaOS-CX Primary Image
-----
Version : FL.10.09.1031
Size    : 872 MB
Date    : 2022-06-14 04:21:23 UTC
SHA-256 : e91e7flade68ab54cd00918673501d6067cdb428d3f4bf6f2beb31f5f8ef322
-----
ArubaOS-CX Secondary Image
-----
Version : FL.10.05.0021
Size    : 642 MB
Date    : 2020-10-29 10:36:02 PDT
SHA-256 : 4c795e8c9eec5952645ded19cf9a2018deb545c7ed0221f32a1a5bd0d64ee5f6
Default Image : primary
Boot Profile Timeout : 5 seconds
```

MOD 2-16

ArubaOS-CX switches have two flash memory modules or partitions for storing switch software image files. AOS-CX refers to those modules as primary and secondary images. Having two flash memory locations helps to manage the switch more easily, and to help with planning firmware upgrades. For example, you can download a new firmware code in the flash memory that is not currently in use and then instruct the switch to boot from it during a maintenance window.

The show images command shows information about the software in the primary and secondary images.

Show system

General system status info

```
Switch# show system
Hostname           : Switch
System Description : GL.10.04.0003
System Contact     :
System Location    :

Vendor            : Aruba
Product Name      : JL635A Aruba 8325-48Y8C 48p 25G 8p 100G swch
Chassis Serial Nbr : TW98KM003P
Base MAC Address   : 9020c2-bced00
AOS-CX Version     : GL.10.04.0003

Time Zone         : UTC
Up Time           : 1 week, 1 day, 11 hours, 20 minutes
CPU Util (%)      : 4
Memory Usage (%)  : 19
```



MOD 2-17

The show system command displays general status information about the system. During remote access, the device platform and version information may not be obvious. In the example shown, you are attached to device named “Switch”, which is an Aruba model 8325-48YC. It has 48x25gbps ports, and 8x100Gbps ports. You see the serial number, base MAC address, and uptime, as well as CPU and memory utilization. This can give you some idea of the general health of the system. By baselining these parameters, you can begin to establish what normal is for your network.

Command context: operator (>) or manager (#)

Authority: operators or administrators

Show running-config

Current configuration, no user info

```
Switch# show running-config ← Add all keyword to also see default values
Current configuration:
!
!Version AOS-CX GL.10.04.0003
!export-password: default
Hostname Switch
!
Router ospf 1
    area 0.0.0.0
Vlan 1,1101,1103,1105,1107
Interface mgmt
    no shutdown
    ip static 10.250.1.3/24
    default-gateway 10.250.1.254
Interface 1/1/29
    no shutdown
    description To-8325-Core-Table01-Port47
<<<Omitted output>>>
```

MOD 2-18

Switch configuration files.

This section introduces you to viewing and manipulating switch configuration files.


The show running-config command displays the current configuration running on the switch. No user information is displayed.

Command context: manager (#)

Authority: administrators

Show running-config all includes configuration commands and values for all features, including the default configuration that is normally not visible in the configuration output.

Configuration File Management



MOD 2- 19

An ArubaOS-CX switch uses RAM and flash memory to save configuration files. You should be aware of RAM and flash memory characteristics, as they relate to the storage of device configuration files.

RAM memory is used to save the configuration that is currently running on the device: the running-config. Remember that RAM is not permanent storage; the contents of RAM are lost when a device reboots or powers down. If you modify the running-config and the switch reboots, your changes are lost. Use the command `show running-config` to view this file.

Flash memory is used to permanently save a configuration file. The contents of flash memory survive during reboot or power off. Use the `show startup-config` command to verify the contents on this file.

Of course, any time you make changes to the running-config, you should copy those changes to the startup-config using the `copy` command. To permanently save the configuration that is running in the device, use the command `copy running-config startup-config`.

Note. Alternatively, you can use the command `write memory` which does the exact same thing as `copy running-config startup-config` command.

If you need to restore a configuration from the startup-config file, use the command `copy startup-config running-config`.

As a best practice, you should export the running and the startup configuration files to an external file server. This will help you to recover from catastrophic device failures or administrative mistakes. To accomplish this task, you can use the copy command, as shown below.

```
Switch# copy {startup-config | running-config}  
{tftp://|sftp://USER@}{IP|HOST}[:PORT][;blocksize=VAL]/FILE
```

For example, to copy the running-config to a TFTP server with IP address 10.253.1.21, use the command copy running-config tftp://10.253.1.21/switch_config.cfg

ArubaOS-CX switches also allow you to copy files to a USB flash drive, which can be directly connected to the switch. To copy your running or startup-config to a USB device, use the following syntax:

```
Switch# copy {running-config | startup-config} usb:/file
```

ArubaOS-CX considers startup and running config checkpoints, you will learn about checkpoints in the next section.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Checkpoint Overview

Features

- Configuration snapshot with relevant metadata
- Create working stages
- Export to other switches of same model
- Running-config and startup-config are considered checkpoints

Two types

- System Generated
- User Generated

Generated after 5 minutes of inactivity after a configuration change
Format CP<YYYYMMDDHHMMSS>



MOD 2- 20

Using the copy commands that you just learned allows you to store configuration files and recover lost configurations in case of a power outage or administrative mistake. This is a good thing, but these files do not include any additional data about the state of networking processes. For a true recovery, a new approach is needed: checkpoints.

A checkpoint is a snapshot of the switch running configuration and its relevant metadata during the time of creation. You can use checkpoints to apply the switch configuration stored within a checkpoint whenever needed, such as revert to a previous working configuration. A checkpoint is flexible since it can be applied to other switches of the same model.

There are two types of checkpoints:

System generated checkpoints are generated after 5 minutes of inactivity after a configuration change. By default, a checkpoint is taken every 300 seconds or 5 minutes. The format that this checkpoint type uses is CP<YYYYMMDDHHMMSS>. For example, CPC20200111000038

User generated checkpoints are manually generated by the administrator and can be assigned whatever name the administrator deems relevant at the time of creation.

Checkpoint Configuration

Create manual checkpoint

```
Switch# copy {running-config | startup-config} checkpoint <checkpoint-name>
```

Verify checkpoint list

```
Switch# show checkpoint list all
```

Copy checkpoint

```
Switch# copy checkpoint <checkpoint-name> {running-config | startup-config  
| <external-file-server>}
```

Restore to running-config

```
Switch# checkpoint rollback <checkpoint-name>
```



MOD 2- 21

Checkpoint configuration

To create a checkpoint, use the copy command.

```
Switch# copy {running-config | startup-config} checkpoint <checkpoint-name>
```

To verify the checkpoint list, use the command show checkpoint list all.

To revert the switch's configuration to a previously saved checkpoint use the following command:

```
Switch# copy checkpoint <checkpoint-name> {running-config | startup-config}
```

Note. You can export a checkpoint file to an external server.

Alternatively, you can use the checkpoint rollback command. This command only restores the configuration to the running-config file.

```
Switch# checkpoint rollback <checkpoint-name>
```

Note. The checkpoint rollback command performs the same task as copy checkpoint running-config command.

Checkpoint Auto Mode

1. Administrator enables checkpoint auto mode.

```
Switch# checkpoint auto <minutes>
```


2. The system generates a checkpoint and starts a timer.

3. Administrator configures the device.

4. Administrator must confirm the changes to be permanent.

```
Switch# checkpoint auto confirm
```

If no confirmation received, system rolls back to checkpoint.



MOD 2- 22

Checkpoint auto mode

There are situations when applying a configuration could lock out your access to the switch. This leaves no other option but to locally connect to the device using the console port and remove the offending configuration lines.

The checkpoint auto mode feature addresses this problem in an elegant way. The solution includes a timer that you set. The system then creates a checkpoint with the current state. You then make your desired configuration changes, which the system applies, after a confirmation. If the configuration lines did not affect your access, you will be able to confirm the change. If your changes locked you out of the switch, confirmation does not occur. The switch thus automatically reverts to the latest snapshot after the time expires.

To use the checkpoint auto mode, issue the following configuration.

```
Switch# checkpoint auto <minutes>
```

```
Switch# checkpoint auto confirm
```

Important: Do not execute the checkpoint auto command until you have verified that your configuration changes are functioning correctly.

Top CPU

CPU utilization info

Switch# **top cpu**

top - 20:40:46 up 17 days, 23:45, 1 user, load average: 0.32, 0.26, 0.24

Tasks: 238 total, 1 running, 194 sleeping, 0 stopped, 0 zombie

%Cpu(s): 2.5 us, 1.6 sy, 0.0 ni, 95.9 id, 0.0 wa, 0.0 hi, 0.0 si, 0.0 st

KiB Mem : 7868164 total, 3013360 free, 13999328 used, 3455476 buff/cache

KiB Swap: 2097148 total, 2095868 free, 1280 used. 6079420 avail Mem

PID	USER	PR	NI	VIRT	RES	SHR	S	%CPU	%MEM	TIME+	COMMAND
26308	admin	20	0	3344	2064	1612	R	11.8	0.0	0:00.03	top -b -n 1 -s -c -o
629	root	20	0	33928	17084	6124	S	5.9	0.2	777:57.14	/usr/sbin/ovsdb-server
3047	root	20	0	791184	84416	32232	S	5.9	1.1	189:43.13	/usr/bin/prometheus
1	root	20	0	7600	6312	2960	S	0.0	0.1	0:28.14	/sbin/init



MOD 2- 23

Switch operation verification

The next three figures introduce you to three commands that you can use to verify the operation of your AOS-CX switch.

Top CPU.

The top cpu command shows detailed CPU utilization information for the AOS-CX switch. You learned how to see high-level CPU utilization with the command show system. The top cpu command provides more detailed information.

Command context: operator (>) or manager (#)

Authority: operators or administrators

Likewise, the top memory command shows memory utilization information for processes running on the switch.

Show events

Event logs since last reboot

```
Switch# show events -r
-----
Event logs from current boot
-----
2019-12-06T10:24:18.679695+00:00 Spine-8325-C hpe-sysmond[3212]:
Event|6303|LOG_INFO|AMM|1/1|Current system memory usage is 27%
2019-12-06T09:09:19.067316+00:00 Spine-8325-C lldpd[2165]:
Event|106|LOG_INFO|AMM|1/1|LLDP neighbor 90:20:c2:bc:7c:00 deleted
on 1/1/35
```

Use **-r** parameter to list most recent events first

MOD 2- 24

The show events command displays the event logs generated by the switch modules since the last reboot. Log analysis is a powerful tool to investigate and troubleshoot system and protocol-related problems. Use the **-r** parameter to list the most recent log events first (ordered from the newest to oldest events).

Command Context: manager (#)

Authority: administrators or auditors

Note: The show event command has other parameters that you can use like **-e**, **-s**, **-a**, **-n**, **-c** and **-d**. You can learn more about how to use these parameters in the Command-Line Interface Guide for AOS-CX document.

Note: This event is a file on the switch SSD drive, there is option to see events before last reboot.

Show interface brief

Available interfaces and status

```
Switch# show interface brief
```

Port	Native VLAN	Mode	Type	Enabled	Status	Reason	Speed (Mb/s)	Description
1/1/1	1	routed	SFP28DAC3	yes	up		25000	--
1/1/2	1	trunk	SFP28DAC3	yes	up		25000	--
1/1/3	1	trunk	SFP28DAC3	yes	up		25000	--

<<<Omitted Output>>>

Command output filtering

```
Switch# show interface brief | include up
```

1/1/1	11	access	1GbT	yes	up		1000	pc1
1/1/2	4	trunk	1GbT	yes	up		1000	ap1
1/1/3	1	access	1GbT	yes	up		1000	--
1/1/26	1	access	SFP+DAC1	yes	up		10000	--
1/1/27	1	trunk	SFP+DAC1	yes	up		10000	--
1/1/28	1	trunk	SFP+DAC1	yes	up		10000	--

MOD 2- 25

The show interface brief command helps you to see what interfaces are available and the current status of these interfaces. This command also briefly shows layer 2 and layer 3 configurations applied to the interface. This is a frequently used command for many network engineers. At a glance, you can tell the status of every interface on the entire device, with a specific focus on the Enabled and Status columns.

In the Enabled column, yes means that the interface is not disabled in the switch configuration – no administrator has disabled the interface. up in the status column means that something is attached to that interface, and there is at least a good layer 1 and layer 2 connection between them.

Command context: operator (>) or manager (#)

Authority: operators or administrators

The pipe (|) command is a CLI session command that filters the output of show commands according to the criteria specified by the parameter include, exclude, count, begin, or redirect.

The pipe (|) command is supported for use with the show command only.

You can use multiple pipe commands with a single show command.

You can use the pipe command with the page command.

Command completion by pressing the <Tab> key is not supported for pipe commands.

This is an example where you see every line in the output of the running-config file that has the word “vlan” in it, with the exceptions of those output lines that contain the word “vlan2”.

```
Switch# show running-config | include "vlan" | exclude "vlan2" | count
```

The count parameter causes the switch to display a count (the number of occurrences) instead of the actual lines that match the criteria in the pipes.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Configuring a Hostname

Name devices for easy identification

Command context: **Config**

Syntax: **hostname** <Hostname>

```
Switch(config)# hostname Access-1  
Access-1(config)#
```

MOD 2-26

On Aruba switches configuration becomes active immediately after entering the command

Configuring a hostname

You should name each network device. This ensures that you can easily identify them, especially in large networks. In a real environment this is one of the first configuration steps that you will do.

Many corporations have documented, standardized naming conventions. For example, one standard might be Building-Floor-Rack_Name. For example, the switches in rack 2, on the 2nd floor of building 9 might, be named 9-2-2_SWA, 9-2-2_SWB, 9-2-2_SWC, and so on.

The figure shows how to name devices using the hostname command.

Command context: Config

Syntax: hostname <Hostname>

Example:

```
Switch(config)# hostname Access-1
```

```
Access-1(config)#
```

Management Port

Features

- Dedicated for the ArubaOS-CX *mgmt* interface
- DHCP is enabled by default
- Static configuration is supported

Verification

```
Switch-1# show interface mgmt
  Address Mode           : static
  Admin State            : up
  Mac Address            : 88:3a:30:97:c7:01
  IPv4 address/subnet-mask : 10.251.1.4/24
  Default gateway IPv4    : 10.251.1.254
  Primary Nameserver      : 10.253.1.21
<<<Omitted output>>>>
```

Configuration

```
Switch-1 (config)# interface mgmt
Switch-1 (config-if-mgmt)# ip static <IP-address/Mask>
Switch-1 (config-if-mgmt)# default-gateway <Default-gateway-IP>
Switch-1 (config-if-mgmt)# nameserver <DNS-Primary-server>
```



MOD 2- 27

The ArubaOS-CX mgmt interface is enabled by default. It receives its IP parameters using DHCP. Use the show interface mgmt command to verify the interface's state.

If you prefer, manually configure IP parameters on the management interface, as shown in the figure.(DHCP is the default setting for the interface addressing).

```
Switch(config)# interface mgmt
```

```
Switch(config-if-mgmt)# ip static <IP-address/Mask>
```

```
Switch(config-if-mgmt)# default-gateway <default-gateway-IP>
```

```
Switch(config-if-mgmt)# nameserver <DNS-primary-server>
```


Managing Interfaces

Enabling Interface

```
Switch(config)# interface 1/1/1  
Switch(config-if)# no shutdown
```

Disabling Interface

```
Switch(config)# interface 1/1/1  
Switch(config-if)# shutdown
```

Interface description

```
Switch(config)# interface 1/1/1  
Switch(config-if)# description To-Core
```

Interface config state

```
Switch# show run interface 1/1/1
```

Core/Aggregation switches (i.e. 8320, 8325) – default port state **disabled**

Access switches (i.e. 6200,6300) – default port state **enabled**

MOD 2- 28

In AOS-CX, all ports are disabled by default on all core and aggregation switches. On access layer switches the port status is enabled by default. You can modify this state using the no shutdown and shutdown commands, as shown in the figure.

Command context: config-if (interface level)

Syntax: shutdown and no shutdown

Here is an example of enabling an interface:

```
Switch(config)# interface 1/1/1
```

```
Switch(config-if)# no shutdown
```

Here is an example of disabling an interface:

```
Switch(config)# interface 1/1/1
```

```
Switch(config-if)# shutdown
```

A description can be configured for each interface, this facilitates management.

Command context: config-if (interface level)

Syntax: description <description>

Here is a configuration example of defining a description for an interface:

```
Switch(config)# interface 1/1/1
```

```
Switch(config-if)# description To-Core-switch
```

You can view the current configuration state for an interface with the command `show run interface 1/1/1`, for example, as shown in the figure.

You can also revert an interface to its default configuration from config-if mode with `default int 1/1/1`, for example, as shown in the figure.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Verifying Interface Status

```
Switch# show interface 1/1/1
```

```
Interface 1/1/1 is down (Administratively down)
```

```
Admin state is down
```

```
State information: Group speed mismatch
```

```
Link transitions: 0
```

```
Description: To-Core-switch
```

```
Hardware: Ethernet, MAC Address: 90:20:c2:bc:3f:00
```

```
MTU 1500
```

```
Type SFP-BT
```

shutdown

description

```
Switch# show interface 1/1/1
```

```
Interface 1/1/1 is up
```

```
Admin state is up
```

```
Link state: up for 15 hours (since Mon Aug 08 20:08:34
```

```
Link transitions: 0
```

```
Description: To-Core-switch
```

```
Hardware: Ethernet, MAC Address: 90:20:c2:bc:3f:00
```

```
MTU 1500
```

```
Type SFP-BT
```

no shutdown



MOD 2- 29

The show interface command shows status and configuration for all interfaces on the switch. In figure's top example, Administratively down means that a shutdown command was configured on that interface. You can also see that an interface description was configured on that interface with the description command.

In the bottom example, notice that the interface is now up. Somebody has entered the no shutdown command. The interface description remains unmodified.

Command context: manager (#)

Syntax: show interface <interface-ID>




Introduction to VLANs

Virtual LANs (VLAN) are a fundamental networking component that help improve security, scalability, and many more features.


CONFIDENTIAL - FOR TRAINING

LANs and VLANs

LAN: Devices in the same broadcast domain



VLAN: Devices in the same broadcast domain



Both scenarios operate in the same way

MOD 2- 31

A local area network (LAN) is a group of devices in the same broadcast domain. The figure shows two LANs, labeled LAN-A and LAN-B. LAN-A switch has two connected devices. Host A connects to Port 1 and Host B connects to port 2. The other switch has similar connections.

Recall that if a switch receives a broadcast on any port, it floods it out all other ports, except the ingress port. Therefore, if Host A sends a broadcast, the LAN-A switch forwards it out all other ports, and so Host B receives the frame. Of course, Hosts D and E do not receive this frame – there is no connection between LAN-A and LAN-B. They are physically separate LANs, connected to physically separate switches.

What if you wanted to connect these three LANs into an internetwork? Simply add a router, which can route unicast and multicast traffic between the LANs.

Now consider a VLAN, which is, just like a physical LAN, a group of devices in the same broadcast domain. Suppose that you have a single, physical Aruba switch named SW1. Hosts A – E connect to ports 1, 2, 11, and 12 on this switch (see the right side of the figure). By default, all these devices are in the same broadcast domain. If Host A sends a broadcast, all other hosts receive the frame. Now suppose that you learned about creating VLANs and created a Virtual LAN named VLAN10. It is as if you have created a small virtual switch inside the physical switch. This virtual switch exists, but it is not connected to any of the physical switch ports.

And therefore, you define or map physical ports 1 and 2 as being members of the red VLAN

10. Similarly, you create VLAN20, and assign ports 11 and 12 as members of the blue VLAN 20. You have now created two separate broadcast domains on a single physical switch. When Host A sends a broadcast, the switch knows to only forward that frame out all ports that are in the same VLAN, with the exception of the ingress port. Only Host B receives the broadcast. If Host D broadcasts, only E receives it.

You have effectively recreated the scenario on the left side of the figure, using a single physical switch. Just like the physical scenario on the left, there is no connectivity between the separate VLANs. No unicast, multicast, nor broadcast traffic shall pass between the VLANs. Of course, you could always connect a router to create an internetwork, just like you did with the physical switches.


But if the scenarios operate in the same way, why bother with VLANs ? Let's talk about the advantages of VLANs and learn how to create and implement them.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

VLANs, Performance, and Security


One large VLAN ☹️

Lots of broadcasts wastes resources
Limited Security options



Multiple small VLANs ☺️

Fewer broadcasts per VLAN
Improved security options



The left-side example in the figure has all hosts on a single VLAN. Having hundreds or thousands of hosts all in the same VLAN certainly simplifies your network design and reduces your configuration burden. You don't even need a router for inter-host communication. You just need the router to get you to the internet. That's nice.

But remember, this is all one broadcast domain. When one host sends a broadcast all other hosts must receive and process it (at least to the point where they realize the broadcast is not for them). Most end systems send frequent broadcasts: to query for services, perform ARP requests, and much more. So now you have all devices receiving frequent broadcasts from all the other devices. This wastes bandwidth on the network, and CPU resources on each device.

Furthermore, simple layer 2 switches have relatively limited security options. By default, everyone on a VLAN can send packets to every other device on that VLAN. Of course, they can all see each other's broadcasts: bad actors might glean useful information from these broadcasts to instigate an attack.

Compare this to the significantly improved design on the right of the figure. Yes, you have a bit more complexity configuring various VLANs, IP addressing, and routing. However, that same large node network is now split into multiple smaller VLANs, say a couple hundred nodes per VLAN. (The actual number of nodes per VLAN depends on network design details and host application usage details that are beyond the scope of this course).

Now each host only sees broadcasts from other members of the VLAN, as opposed to all the hosts in the LAN. Far fewer CPU and bandwidth resources are wasted. You also get better isolation of problems. An issue on one VLAN will be less likely to affect other VLANs, depending on the nature of the issue, of course.

You know that by default folks in VLAN 10 cannot communicate with those in other VLANs. You must configure your routing services to route between those VLANs. And those router functions can include access control lists (ACLs) and other security features to secure that network. As you progress in your training, you will learn how Aruba makes this security much easier, with centralized policies pushed down to appropriate layer 3 devices, and other automated processes that can help maximize your security with minimal effort.

OK, that's why you use VLANs. Let's explore how to configure them.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Creating VLAN

Create VLAN 10

```
SW1(config)# vlan 10
SW1(config-vlan-10)#
```

Create VLAN 2,3,4,5 and 10

```
SW1(config)# vlan 2-5,10
```

Remove VLAN 10

```
SW1(config)# no vlan 10
```


Disable VLAN 10

```
SW1(config)# vlan 10
SW1(config-vlan-10)# shutdown
```

Name VLAN 10

```
SW1(config)# vlan 10
SW1(config-vlan-10)# name Sales
```

Default: all ports map to VLAN 1



VLAN 10 exists, but nothing is connected

MOD 2- 33

AOS-CX refers to VLANs by their VLAN-ID - a number between 1 and 4094. In AOS-CX VLAN 1 is created by default and cannot be removed; and by default, all ports are members of this VLAN. This is a common default for many switches.

The `vlan` command creates a VLAN, which is enabled by default. In the figure's top example, you create VLAN 10. You can also create a series of VLANs with a single command, as shown in the next example.

To remove the VLAN, use the `no vlan` command, like `no vlan 10` to delete VLAN 10. Or, instead of deleting it, you could use the `shutdown` command to disable the VLAN.

It is also a good idea to name the VLAN, as shown in the figure, with the `name` command.

Remember, when you use the command `vlan 10` to create a VLAN, it is as if you have created a virtual switch inside of the actual, physical switch. Although it exists, no ports have yet been defined as members of this VLAN. No devices are yet attached to this virtual switch, as shown in the figure.

The next thing you need to do is associate or map physical interfaces as members of the VLAN.

Access Port

Define VLANs


```
SW1(config)# vlan 10
SW1(config-vlan-10)# name Sales
SW1(config-vlan-10)# exit
SW1(config)# vlan 20
SW1(config-vlan-20)# name Service
SW1(config-vlan-20)# exit
```

Map ports to VLANs

```
SW1(config)# interface 1/1/1-1/1/2
SW1(config-if-<1/1/1-1/1/2># vlan access 10
SW1(config-if-<1/1/1-1/1/2># exit
SW1(config)# interface 1/1/11-1/1/12
SW1(config-if-<1/1/4-1/1/5># vlan access 20
SW1(config-if-<1/1/4-1/1/5># exit
```

```
SW1# show vlan
```

VLAN	Name	Status	Reason	Type	Interfaces
1	DEFAULT_VLAN	up	ok	default	<output omitted>
10	Sales	up	ok	static	1/1/1-1/1/2
20	Service	up	ok	static	1/1/11-1/1/12



The rest of the ports are still mapped to VLAN1

MOD 2- 34

The figure shows SW1, with four connected devices, all in VLAN 1 by default.

- You define VLAN 10, with a name of Sales.
- You define VLAN 20, with a name of Service.

Now you need to assign or map ports to these VLANs.

From the global configuration context, you choose to configure a range of interfaces at the same time – ports 1/1/1 through 1/1/2. Then you assign these ports to VLAN 10 (note that you can perform the assignment with just a single port, a range of ports, or a combination of ranges and individual ports as shown in the figure).

You repeat these steps for ports 1/1/11 – 12, as shown.

Use show vlan to verify VLAN creation and mapping.

Understand that only one VLAN ID can be assigned to an interface (edge port). For example, interface 1/1/1 can't be a member of both VLAN 10 and VLAN 20 at the same time. This would be like you trying to be in the sales room for a meeting and in the engineering room for another meeting at the same time.


ArubaOS-CX switch interface modes can be layer 2 or layer 3. To convert an interface to layer 3 use command routing. If you want to convert an interface to layer 2 mode, use

command no routing.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Extending VLAN Across Multiple Switches

Problem You must extend VLANs over multiple switches using one port



You used TWO ports for inter-switch VLAN connections – not scalable!

MOD 2- 35


The figure shows two switches, each with ports in VLANs 10 and 20. Now you want to connect these two physical switches together (SW1 and SW2), so you connect port 24 of SW1 to port 24 of SW2. What happens? Will all devices be able to communicate?

Recall that by default ports are members of VLAN1, including port 24 on both switches. This means that port 24 is not connected to VLANs 10 and 20, and so cannot carry that traffic. Effectively, you still do not have connectivity between the switches for VLANs 10 and 20. How can you fix this?

One way is to connect two more pairs of port: Map port 22 on both switches to VLAN 10, and port 23 as to VLAN 20. Now all VLANs can communicate. A broadcast from Host A is flooded out all ports in VLAN 10 (except ingress port), so all members of VLAN 10 receive the broadcast. This is good, but there is a problem. You used up 2 physical ports on each switch. What if you had 100 VLANs? We'd need 100 physical switch ports just for interconnects – that is not scalable and typically not feasible.

We need a way to use one single physical port to connect multiple VLANs.

802.1Q Tagging



One trunk link carry traffic of multiple VLANs

MOD 2- 36

Look at the standard Ethernet frame shown in the figure. Suppose that this was a broadcast frame from Host A, flooded out all ports in VLAN 10, including port 24. There is no field in this frame to inform receiving switches about the intended destination VLAN. If SW2 could talk, it might ask, “What do I do with this? Is this for VLAN 10? VLAN 20? VLAN1?”

We need a tagging mechanism, as defined in the IEEE 802.1Q standard to deal with this issue. This standard adds an additional field between the length and payload fields in the Ethernet header: the 802.1Q Tag field. The most important part of this field is the VLAN tag or VLAN-ID (or simply VID). Before SW1 floods this broadcast frame across this specially-defined trunk port, it adds this tag to the frame.

Thus, SW2 receives the frame, looks at the tag and knows, “Ahhh, this frame is for VLAN 10”. It strips off the special tag and forwards a standard Ethernet frame out all ports in VLAN 10 to Hosts E and F. From Hosts A, E, and F’s perspective, the tagging is transparent since they only deal with or see the normal 802.3 Ethernet frame.


Next Host G sends broadcast. SW2 floods a standard Ethernet frame out its local ports. Before flooding the frame out port 24, it’s a tag, “This frame is for VLAN 20”. SW1 receives this frame, sees that it is for VLAN 20, strips off the tag, and forwards a standard Ethernet frame out all ports in VLAN 20 (11 and 12).

A trunk connection is like a virtual cable inside a physical cable, similar to a VLAN being virtual switch inside a physical switch. Thus, you can now extend the VLAN concept across

multiple switches using a single physical port.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

VLAN Trunk Port



In ArubaOS-CX VLAN-1 is the native vlan by default

MOD 2- 37

The 802.1Q standard allows you to select a single VLAN on which no tag or VID is included in the layer 2 header. This means that a standard, untagged Ethernet frame is sent out of the port. This VLAN is known as the native VLAN or untagged VLAN. By default, AOS-CX VLAN 1 is the native VLAN. You can modify this by using `vlan trunk native <vlan-id>` command.

Note: The VLAN-ID field is 12_bits long; this length allows a switch to carry traffic for up to 4,094 VLANs. To avoid inconsistencies Aruba recommends you use the same native VLAN in both sides of a trunk connection.

Configure VLAN Trunks – Allowed VLANs

Configure trunk port

```
SW1(config)# interface 1/1/24  
SW1(config-if)# vlan trunk allow 1,10,20
```

Validate trunk port

```
SW1# show vlan port 1/1/24
```

VLAN	Name	Mode	Mapping
1	DEFAULT_VLAN_1	native-untagged	port
10	Sales	trunk	port
20	Service	trunk	port



MOD 2- 38

In ArubaOS-CX, use the command `vlan trunk allowed` to allow a VLAN-ID(s) to traverse a trunk interface. Multiple VLAN-IDs can be assigned to a trunk interface. In the figure, interface 1/1/25 is configured as a trunk link, and only allows VLANs 1, 10, and 20. Other VLANs will not traverse this link.

To verify the Trunk interface, use the command `show vlan port`.

Note: This example is focused on a trunk port configuration. It is assumed that VLANs are already created, and access ports are mapped to them.

Note: There is a command `vlan allow all` allows all configured VLANs to traverse the trunk port. You may encounter folks who like to use this command, because it simplifies their life – whenever you create a new VLAN, it is automatically allowed to traverse the trunk link without you explicitly allowing it. This is a bad idea, as it can cause security weaknesses and enable bad actors to potentially probe deeper into your network. Unless you are in a non-production lab environment, avoid using this command.

Native VLAN

Modify the native VLAN

```
SW1(config)# interface 1/1/24  
SW1(config-if)# vlan trunk native 10
```

Verifying Trunk port

```
SW1# show vlan port 1/1/24
```

VLAN	Name	Mode	Mapping
1	DEFAULT_VLAN_1	trunk	port
10	Sales	native-untagged	port
20	Service	trunk	port

Peers must have the same Native VLAN

MOD 2- 39

Recall that by default, all VLAN traffic is tagged to traverse a trunk link except the native VLAN, which defaults to VLAN 1. However, you can change the native VLAN as desired. As in the example shown in the figure, use the `vlan trunk native` command to change the untagged VLAN: in this case it was changed to 10. This configuration must match on both switches, so be sure to apply this command to the appropriate port of the attached switch.

Verify your configuration with the command `show vlan port`.

Note: Some customers require all VLANs to be tagged, including the native VLAN, and AOS-CX can support this configuration. As typical, peer switch configurations should match, or connections between devices might fail.

Note: You can use the `show lldp neighbor-info` command to see which switch is connected to port 1/1/24 of the switch you are currently configuring.

MAC Address Table

Used by L2 switches to make forwarding decisions

Forward frames based on Destination MAC


Build MAC table based on Source MAC

Destination MAC	Source MAC	Len	Payload	FCS

```
SW1# show mac-address
MAC age-time           : 300 seconds
Number of MAC addresses : 3
```

MAC Address	VLAN	Type	Port
90:20:c2:bc:ee:00	10	dynamic	1/1/1
00:0b:86:b4:eb:37	10	dynamic	1/1/2
30:8d:99:ef:84:56	20	dynamic	1/1/11

Entry timeout = 300s



Layer 2 switches use the MAC address table to make forwarding decisions. The switch builds this table automatically, based on the source MAC address of the frames that it receives from connected devices. How does the switch do this?

Suppose that you power up switch SW1. Its MAC address table is initially empty. You can verify this with the show mac-address command.

Then Host A transmits a frame to Host B. SW 1 receives this frame, which contains a source MAC address of 90:...:00 on port 1; and port 1 is defined as being a member of VLAN 10. Thus, the switch adds the entry to the MAC address table (the first entry as shown in the figure). The switch basically mapped the MAC address to the switch's VLAN and port as an entry in the MAC address table.

At this point, this is the only entry in the table. The switch does not yet know where to forward frames to other devices, like destination 00:...:37, so it floods this frame out all the remaining ports in VLAN 10. All VLAN 10 hosts receive this frame and see the destination MAC address – most discard the frame since the destination MAC address doesn't match their own MAC address, with the exception of Host B. "This is my MAC address, and I will respond to Host A's message".


And so Host B responds to Host A with source MAC of 00:...:37, and destination MAC = 90:...:00. The switch then adds the second entry you see in the table MAC 00:0b:86:b4:eb:37 is in VLAN 10, connected to port 2. A similar process will happen with Host C on port 11. 127

Therefore, switches automatically build the MAC table based on source MAC addresses that they see entering the switch ports, and forward frames based on destination MAC addresses in these frames.

By default, table entries are maintained for 300 seconds (5 minutes). You can verify the MAC address table in ArubaOS-CX using the `show mac-address` command.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Multi-Layer Interface Types: Layer 3 Switch Ports



MOD 2-41

You just learned about layer 2 switchports and how to assign VLANs to them, often to connect to edge devices (access ports) or to interconnect switches (trunk ports). Because AOS-CX switches are multi-layer switches, you can create logical layer 3 interfaces for those VLAN broadcast domains. These interfaces are called switch virtual interfaces (SVIs), which is basically a virtual layer 3 routed interface that exists only inside the switch and is specific to the respective VLAN.

Suppose you create the logical interface SVI 1 using the `interface vlan 1` command. Once you configure an IP address on this SVI interface, it connects to the internal routing construct for VLAN 10. The SVI interface then services routable traffic from VLAN 1 to other destination networks. You define SVI 2 similarly. With some additional configuration, your switch can route traffic between your VLANs.

Now, suppose you need to connect your multi-layer switch to an external router, perhaps using port 1/1/24, as shown in the figure, and you do not want to use a trunk port between the switch and the router. By executing the routing command on the switch interface, this turns the physical port (1/1/24) into a layer 3 interface. This allows you to assign an IP address to the switch's physical interface; and the AOS-CX switch automatically adds the IP network address as a directly connected route in its routing table.

Note: Interfaces on Aruba 6400 and higher series switches are configured as routed interfaces by default. 6300 and smaller switch interfaces are configured as layer 2 switch ports by default. The Aruba 6200 supports SVIs, but interfaces cannot be natively configured

as a routed port.

Let's look at configuration...

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Two Methods of IP Address Configuration

Assign IP address to a VLAN interface (SVI)

```
Switch(config)# interface vlan 3  
Switch(config-if-vlan)# ip address 10.1.3.1/24
```

Assign IP address to a physical port

```
Switch(config)# interface 1/1/5  
Switch(config-if)# routing  
Switch(config-if)# ip address 10.1.5.1/24
```

Verification

```
Switch(config)# show ip interface brief
```

Interface	IP Address	Interface Status
1/1/5	10.1.5.1/24	link/admin up/up
vlan3	10.1.3.4/24	up/up



MOD 2-42


The figure shows a typical SVI addressing configuration. You create the SVI interface by simply executing the interface vlan x from global config mode. This places you in the context where you can configure the interface properties, like the IP address of 10.1.3.1 with a /24 subnet mask (255.255.255.0). By default, creating an SVI interface automatically enables it.

To configure a physical layer 3 interface, enter config interface mode as shown. First, convert it to a layer 3 interface with the routing command. Then, configure the IP address and subnet mask, just like you did with the SVI 3 interface.

Validate your configuration by executing the show ip interface brief command.

Address Resolution Protocol (ARP)

Maps layer-3 IP addresses to layer-2 MAC addresses



MOD 2-43

Imagine that you are in a conference room and are leading a discussion with people that you have just met. You don't yet know everyone's name. Someone enters the room and hands you a note that says, "Important message for Alvin Rogers". You do not know which person is Alvin, so you say, "Excuse me everyone. Who is Alvin Rogers?". Everyone hears your request, but only Alvin responds, "I'm Alvin". You thus can associate the name Alvin to an actual person seated at a particular chair in the room. This is very much like the Address Resolution Protocol (ARP), which maps layer 3 IP addresses to layer 2 MAC addresses.

In the figure, Host A knows that it must communicate with the host at IP address 10.1.20.200. To do this, it must build an Ethernet frame, which requires its own source MAC address (90::00), and Host B's destination MAC address, which it does not yet know.

To learn the destination's MAC address, Host A broadcasts an ARP request, where the destination MAC address ff:ff:ff:ff:ff:ff. (In the IP header, the destination IP address is 255.255.255.255). You previously learned that switches forward broadcasts out all ports in the same VLAN (except for ingress port). Thus, all hosts in the VLAN receive this frame and ignore it, "I'm not 10.1.20.200", with the exception of Host B.

Host B responds with an ARP reply, which is a unicast message specifically to Host A's MAC address, "I am 10.1.20.100, and my MAC address is 00::37". Host A receives this reply and creates an entry in its ARP table, sometimes called an ARP cache. Host A maps the IP address of 10.1.20.200 to 00:0b:86:b4:eb:37. Thus, the next time Host A must communicate with Host B, it doesn't need to broadcast an ARP message, but instead uses the information it

learned and placed in its ARP cache.

It is important to point out that when Host B received the ARP request from Host A, it learned about Host A's MAC address and IP address. Thus, Host B creates an entry in its own ARP table, mapping 10.1.20.100 to MAC address 90:20:c2:bc:ee:00.

From a Windows PC's command prompt, use the command `arp -a` to view the mappings of IP addresses to MAC addresses that the PC has learned.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY




Discovering the Network

Let's explore a few key tools to make it easy to discover, diagnose, and document nearly any network.

CONFIDENTIAL - FOR TRAINING

LLDP and LLDP-MED



MOD 2- 46

LLDP is an open-standard link layer protocol used by network devices to advertise their identity, capabilities, and neighbors on a local area network. Based on IEEE 802 technology, you use it verify layer-2 connectivity between devices, discover an unknown network topology, and help build network documentation.

The MED is an enhancement of LLDP, known as LLDP-MED. It helps to provision devices: determine correct power settings for connected VoIP phones and APs, auto-assign VLAN and QoS settings, and thus enable plug-and-play networking.

Features include:

You can define LAN policies based on the connected device: VLAN, layer 2 802.1p class of service (CoS) values, layer 3 differential service code points (DSCP) settings, and more.

You can discover the location of the device for VoIP and enhanced 911 services.

Extended and automated power management of power-over-Ethernet (PoE) edge devices.


Inventory management allows you to track your network devices, and determine device-specific information, like the manufacturer, software and hardware versions, serial or asset number.

Zero-touch provisioning (ZTP) provides support for quickly deploying new devices in a

network.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

LLDP Operation



MOD 2-47

The IEEE 802.1AB LLDP provides a standards-based method for network devices to discover each other and exchange information about their capabilities. Devices send LLDP data packets out all interfaces with outbound LLDP enabled. Devices accept LLDP advertisements from L2-connected neighbors on ports with inbound LLDP enabled. Devices accept LLDP advertisements from layer 2 connected neighbors on ports where inbound LLDP is enabled. See the figure for an illustration of this process.

Inbound packets from neighbor devices are stored in a special LLDP management information base (MIB). This information can then be queried by other devices through SNMP. Network management tools can use LLDP information to create accurate physical network topologies. LLDP helps to determine which devices are neighbors and through which interfaces they connect. LLDP operates at layer 2 and requires an LLDP agent to be active on each interface that sends and receives LLDP advertisements.

LLDP advertisements can contain a variable number of TLV (type, length, value) information elements. Each TLV describes a single attribute of a device such as: system capabilities, management IP address, device ID, and the port ID.

Packet boundaries

LLDP packets travel only to directly-connected peers. These packets are never forwarded to any other devices, regardless of whether they are LLDP-enabled. An intervening hub or repeater forwards the LLDP packets it receives in the same manner as any other multicast.

packets it receives. Therefore, two LLDP switches joined by a hub or repeater handle LLDP traffic in the same way that they would if directly connected.

Any intervening IEEE 802.1D device or layer-3 device that is either LLDP-unaware or has disabled LLDP operation will not forward (drops) the LLDP frames.

LLDP-MED

LLDP-MED in the switches uses the standard LLDP commands described earlier in this section, with some additional extensions that were not in the original LLDP standard, and also introduces new features unique to LLDP-MED operation.

Note: Cisco's version of the standard LLDP protocol is the Cisco Discovery Protocol (CDP). CDP is supported in read-only mode, to support connectivity of Cisco Voice over IP (VoIP) phones. Cisco's CDP protocol is supported on the ArubaOS-CX switches by default.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

LLDP Configuration

```
lldp
interface 1/1/1
  lldp receive
  lldp transmit
```

Default on AOS-CX

```
Switch# show lldp configuration 1/1/1
LLDP Global Configuration
=====
LLDP Enabled : Yes
LLDP Transmit Interval : 8
LLDP Hold Time Multiplier : 4
LLDP Transmit Delay Interval : 2
LLDP Reinit Time Interval : 2
LLDP Port Configuration
=====
PORT TX-ENABLED RX-ENABLED
-----
1/1/1 Yes Yes
```

```
switch# show lldp tlv
TLVs Advertised
=====
Management Address
Port Description
Port VLAN-ID
System Capabilities
System Description
System Name
```



MOD 2-48

LLDP is enabled by default, both globally and on the switch interfaces for AOS-CX switches. You can configure LLDP, if needed, by using the syntax shown in the figure, as well as validate what you configured. Aruba gateways support LLDP, but it is disabled by default.

A portion of the output is not shown in the figure. At the bottom of the show LLDP configuration command you would see the LLDP advertised TLVs. But you can also use the show lldp tlv command to view this information. It is good that the switch advertises information about itself; however, with few exceptions, the switch does not advertise TLVs for LLDP-MED. You would have to manually enable the TLV values that are of interest to you.

Note: With LLDP-MED, you can use the port VLAN ID (native VLAN) to detect native VLAN configuration inconsistencies between connected networking devices.

LLDP Discovery

```
Access-1# show lldp neighbor-info 1/1/21
```

```
Port : 1/1/21
Neighbor Entries : 1
Neighbor Entries Deleted : 0
Neighbor Entries Dropped : 0
Neighbor Entries Aged-Out : 0
Neighbor System-Name : Core-1
Neighbor System-Description : Aruba JL635A GL.10.04.0030
Neighbor Chassis-ID : 90:20:c2:c0:ae:00
Neighbor Management-Address : 90:20:c2:c0:ae:00
Chassis Capabilities Available : Bridge, Router
Chassis Capabilities Enabled : Bridge, Router
Neighbor Port-ID : 1/1/20
Neighbor Port-Desc : 1/1/20
Neighbor Port VLAN ID : 1
TTL : 120
```




MOD 2- 49

You can easily use LLDP for neighbor discovery as shown in the figure. You should examine the Port VLAN ID (native VLAN) to detect native VLAN inconsistencies between devices.

ICMP

Used by network devices to send error and operational info related to IP communications



MOD 2- 50

Suppose that a new printer is added into the network and a laptop wants to use this device. The first troubleshooting step is to verify connectivity between these two devices, to ensure that the network is properly configured. Internet Control Message Protocol (ICMP) is the one IP protocol that helps you to address this issue.

ICMP is used by network devices to send error and operational information related to IP communications: it's basically a support protocol for IP. In the figure Device B sends an ICMP message to Device A. The IP header of the IP packet has a "1" in the Protocol field. Each IP protocol has a number assigned to it, just like TCP (6) and UDB (17) are assigned protocol numbers. From this perspective, ICMP seems to sit or operation at layer 4, but its purpose is largely to function at layer 3. Think of ICMP as a kind of an extension to the IP protocol to Layer 3. Obviously, ICMP uses neither TCP nor UDP headers: however, IP and ICMP communicate directly. In other words, if IP has a problem, ICMP informs other devices about this problem.

What type of problem or information might this be? The ICMP Type (with help from the Code field) indicates the type of message, as indicated by the figure.

With most operating systems, including AOS-CX switches, the ping command used to test network connectivity. When executing ping, you specify a destination to test. The resulting ICMP packet consists of an ICMP Type 8 echo request, and the corresponding reply a Type 0 Echo Reply. A type 3 indicates that a router received an IP packet to be routed but lacked a matching route in its route table: a destination unreachable message is displayed. Or perhaps

the router knows that another router has a better path , and so indicates this with a redirect message.

Every IP packet sent includes some time-to-live (TTL) value. Each router along the packet's path decrements the TTL by 1. The router that decrements TTL to 0 will discard the packet and inform the original source that its packet was discarded because its time-to-live had expired. This ensures that if a routing loop exists, the packet doesn't continually rotate in the loop for infinity.

Let's look at a couple more of these ICMP message types in more detail.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Ping



```
Switch(config)# ping 10.254.1.23
PING 10.254.1.23 (10.254.1.23) from 10.12.1.2 : 100(128) bytes of data.
108 bytes from 10.254.1.23: icmp_seq=1 ttl=61 time=4.72 ms
108 bytes from 10.254.1.23: icmp_seq=2 ttl=61 time=3.10 ms
108 bytes from 10.254.1.23: icmp_seq=3 ttl=61 time=2.98 ms
```



MOD 2- 51

The ping command sends ICMP echo request messages to a destination device. The remote device responds to each received echo request with an echo reply message. When the source device receives the messages back, it understands that the network was properly configured since the communication succeeded between the two devices.

In AOS-CX the ping command must specify the IPv4 or hostname.

Syntax: ping <IPv4-address or hostname>

Command context: operator (>) manager (#)

Authority: operators or administrators

The ping command is supported in most of the operating systems in the market today. The example ping session shown in the figure is from the command prompt of an AOS-CX switch.

Traceroute



```
Switch# traceroute 10.11.12.104
traceroute to 10.11.12.104 (10.11.12.104), 1 hops min, 30 hops max, 3
sec. timeout, 3 probes
 1  10.11.11.1      0.202ms  0.099ms  0.504ms
 2  10.11.0.2      0.188ms  0.101ms  0.099ms
 3  10.11.12.104   0.142ms  0.087ms  0.122ms
```



MOD 2- 52

Traceroute is a variant of the ICMP protocol and it is a useful troubleshooting tool when you deploy a layer 3 technology, such as an IP routing protocol. This command tracks the path that a packet takes to reach its destination, listing the layer 3 devices that are in the path (the hops the packet takes to reach the destination).

In AOS-CX the traceroute command must specify an IPv4 address or hostname.

Syntax: `traceroute <IPv4-address or hostname>`

Command context: operator (>) manager (#)

Authority: operators or administrators

In the example shown in the figure, the switch tries to establish connectivity to the destination 10.11.12.104, via two intermediate routers (10.11.11.1 and 10.11.0.2). You will learn more about routing and layer 3 in the next chapters.

Note: Many firewalls block ping and trace commands. If you are attempting to use ping and trace commands to test connectivity through a firewall, the tests may fail. This does not necessarily mean that anything is wrong; it just means that the firewall is dropping your ping and trace packets as a security measure.

Note: Intermediate route control plane (software) in the path must generate an ICMP TTL expired packet when the TTL is decremented to 0. This is because the data plane drops the

packet, while the control plane sends notification.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Link Aggregation

Link aggregation can be used to improve switch-to-switch performance and reliability. You might also leverage this feature for switch-to-server connections.

CONFIDENTIAL - INTERNAL

Link Aggregation: Overview

Bundle multiple physical links into one logical link - a Link Aggregation Group (LAG)

Virtual Interface controls physical ports

Protocols and processes refer now to LAG



Performance Increase bandwidth – traffic is distributed across member ports

Reliability One member port fails, traffic flows over other member ports

MOD 2- 54

Ethernet link aggregation bundles multiple physical Ethernet links into one logical link, called a link aggregation group (LAG). The figure shows two devices connected by three physical Ethernet links, bundles into one virtual link called LAG 1. Link aggregation is supported on many products like servers, storage, routers, and firewalls, but is also commonly used to interconnect switches.


The benefits of link aggregation include:

- **Performance:** You get increased bandwidth because all traffic is distributed across the member ports.
- **Reliability:** Member ports dynamically back each other up. When a member port fails, its traffic automatically flows over other member ports.

Enabling Link aggregation on a switch creates a virtual interface. You then assign physical ports to be members of that virtual interface. The switch's various protocols and processes will only perceive and refer to the virtual interface and not the actual physical interfaces of the LAG: they no longer perceive the individual physical interface members.

Because it is seen as a logical interface, understand that broadcast and multicast traffic is sent across a physical link in the bundle, but the remote peer sees this as coming from a logical interface, and thus ensures that the broadcast or multicast is not sent back to the source on one of the other physical links in the LAG in order to prevent the creation of layer 2 loops.

Link Aggregation: Requirements



MOD 2- 55

Configure LAG Interfaces with matching parameters is necessary to avoid errors and misconfiguration. Things that need to match between the two devices comprising a LAG include duplexing (full or half), the link speed, and the media type (the type of cable connector).

AOS-CX displays a warning if you attempt to map mismatched interfaces to a LAG: for example, you should not add a 10 Gb/s link to a LAG group of 1 Gb/s links.

AOS-CX allows each LAG to have up to 8 individual member ports. Use the command `show capacities` to verify your switch capacity.


Now that you understand what a LAG is, you will about two operating modes used for LAGs:

Static LAG, which uses no protocol

Dynamic LAG, which uses the Link Aggregation Control Protocol (LACP)

Static LAG

- Devices do not exchange control messages
- LAG configured independently on each peer



Aruba does **NOT** recommend this implementation

MOD 2- 56

When using the static LAG mode devices do not exchange any control information (left-side of figure). In other words, there is no signaling between switch peers about LAG. You simply configure LAG independently on each peer. If your configuration is good, then the LAG works. However, the switch peers have no knowledge of who they are connected to, or whether they are connected to the same peer, or connected correctly.

This mode is not recommended because a misconfiguration on one side is not detected by the peer. This can lead to unexpected behaviors, which can be challenging to detect and troubleshoot. Another scenario that could cause problems is shown in the figure's right-hand example. Switch Access-1 is not aware that you have erroneously connected LAG member ports to two separate physical switches.


Configuring a Static Link Aggregation Group

```
Switch(config)# interface lag <id>
Switch(config-lag-if)# no routing
Switch(config)# interface {member/slot/port}
Switch(config-if)# lag <id>
```

Create LAG interface

Disable routing on LAG

Map the interface to LAG



```
Core-1(config)# interface lag 20
Core-1(config-lag-if)# no routing
Core-1(config)# interface 1/1/43
Core-1(config-if)# lag 20
Core-1(config)# interface 1/1/44
Core-1(config-if)# lag 20
```

```
Core-2(config)# interface lag 20
Core-2(config-lag-if)# no routing
Core-2(config)# interface 1/1/43
Core-2(config-if)# lag 20
Core-2(config)# interface 1/1/44
Core-2(config-if)# lag 20
```



MOD 2- 57

The following steps create a layer 2 Dynamic Aggregation Group:

1. Create a LAG interface, represented by an identifier – any unused value between 1 and 256:

```
Switch(config)# interface lag <id>
```

2. Disable routing (Layer 3 capabilities) in the new LAG interface. This step limits the LAG to only process layer 2 frame headers:

```
Switch(config-lag-if)# no routing
```

3. Map a port member to the LAG.

```
Switch(config)# interface 1/1/1
```


```
Switch(config-if)# lag <id>
```

Note: Member ports lose their port VLAN configuration when added to a LAG. A new LAG interface is an access port in VLAN 1 by default. Any protocol information from here onward would be done on the virtual LAG interface and not the physical interfaces that comprise the LAG.

802.3ad: Dynamic LAG

LACP Standard protocol – Recommended

- Peers exchange periodical LACP Data Unit messages
- Switches know about LACP peers



- System ID: Identifies the LACP peer
- Operational key: Identifies LAG instance
- Include **Flags**

• LACP_Activity	• Synchronization
• LACP_Timeout	• Collecting
• Aggregate	• Distributing



MOD 2- 58

With IEEE 802.3ad peer devices that implement dynamic LAG exchange control messages to establish and maintain the LAG. This mechanism will also detect link failures and ensure that LAG port members terminate on the same device and meet the LAG requirements.

The 802.3ad implementation is commonly known as the Link Aggregation Control Protocol (LACP). The two peers periodically exchange LACP messages called LACP data units (LACP-DU). These messages include:

System ID, which uniquely identifies the switch.

Operational key, which uniquely identifies the Link aggregation group.

Note: Dynamic LAG or LACP is the recommended method to implement link aggregation, to avoid unexpected network problems.

There are some flags that are included in the LACP-DUs. The bottom content section in the figure shows the meaning of these flags.


With 802.3ad standard you can bundle up to 16 ports; however, only 8 active ports can be used at any given time.

802.3ad Operational Modes

Passive Device waits to receive a LACP Data Unit

Active Device transmits LACP Data Unit

Device A port mode	Device B port mode	LAG formed?
Passive	Passive	No
Passive	Active	Yes
Active	Passive	
Active	Active	



MOD 2- 59

LACP can be configured in one of two modes, which controls how peer negotiation proceeds.

Passive mode, the device passively waits to receive LACP-DU messages from the peer to dynamically create the LAG. This mode places the LAG in a listening state. It is as if the switch is thinking, "I'll just sit here and wait until I hear from my peer." If you configure the peer switch to also be in passive mode, then it is thinking the same thing. Both switches wait to hear from the other, and so a functional LAG connection is never formed. At least one peer must be in active mode.

Active mode, the device actively transmits LACP-DU messages over its member ports: "Hey! Let's form a LAG". Whether the peer is passive or active, it responds, negotiation continues, and the LAG successfully forms if the requirements are met. This is shown in the figure.

The rate LACP-DU messages are sent is configurable, there are 2 settings: fast and slow. Slow which is the default is a recommended setting.


Let's talk about configuration.

Configuring a Dynamic Aggregation Group

```
Switch(config)# interface lag <id>
Switch(config-lag-if)# no routing
Switch(config-lag-if)# lacp mode {active | passive}
Switch(config)# interface {member/slot/port}
Switch(config-if)# lag <id>
```

Annotations:

- Create LAG interface
- Disable routing on LAG
- LACP mode
- Map the interface to LAG



```
Core-1(config)# interface lag 20
Core-1(config-lag-if)# no routing
Core-1(config-lag-if)# lacp mode active
Core-1(config)# interface 1/1/43
Core-1(config-if)# lag 20
Core-1(config)# interface 1/1/44
Core-1(config-if)# lag 20
```

```
Core-2(config)# interface lag 20
Core-2(config-lag-if)# no routing
Core-2(config-lag-if)# lacp mode active
Core-2(config)# interface 1/1/43
Core-2(config-if)# lag 20
Core-2(config)# interface 1/1/44
Core-2(config-if)# lag 20
```



MOD 2-60


The configuration for dynamic LAG is the same as the one used for a static LAG, except for an extra line where you need to specify the LACP mode. Fairly simple, yes?

Important: Dynamic link aggregation using LACP is the recommended configuration; as opposed to static link aggregation, which is not recommended.

Verifying LAG Interface

Switch# **show interface lag**

```
Aggregate lag1 is up
Admin state is up
Description :
MAC Address           : 88:3a:30:97:59:c0
Aggregated-interfaces : 1/1/25 1/1/26 2/1/25 2/1/26
Aggregation-key       : 1
Aggregate mode        : active
Speed                 : 40000 Mb/s
L3 Counters: Rx Disabled, Tx Disabled
<<<<Omitted output>>>>
```



Switch# **show interface lag brief**

Port	Native VLAN	Mode	Type	Enabled	Status	Reason	Speed (Mb/s)
lag1	1	trunk	--	yes	up	--	40000



MOD 2- 61

Use the show interface lag to verify your LAG configuration and status (see figure). You can also use show interface lag brief to see the LAG's total available bandwidth. The speed reflects the actual active member ports. This could be 20,000, 30,000, or 40,000 depending on active 10GbE ports in the LAG.

Verifying LAG Interface

```
Switch(config)# show lacp interfaces
```

Actor details of all interfaces:

Intf	Aggr Name	Port Id	Port Pri	State	System-ID	System Pri	Aggr Key	Forwarding State
1/1/27	lag1	28	1	ASFNCD	88:3a:30:a5:d5:80	65534	1	up
1/1/28	lag1	29	1	ASFNCD	88:3a:30:a5:d5:80	65534	1	up

Partner details of all interfaces:

Intf	Aggr Name	Port Id	Port Pri	State	System-ID	System Pri	Aggr Key
1/1/27	lag1	28	1	ASFNCD	88:3a:30:a5:26:80	65534	1
1/1/28	lag1	29	1	ASFNCD	88:3a:30:a5:26:80	65534	1



MOD 2- 62

You can obtain detailed information on the LAG by using the `show lacp interfaces` command. The output will provide valuable troubleshooting information such as port forwarding state, the partner system mac-address, and the LACP flags. Flags indicate the state of link aggregation on each physical interfaces bundled into LAG.

Note: The Partner output in the `show lacp interfaces` command is based on the received LACP information from the partner.

Port 2 Connects to Invalid LAG ID on Partner Side

```
Switch(config)# show lacp interfaces
```

Actor details of all interfaces:

Intf	Aggr Name	Port Id	Port Pri	State	System-ID	System Pri	Aggr Key	Forwarding State
1/1/27	lag256	28	1	ALFNCD	8c:85:c1:65:0e:c0	65534	256	up
1/1/28	lag256	29	1	ALFO	8c:85:c1:65:0e:c0	65534	256	lacp-block

Partner details of all interfaces:

Intf	Aggr Name	Port Id	Port Pri	State	System-ID	System Pri	Aggr Key
1/1/27	lag256	1	1	ALFNCD	02:01:00:00:01:00	65534	1
1/1/28	lag256	1001	1	ALFN	02:01:00:00:01:00	65534	2



MOD 2-63

Troubleshooting LAB problems

Let's look at some sample output of the show lacp interfaces command to troubleshoot problems with a LAG.

The figure shows the first example. In this example, the second local port in LAB256 (port 1/1/28) is not connected to the same logical LAG on the partner side as the local port 1/1/27. In other words, the Partner has two LAG interfaces, and this switch's physical ports in LAG256 are connected to two different LAGs on the partner, which is a misconfiguration.

Port 1/1/25 Connects to Different LACP System ID

```
Switch(config)# show lacp interfaces
```

Actor details of all interfaces:

Intf	Aggr Name	Port Id	Port Pri	State	System-ID	System Pri	Aggr Key	Forwarding State
1/1/25	lag256	26	1	ALFO	8c:85:c1:65:0e:c0	65534	256	lACP-block
1/1/27	lag256	28	1	ALFNCD	8c:85:c1:65:0e:c0	65534	256	up
1/1/28	lag256	29	1	ALFO	8c:85:c1:65:0e:c0	65534	256	lACP-block

Partner details of all interfaces:

Intf	Aggr Name	Port Id	Port Pri	State	System-ID	System Pri	Aggr Key
1/1/25	lag256	27	1	ALFN	8c:85:c1:66:f0:40	65534	256
1/1/27	lag256	1	1	ALFNCD	02:01:00:00:01:00	65534	1
1/1/28	lag256	1001	1	ALFN	02:01:00:00:01:00	65534	2



MOD 2- 64

In this example, local port 1/1/25 is connected to a different LACP system ID; in other words, the local LAG has one port connected to one partner and another port, in the same LAG, connected to a completely different partner. This is obviously a misconfiguration.

Note: When the peer does not send LACP, AOS-CX blocks the port. This behavior may be different on other platforms.

Port Connects to Different LACP System ID w/Same Key ID

```
Switch(config)# show lacp interfaces
```

Actor details of all interfaces:

Intf	Aggr Name	Port Id	Port Pri	State	System-ID	System Pri	Aggr Key	Forwarding State
1/1/25	lag256	26	1	ALFO	8c:85:c1:65:0e:c0	65534	256	lacp-block
1/1/27	lag256	28	1	ALFNCD	8c:85:c1:65:0e:c0	65534	256	up
1/1/28	lag256	29	1	ALFO	8c:85:c1:65:0e:c0	65534	256	lacp-block

Partner details of all interfaces:

Intf	Aggr Name	Port Id	Port Pri	State	System-ID	System Pri	Aggr Key
1/1/25	lag256	27	1	ALFN	8c:85:c1:66:f0:40	65534	1
1/1/27	lag256	1	1	ALFNCD	02:01:00:00:01:00	65534	1
1/1/28	lag256	1001	1	ALFN	02:01:00:00:01:00	65534	2

MOD 2-65

This example in the figure is similar. In this example, ports 1/1/27 and 1/1/28 are connected to one switch, where the ports are connected to two different LAGs on the same partner. On top of this, port 1/1/25 is connected to a completely different partner, but the key (LAG ID) matches that of what is configured on the local switch. Obviously, someone really messed up the wiring when trying to cable the switches together.

Port Does Not Have LACP Enabled on Peer Side

```
Switch(config)# show lacp interfaces
```

Actor details of all interfaces:

Intf	Aggr Name	Port Id	Port Pri	State	System-ID	System Pri	Aggr Key	Forwarding State
1/1/25	lag256	26	1	ALFOE	8c:85:c1:65:0e:c0	65534	256	lacp-block
1/1/27	lag256	28	1	ALFO	8c:85:c1:65:0e:c0	65534	256	lacp-block
1/1/28	lag256	29	1	ALFNCD	8c:85:c1:65:0e:c0	65534	256	up

Partner details of all interfaces:

Intf	Aggr Name	Port Id	Port Pri	State	System-ID	System Pri	Aggr Key
1/1/25	lag256	0	0	PLFOEX	00:00:00:00:00:00	0	0
1/1/27	lag256	1	1	ALFN	02:01:00:00:01:00	65534	1
1/1/28	lag256	1001	1	ALFNCD	02:01:00:00:01:00	65534	2

MOD 2-66


In this example the Partner port does not have LACP enabled.

Note: When the peer does not send LACP messages, AOS-CX switches will block the port.

Load Sharing Algorithm

Link decision: Hash algorithm uses either:

- Source and destination IP
- Source and destination MAC



MOD 2-67

Peer LAG devices must balance the traffic load across multiple physical interfaces. How does the switch decide which member port should be used for a given packet? It uses a Hash algorithm.

Load sharing algorithm

A hash algorithm is a mathematical function that is applied to an input (x). Interestingly, if you see the output (y) of this function, you cannot derive or infer input (x). Therefore, it is referred to as a one-way function. Another characteristic of a hash is that if the same input (x) has been entered, the result will be always the same. So, what does the switch use as inputs to this function to determine which link to use in the LAG?

The switch uses frame and/or packet header information as hash function inputs, and the result is the port member to be used for that particular frame. Depending on switch algorithm used, inputs for the hash could be:

- Layer 4 TCP/UDP ports
- Layer 3 source and destination IP addresses
- Layer 2 source and destination MAC addresses

Note: Layer 3 is the default on AOS-CX switches and is typically recommended to use in most situations.

The figure shows how the layer 2 source and destination MAC addresses are taken as inputs

for the algorithm. The output is to use physical port 1 for all frames with this source and destination MAC address combination: always. Only a link failure will cause another port in the LAG to be used for these frames.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Configure Link Aggregation Load Sharing

Verifying Hash algorithm

```
Switch# show lacp aggregates

Aggregate name      : lag1
Interfaces          : 2/1/26 1/1/25 2/1/25 1/1/26
Heartbeat rate      : Fast
Hash                : l3-src-dst
Aggregate mode      : Active
```

Changing algorithm

```
Switch(config)# interface lag <ID>
Switch(config-lag-if)# hash l4-src-dst
```

Default setting is l3-src-dst

MOD 2-68

With AOS-CX switches, the default input information for the hash algorithm are the Source and Destination IP addresses. You can verify with the `show lacp aggregates` command. Notice that the highlighted hash information is `l3-src-dst`. You know the `l3` means layer 3, which uses the source and destination addresses.

In some cases, the hash algorithm's use of source and destination IP addresses may not properly distribute the load equally among port members. This could lead to traffic congestion and poor performance.

In that case, you can modify the hash input data. As shown in the figure, you have decided to use layer 4 source and destination port as inputs to the hash function. This might help to avoid port member oversubscription when there are multiple connections between the same two devices, for example.



Knowledge Check

Self-check on key learning points



How about a knowledge check?

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Question #1

Which statements below accurately describe initial switch access and setup? (Select two.)

- a. You can make an out-of-band Management (OOBM) connection to a switch using a serial connection.
- b. One advantage of making an out-of-band Management (OOBM) connection to a switch using the Management Port is that it enhances redundancy for your data traffic.
- c. One advantage of making an OOBM connection to a switch using the Management Port is that it can mitigate bad actors trying to gain access to your switches.
- d. You connect to a switch via the console port, you must configure that port with an IP address.

a

Knowledge Check


MOD 2-70

A,C

Question #2

How would you properly identify the port indicated in the figure?

- a. Interface 1/48
- b. Interface 1/1/48
- c. Interface 48/1/1
- d. Interface 1/48/1



Knowledge Check

MOD 2-71

B

Question #3

When entering Command Line Interface (CLI) commands on a switch, why might you use the <Tab> key?

- a. To move up one level in the CLI prompt mode hierarchy
- b. To access the context-sensitive help feature
- c. To exit out of global configuration mode
- d. To complete a CLI command that is partially entered

a

Knowledge Check

MOD 2-72

D

Question #4

Your colleague asks you about managing configuration files on an Aruba switch. Which statements below would improve their understanding? (Select three.)

- a. The running-config is stored in RAM.
- b. The startup-config flash memory.
- c. When the switch reboots unsaved changes to the startup-config are lost.
- d. ArubaOS-CX considers both the startup-config and the running-config to be “checkpoints”.

a

Knowledge Check

MOD 2-73

A,B,D

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Question #5

Which of the options below are advantages of Virtual LANs (VLAN) as compared to physical LANs? (Select two.)

- a. VLANs segment the broadcast domain to improve performance.
- b. VLANs are more scalable.
- c. Routing between VLANs does not require a router.
- d. Multiple VLANs can be defined on a single switch.

a

Knowledge Check

MOD 2-74

B,D

Question #6

You use trunk ports to extend the range of VLANs supported on a switch.

- a. True
- b. False

a

Knowledge Check

MOD 2-75

F

Question #7

Which of the following are true about the Link Layer Discover Protocol (LLDP)? (Select two.)

- a. It helps you discover, document, and troubleshoot networks.
- b. LLDP has a Media Endpoint Discovery (MED) enhancement to help with “plug-and-play” networking.
- c. LLDP operates at Layers 2 and 3 of the OSI model.
- d. LLDP is enabled by default on both Aruba AOS-CX switches and Aruba Gateways.

a

Knowledge Check

MOD 2- 76

A,B

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Question #8

Choose the options below that accurately describe Link Aggregation.

- a. It enhances both performance and reliability.
- b. Member ports in a Link Aggregation Group (LAG) should have matching duplex, speed, and media.
- c. Static LAG configuration is preferred over other methods, due to its simplicity.
- d. Peers that use the Link Aggregation Control Protocol (LACP) can automatically form a LAG between them.

a

Knowledge Check

MOD 2- 77

A,B,D



Lab Activity

Consider the lab activity.

CONFIDENTIAL - FOR TRAINING

Lab Tasks

Lab 02.01: Switching Fundamentals

Task 1 – 2

- Initial Edge Switch Access
- Configure Edge switch uplink to Aggregation switch

Task 3 – 4

- Configure PC port as VLAN Access port
- Configure Edge switch uplink to Aggregation switch

Task 5

- Practice VLAN port configuration

Lab 02.02: Link Aggregation

Task 1 – 2

- Configure Edge1 with Link-Aggregation to Edge2
- Practice and verify the LAG configuration on Edge2

Task 3 – 4

- Configure LAG as a VLAN trunk
- Link-Aggregation Troubleshooting

Task 2.5

- Verify Link-Agg Failover (Optional)
- Static LAG Configuration (Optional)

**Module 3 –
Basic IP
Network
Setup**

MOD 2- 79

The figure provides a brief review of lab tasks. Please see your lab guide for details. When complete, you will continue with Module 3— Basic IP Network Setup.

Rev. 22.41




Basic IP Setup

Aruba Campus Access
Fundamentals

Welcome back to Module 3: Basic IP Setup.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY



Objectives

- Discuss Inter-VLAN routing
- Explain DHCP relay
- Discuss static IP routing
- Configure single-area OSPF

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

- Define networking and fundamental concepts including Local Area and Wide Area Networks, along with typical use cases and components, and explain how devices communicate over these networks using protocols, and how this communication is organized into a seven-layer Open Systems Interconnection (OSI) model.
- Explain various transmission media types: copper or fiber optic cables, or radio waves. You learn how each media type represents digital 1's and 0's, leading to an exploration of binary numbers.
- Compare unicast, multicast, and broadcast traffic types.
- Understand TCP/IP protocol stack.

IP Addressing


Proper IP addressing is vital for hosts, and routers...and you.

CONFIDENTIAL - FOR TRAINING

IP Addressing

Layer 3 identifier for each network device

Feature	IPv4	IPv6
Size of address	32 bits – 4 octets	128 bits – 16 octets
Example address	192.168.10.250	fe80:0000:0000:0000:844d:a9ff:febb:62c3
Possible addresses	4.29 Billion	340 Undecillion (3.4×10^{38})



MOD 3- 4

The Internet Protocol (IP) provides an identification or address for each device in a network. Basic routing decisions are based on analyzing these layer 3 addresses. There are two versions of IP: IP version 4 (IPv4) and version 6 (IPv6). The main difference between them is the format of the IP address and the addressing space. Version 4 can allocate approximately 4.29 billion addresses, while version 6 can allocate 340 undecillion addresses. This module is focused on IPv4 addressing and routing.

An IPv4 address has 32-bits and is expressed in a dotted decimal notation. The dots divide the address in four 8-bit (1-byte) octets. This makes it easy for humans to work with IPv4 addresses. Eight bits enables values from 0 to 255.

The figure shows three hosts in the same LAN, each with a unique IP address. For example, when Host C sends something to Host A, the IP header contains source address 10.0.1.3 and destination address 10.0.1.1.


Each IP address has four octets, organized into two main or logical parts:

Network portion

Host portion

Let's explore IPv4 addressing in more depth.

Home Address vs. IP Address



MOD 3-5

Your home address has a street name and a house number to help the mailperson find it. Likewise, IP addresses have a network portion and a host portion, to help routers deliver data to a specific device.

Everyone on the same network has the same network number (like the street where you live). Perhaps everyone on VLAN 1 has network 10.0.1, and one specific host on that network has a fourth octet of 100. But how do we know which portion of an IP address is for the network, and which portion is for the host?

Subnet Mask

Defines the line between network and host portion of IP address →

Network Portion

Host Portion

Subnet Mask (Binary)

11111111 . 11111111 . 11111111

00000000

Subnet Mask (Decimal)

255 . 255 . 255

0

IPv4 Address

10 . 0 . 1

100




MOD 3-6

Every IP address has an associated subnet mask. This mask defines the line between network portion and host portion of the IP address. The subnet mask is a consecutive sequence of ones followed by a consecutive sequence of zeros. Where the ones end and the zeros begin is the line between network and host portions of the address. Looking at the number in binary makes it very clear where this line is, but it is far more convenient for us humans to write this in decimal. Regardless of how you write it, the mask is merely an instruction, telling devices how to perceive an associated IP address.

Suppose an endpoint, switch, or router receives a packet with 10.0.1.100 as the destination IP address. It compares the address to the mask. Those portions of the IP address that line up with a binary 1 are part of the network number. Those portions of the IP address that line up with the mask's binary 0's are part of the host portion. The subnet mask also defines the size of a network: the number of hosts it can contain. Let's look a couple of examples.


IPv4 Network Example 1



How many hosts on the subnet?

Number of Binary 0s	Total Addresses
8	$2^8 = 256$

10.0.1.0 to 10.0.1.255



MOD 3-7

The example in the figure has IP address 10.0.1.100 with mask 255.255.255.0. Thus, the first three octets represent the network, and the last octet represents the host portion. So, how many hosts can you have on this subnet?

Well, 24 bits of this 32-bit IP address are used to specify the network, so only 8 bits remain to specify hosts on that network. Therefore, 2 to the power of 8 equals number of available hosts. That's 256 unique numbers: from 0 to 255. Just to understand this a bit more, remember your binary to decimal conversion. If you have 11111111 in binary, then $128+64+32+16+8+4+2+1 = 255$.

IPv4 Network Example 2

Network Portion	Host Portion
172 . 16 .	2 . 40
255 . 255 .	0 . 0

How many hosts on the subnet?


Number of Binary 0s	Total Addresses
16	$2^{16} = 65536$
172.16.0.0 to 172.16.255.255	



MOD 3-8

The example in the figure is 172.16.2.40, with a subnet mask of 255.255.0.0. In this example, only 16 bits are used to specify the network, so you have 16 bits remaining to specify hosts. And 2 to the 16th power results in 65, 536 unique hosts: addresses from 172.16.0.0 through 172.16.255.255, with 172.16.2. 40 somewhere in the middle.

IPv4 Address Types



MOD 3-9

There are three types of IP addresses, each with a specific function:


- Network identifier (ID), which uniquely identifies the IP network
- Broadcast address
- Host addresses for the network

The network ID identifies the network itself, and all endpoints within it. This is always the first IP address in the host range, where all host bits are zeros.

The broadcast address is the last IP address in an IP address range, where all host bits are set to binary ones. This allows broadcast communication within the network, like saying, "Attention everyone on this subnet".

Therefore, since one number (first) is reserved as a kind of "name" for the network itself, and another (last) is reserved for the broadcast of that network, all numbers in between those two can actually be used to assign the host unique addresses. Remember, every endpoint in an IP network must have a unique host address.

IPv4 Address Example



MOD 3-10

Let's look at an example in the figure. Network 10.0.1.0 has the subnet mask 255.255.255.0. The first three octets represent the network identifier and the last one represents the broadcast address of the network.

The host octet's first value is 0. Therefore, 10.0.1.0 represents the network ID, which serves as a kind of "network name". It is like saying, "You all live on Main Street".

The host octet's last value is 255, so 10.0.1.255 represents the broadcast address of the network. This is the address that says, "Attention, everyone on Main Street."

The addresses between the network ID and the broadcast address are assigned to devices on the network: 10.0.1.1 to 10.0.1.254. These would correlate to the actual street addresses for people that live on Main Street.

Prefix Length

Subnet Mask (Decimal)
255.0.0.0
255.255.0.0
255.255.255.0

10.0.1.100 255.255.255.0 = 10.0.1.100/24




MOD 3- 11

Look at these typical subnet masks shown in the figure. Of the 32 bits in the top example, how many bits are binary 1s? How about the other two examples? That's right! 8, 16, and 24 bits respectively.

This gives us another, simpler way to specify the subnet mask. Instead of saying "The subnet mask is 255.255.255.0", you can simply say, "The mask is a /24". This is referred to as prefix notation.

The prefix length is written next to the IP address, as in the example shown at the bottom of the figure: IP address / prefix-length, or 10.0.1.100/24, for example.

Using IPv4 Addressing



1. Network/mask is shown one time
2. Unique 4th octet shown next to host
3. Router is assigned .254
4. L2 switches often not shown in L3 diagrams



MOD 3-12


You previously learned about single, large VLAN challenges and how to split them up into multiple subnetworks (subnets). The figure shows VLAN 1, subnet 10.0.1.0/24, with over 250 hosts. Notice four things about this simple figure:

1. The network ID and mask is shown one time: 10.0.1.0/24.
2. The unique fourth octet for each host is shown next to the host. You know that every host on the same network has the same network number (10.0.1.0) and therefore you know that the entire address of the first host is 10.0.1.1, and the next host might be 10.0.1.2 and so on, perhaps up to 10.0.1.253.
3. A router is connected to this VLAN and has been assigned the address of 10.0.1.254. Many seasoned network engineers assign either the first available address (.1) or the last available address (.254) to a connected router in the VLAN.
4. A layer 2 switch is not shown in this diagram. In discussions focused on layer 3 routing, layer 2 switches are depicted as “a wire”. The switch does exist, and each host connects to a unique port on the switch. But since layer 2 switches are transparent to layer 3 communications, it makes sense to omit it from the diagram. It simplifies the figure and focuses your mind on layer 3 concepts.

Similarly, you might show VLAN 2, subnet 10.0.2.0/24, and so on.

IP Routing

Objective Device communications between different networks



MOD 3-13

You have learned how devices communicate within the same network (i.e. broadcast domain), using layer 2 switching. Layer 2 switches forward frames to devices in the same LAN by processing layer 2 frame headers. Recall that switches build a MAC address table based on source MAC addresses and VLANs, and forward frames based on destination MAC addresses.


Now you will learn how layer 3 routers transmit packets between different VLANs or broadcast domains, using the layer 3 IP address destination addresses to make their forwarding decisions. An example is shown in the figure, where Host A in VLAN10 wants to communicate with Host B in VLAN 20.

Layer 3 Analysis

Routing: Layer 3 lookup information to take decisions

- 6 Presentation Layer
- 5 Session Layer
- 4 Transport Layer
- 3 Network Layer
- 2 Data Link Layer
- 1 Physical Layer

Layer 3 Header	
Source IP	Destination IP
10.1.10.20	192.168.1.100




MOD 3- 14

The ability to analyze the layer 3 header and use the destination address for packet forwarding is known as routing. Layer 3 devices, like Aruba gateways, routers, firewalls, and multilayer switches, perform routing: they analyze layer 3 destination IP addresses, select the best path to get from original source to ultimate destination, and then forward packets along that path.

IP Route

Objective Device communications between different networks



Host can add routes into its table, but this solution does not scale

MOD 3-15

When a device communicates with others in a different network, it must know which local network device on its broadcast domain can route the traffic toward the destination network. This information is provided to computers by building an IP routing table.

For endpoints, this information must be manually added, in the form of so-called static or connected routes. Routers and multilayer switches can use manually-added static and/or connected routes or they can dynamically learn the best routes to each destination through the use of a dynamic routing. You will soon learn about dynamic routing protocols in a later module.

A static route must specify the following information, as shown in the figure for Core-1 and Core-2:

Destination IP address: where you want to go

Subnet mask: How that destination IP address is split into a network portion and a host portion

Next hop IP address: The IP address of the router on your network that can route packets to the next hop in order to reach the destination.

Look at the scenario shown in the figure, which shows three networks connected by two routers. Core-1 connects Networks A and B, while Core-2 can route between networks B and

C.


If Host A (in network B) must communicate with Server-1 in Network A, it must use the first route in its routing table: 10.0.0.0/8: "To get to destination 10.0.0.1, which has a mask of 255.0.0.0 (/8), you must use next-hop router at IP address 172.16.0.1.

What if Host A wants to send packets to Server-2? The second route in its routing table, for 192.168.0.0/24, says to sent packets to next-hop router at IP address 172.16.0.2, which is Core-2.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

IP Route

Objective Device communications between different networks



Host can add routes into its table, but this solution does not scale

MOD 3-16

When a device must communication with others in a different network, it must know which local network device on its broadcast domain can route the traffic toward the destination network. This information is provided to computers using IP Routes.

For endpoints, this information must be manually added, in the form of a so-called “static route”. However, routers and multilayer switches can use manually-added static routes or they can dynamically and automatically determine the best routes to each destination, using a routing protocol. You will soon learn about dynamic routing protocols in a later module.

A static route must specify the following information, as shown in the figure for Router 1 and Router 2:

Destination IP address: where you want to go

Subnet mask: How that destination IP address is split into a network portion and a host portion

Next hop IP address: The IP address of the router on your network that can route packets to the destination.

Look at the scenario shown in the figure, which shows three networks connected by two routers. Core-1 connects Networks A and B, while Core-2 can route between networks B and C.

If Host A (in network B) must communicate with Server-1 in Network A, it must use Route 1. Route 1 says, "To get to destination 10.0.0.1, which has a mask of 255.0.0.0 (/8), you must use next-hop router at IP address 172.16.0.1.

What if Host A must use Server-2? Route 2 says, to arrive at 192.168.0.1, which uses a 255.255.255.0 (/24) mask, you must send the packets to next-hop router at IP address 172.16.0.2.


How did these routes get into Host A's routing table? For a connected route, the device examines the IP address and subnet mask configured for its Ethernet adapter and thus knows the local network to which it is connected. For remote networks, the destination network must be added manually. This is referred to as a static route.

Note: The problem with static routes is that they don't scale well: imagine a campus network that might have hundreds or thousands of destination networks. Manually adding these to every endpoint would be problematic and difficult to manage.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Default Gateway

Objective Special route that indicates all possible destinations
Optimizes host routing information



Endpoints typically have only one DG

MOD 3-17

A default gateway (DG) is the device that routes all network destinations for the endpoint devices in a broadcast domain or network. It is like telling a host, “To get to everywhere in the known universe, go to this next-hop address”. The default gateway optimizes and simplifies endpoint routing decisions, since only a single, static route is required on the endpoint devices. This is commonly referred to as a default route and the routing device as the default gateway.

For example, assume that all the people in a neighborhood have fences around their homes. Inside the fence is your network. You can communicate with people inside your fence. If you want to move beyond your fence, you need a gate. A router performs the function of a gate: the default gateway is the gate that you use to leave (exit) your network device.

The figure shows how router Core-1 acts as the DG for all devices in Network B. These devices typically only need to define the default gateway IP address.

Note: Different endpoints in the same subnet could have different default gateways, if there is more than one router on the network. However, this is very uncommon.

On a Microsoft Windows device, a static IP address and default gateway IP address can be defined at the same place. For Windows 10 devices you can follow these instructions:

Click the Windows start button and type Control Panel.

Navigate to Network and Internet → Network and Sharing Center → Change adapter settings.

Right click the Network Interface Card (NIC) that you want to set up and select Properties.

Double click Internet Protocol Version 4 (TCP/IPv4).

Enter the proper information for the IP address, subnet mask, default gateway IP address. Optionally you can enter DNS server information.

Click OK twice.

To verify the previous set up, you perform the following:

Click the Windows start button and type command prompt.

Type ipconfig command.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY


Inter-VLAN Routing



You must be able to connect your various VLANs together using a routing device. Let's explore Inter-VLAN routing.

CONFIDENTIAL - FOR TRAINING

Multi-Layer Interface Types



MOD 3-19

A switch port can only run layer 2 features and protocols, like creating and assigning VLANs, for example. The figure shows ports 1 – 4 used as layer 2 interfaces. They connect to end systems and accept layer 2 frames as being a member of a layer 2 VLAN. The layer 2 frames from these ports forward it based on their destination layer 2 MAC addresses. You learned that on some AOS-CX switches, the ports are layer 3 interfaces by default, and therefore you must configure those switch ports, like ports 1 to 4 with the command `no routing` on the respective interfaces.

But what if you want to route between these VLANs? The AOS-CX switches are multi-layer switches: they have both internal layer 2 switching functions and internal layer 3 routing capabilities. You need a way to connect each layer 2 VLAN to the internal routing functions. To do this, you must either define the port as a routed port or create switch virtual interfaces (SVIs) for the VLANs. An SVI is a virtual layer 3, routed interface that exists only inside the switch as a virtual (logical) construct.

Suppose that you create interface SVI 10. Because it is an SVI, by definition, it connects to the internal routing construct. Because it is SVI “10”, by definition, it connects to VLAN 10, and so services routable traffic from VLAN 10 to other destination VLAN networks. Similarly, you might define SVI 20. With some routing configuration, which you will soon learn, your switch can now route traffic between your VLANs.

Now suppose that you need to connect your multi-layer switch to an external router, perhaps using port 24, as shown in the figure. Since all ports are layer 3 interfaces by default, port 24

connects to the internal routing functions by default. You merely need to configure it with typical layer 3 parameters, such as an IP address. You will soon learn about these concepts and syntax for creating SVI interfaces.

The SVI's are virtual layer 3 interfaces, for internal routing, and port 24 is a physical layer 3 interface, for external routing. Both are layer 3 interfaces, and so perform routing functions. They accept routable layer 3 packets and forward them based on their destination IP address.

At this point, you now know about three important interface types:

Layer 2 switch ports

Layer 3 SVI's

Layer 3 physical routed ports.

You are ready to learn about another very important interface type when dealing with inter-VLAN routing: a trunk port.


CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Inter-VLAN Routing


VLAN = Broadcast domain

Routers connect VLANs

Older solution: Switch and Router



Newer solution: Multilayer switch



MOD 3-20

You learned about VLANs in Module 2. A VLAN is a broadcast domain, with a unique IP network number. In other words, all devices in the same VLAN have the same network address. Examine the figure. All devices in the Sales VLAN are associated with network 10.1.10.x, where x is some unique host value. Therefore, IP addresses might be 10.1.10.100, 10.1.10.101, 10.1.10.102, and so on. Everyone in the HR VLAN might have host addresses like 192.168.20.100, 192.168.20.101, and so on.

Recall that devices in different VLANs cannot communicate directly unless you connect them with a router. Inter-VLAN routing connects separate VLANs into a routed internetwork of communicating devices.

Many years ago (1990s), multi-layer switches did not exist. We used layer 2 switches for host connectivity, and then routed between them with an external router, as shown in the right-side of the figure. This was called a router-on-a-stick configuration. You might still see this setup today, when a customer performs routing between network segments on an external firewall.

A potential problem with this deployment is that the switch-to-router link can become oversubscribed, although a LAG could alleviate this bandwidth problem to some extent. Performance can also be suboptimal, because sending frames to the router requires an additional switching operation.


Multilayer switches are more efficient devices. The switching and routing functions of the device are connected via a high-speed internal backplane. Initial routing decisions and other

processes all happen in the box, which reduces reduce latency and increases performance. All AOS-CX switches are multilayer switches which have routing enabled by default. The exceptions are the 6000/6100 series, which do not support routing, they are strictly layer 2 switches.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

IP Routing Table

Objective Used by routing devices to select the best path to a destination



Routing devices (routers and multilayer switches) build and maintain a routing table that informs them of the best path to any given destination. One option to build the routing table is where you manually add entries to the route table in the form of static routes. Alternatively, you can configure a routing protocol, which automatically builds and maintains this table. Typically, entries in the routing table do not expire unless a change in the topology causes an update. This differs from the MAC address table in layer 2 switches where an entry expires after 5 minutes if the switch stops receiving traffic from the endpoint.

You will learn about all the entries that can exist in an actual route table soon. Meanwhile, the figure shows a slightly simplified view of the route table on the Core-1 multi-layer switch. There are three networks, 10.0.0.0/8, 172.16.0.0/16, and 192.168.0.0/24. Recall that the subnet mask determines the network portion of an IP address.

Now let us analyze Core-1's routing table. The first entry says, "To get to any host on network 192.168.0.0/24, send packets to next hop router 172.16.0.253 (Core-2). To get to that next-hop address, forward the packet out local VLAN172".

The next entry says, "To get to any host on network 10.0.0.0/8, there is no next-hop...I'm directly connected to that network. Simply forward the packet out my local VLAN interface 10". Lastly, to connect to the network 172.16.0.0/16 Core-1 is directly connected to this network and therefore there is no need to assign a next-hop IP address: simply forward the packet out the VLAN interface 172.

Consider at this moment that you configure only Core-1. Core-1 will properly route the packet and send it to Core-2. Core-2 will have no problem to send the packets to Server-2, since this network is locally connected to Core-2. Therefore, communication from Server-1 to Server-2 is successful. However, the response from Server-2 is sent to Core-2. Core-2 receives the packet but since the destination (10.0.0.2) is not in its routing table, the packet will be dropped. Therefore, bidirectional communication cannot be done by only configuring routing on Core-1. To solve the problem, remember that Core-2 will need to be configured with a route to reach the non-direct connected network (10.0.0.0/8). Try to think about how this route would look like:

Destination = 10.0.0.0/8

Next-hop = 172.16.0.254

Interface = VLAN that is connected to network 172.16.0.0 (in this example, VLAN172)


This route would need to appear in Core-2's routing table (statically or dynamically) to allow for bi-directional communication between the two servers.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

DHCP Relay


DHCP Service

- Dynamically assign IP parameters to hosts
- DHCP server and Host must be in same network



DHCP helper address

- DHCP server and host can be in different networks
- Router converts broadcast to unicast



AOS-CX configuration

```
Switch(config)# interface vlan 1  
Switch(config-if)# ip helper-address 192.168.10.1
```



MOD 3- 22

You have learned to configure static IP addressing parameters on an endpoint. However, configuring static IP address assignment is tedious and not scalable. Typically, most endpoints automatically get IP addressing information using the Dynamic Host Configuration Protocol (DHCP).

To obtain IP addressing information dynamically, endpoints broadcast a DHCP request, “Hey everyone, I need an IP address, a subnet mask, a default gateway, and a DNS server address”. Because it is a broadcast, the host and server must be on the same subnet. Remember, a router defines the edge of a broadcast domain, and does not forward broadcast. DHCP servers in different broadcast domains (different VLANs) do not receive the request, and so no address is assigned.

Now wait a minute. If an organization deploys thousands of broadcast domains (VLANs), then you would need thousands of DHCP servers – one per VLAN! This is not realistic. Instead, you need a centralized DHCP solution for all your VLANs.

From a routing perspective, the solution to the issue of endpoints being in separate VLANs from the DHCP servers is to configure DHCP relay. Using DHCP relay, a DHCP helper address is configured on each router interface where DHCP broadcast requests would be received. That router interface typically serves as the default gateway for endpoints. The following process describes how the solution works, when a router is properly configured with a DHCP helper address.

The client broadcasts a normal DHCP query.

The router that is on the client's network (the client's default gateway) receives this broadcasted DHCP query.

After receiving the DHCP broadcast which would be normally discarded, a new DHCP unicast packet is created and forwarded to the DHCP server. This new packet has source IP address set to the IP address of the SVI interface on which the broadcast was originally received. This new DHCP unicast packet also sets the destination IP address to what is specified in the `ip helper-address` command (192.168.10.1, in this example). The destination address is the IP address of the DHCP server. Now that the message is a unicast packet, the router can forward it as it would any other unicast packet: towards its final destination, the DHCP Server.

Thus, the DHCP Server receives the DHCP request and replies with a DHCP offer message that contains the IP addressing details. This unicast message is sent to the requesting host, via the DHCP helper router. This router, in turn, forwards the response to the requesting endpoint device.

In AOS-CX, use the `ip helper-address` command to define the remote DHCP server IP address. Up to eight addresses can be defined. When more than one DHCP server has been defined on the layer 3 interface, the switch sends the request to all servers. This is a common practice in company networks, since multiple DHCP servers provide a redundant and scalable solution.

Here is the configuration for defining IP helper for an SVI:

```
Switch(config)# interface vlan 1 <VLAN-ID>
```

```
Switch(config-if)# ip helper-address <IP-Address>
```

If you had a layer 3 routed interface, you would execute the `ip helper-address` command under the layer 3 interface context. The bottom of the figure

illustrates how to do this for Network B SVI interface on the multilayer switch.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Static IP Routing



The simplest form of routing is probably static IP routing.

CONFIDENTIAL - FOR TRAINING

Route Types in ArubaOS-CX



- **Connected:** subnet is physically connected to the device
- **Local:** local IP address configured on the switch
- **Static:** route added manually as a static route
- **RIPv2:** route learned via RIPv2
- **OSPF:** route learned via OSPF (preferred over RIPv2)
- **BGP:** route learned via BGP



MOD 3- 24

You previously learned about the IP routing table, and that a routing device makes layer 3 forwarding decisions based on the best-path entries in this table. The entries in the routing table are populated manually or through a routing protocol. AOS-CX switches support multiple route types, as shown in the figure.

Let's look at some examples.

Route Types: Connected and Local Entries

```
SVI Configuration:  
Switch(config)# interface vlan 1  
Switch(config-if)# ip address 10.1.1.1/24  
Switch(config)# interface 1/1/1  
Switch(config-if)# vlan access 1
```



```
Switch# show ip route  
  
'[x/y]' denotes [distance/metric]  
  
10.1.1.0/24, vrf default  
    via vlan1, [0/0], connected  
10.1.1.1/32, vrf default  
    via vlan1, [0/0], local
```

MOD 3-25

Consider the SVI configuration in the figure. An SVI is configured for VLAN1 with an IP address, and it is assigned to port 1/1/1.

This configuration generates some entries in the routing table, as seen here with `show ip route`. You see two entries in the routing table.

- 10.1.1.0/24 is a connected route type: the subnet physically connects to the switch and therefore there is no need for a next hop device.
- 10.1.1.1/32 is a local route: the configured IP address participates in the switch internal routing process. Recall that local route types indicate an SVI, as shown here, or a loopback interface.

Note. You may notice that the local entry uses a /32 subnet mask. This is a special mask that indicates that subnetting cannot happen (no divisions are permitted) in the address space provided. This mask is useful to refer to individual hosts.


Static Routes

- Needed when destination is not connected to the switch
- Administrator manually defines the next hop

```
Switch(config)# ip route <remote-network/mask> <next-hop>
```

```
Core-1(config)# ip route 10.1.20.0/24 10.1.12.2
```

```
Core-2(config)# ip route 10.1.10.0/24 10.1.12.1
```



```
Core-1# show ip route static  
<<<Omitted output>>>
```

```
10.1.20.0/24, vrf default  
via 10.1.12.2, [1/0], static
```

MOD 3-26

You need a static route when the destination network is not directly or physically connected to the switch or router, and no routing protocol (like OSPF) advertises that route. With a static route, you manually define the destination network and the next hop in the path to the destination. This next hop IP address must be directly connected to the device you are configuring. Use the `ip route` command, as shown in the figure.

Suppose that endpoints in subnet 10.1.10.0 must communicate with others in 10.1.20.0. Core-1 switch has two physically connected networks: 10.1.10.0/24 and 10.1.12.0/30. However, the switch is not physically connected to the 10.1.20.0/24 subnet. Therefore, you add the route shown in the figure for Core-1.

Core-2 switch also needs a static route to complete the path. It only has entries in its routing table for 10.1.20.0/24 and 10.1.12.0/30 networks. You must add a new route for Core-2 for network 10.1.10.0, as shown.

Verify your efforts with the `show ip route static` command.

Administrative Distance

- A trust rating for route entries
- Used if a route is learned by multiple methods
- Lower distance is preferred

```
Core-1# show ip route

Displaying ipv4 routes selected for forwarding

'[x/y]' denotes [distance/metric]

10.1.20.0/24, vrf default
    via 10.1.12.2, [1/0], static
10.1.20.0/24, vrf default
    via 10.1.12.2, [110/200], ospf
```

Populating method	Admin Distance
Directly connected	0
Static Route	1
External BGP	20
OSPF	110
Internal BGP	200
Unknown or unused	255

MOD 3-27

Sometimes a router learns a route from two different sources. Maybe a router is running BGP and OSPF, and you have configured some static routes. Perhaps all three methods have learned about the network 10.1.20.0/24. Which source for routing information should the router trust? The BGP route, the OSPF route, or the static route(s)? The protocol or method with the lowest administrative distance is used.

This is very much like humans learning information. You are hopefully very close to your mother...you trust her. You might say she has a very low administrative distance. If some stranger gives you conflicting advice, you may choose to listen to your mother. The stranger has a higher administrative distance.

The table shows the AOS-CX administrative distances for several routing protocol.

In the figure, Core-1 has learned about network 10.1.20.0 from OSPF and from a static route. Core-1 uses the path specified by the static route, due to the lower, more trusted administrative distance. You can see the administrative distance value in the routing entry: it is the first number located in the square brackets "[]". For example, the static entry has an administrative distance of 1 and OSPF a distance of 110.

OSPF



The most common dynamic routing protocol used internally by companies is Open Shortest Path First (OSPF).

CONFIDENTIAL - FOR TRAINING

Scalability Issues

Static routes

- Suitable for simple networks
- No automated failover
- Human factor can cause outages
- Maintenance could be challenge

Dynamic routes

- Scale to any size
- Offer automatic failover
- Router exchanges improve availability
- Maintenance is simple



MOD 3- 29

Introduction to dynamic routing

Working with static routes is simple and easy for small networks that have a few subnets. However, when an organization has hundreds or even thousands of subnets, static routing is not efficient to manage. Since there is no automatic route advertisement, you must manually configure all the routing. And, there's no dynamic failover mechanism if the primary path fails. With static routing, you would have to manually configure these policies.


The human factor is also a big consideration. You can cause serious problems by misconfiguring a static route or adding it to the wrong router. This could result in layer 3 route loops, traffic blackholing, and lost connectivity.

With static routes, recovering from routing device failures would be slow. You must manually configure new alternative paths. This is tedious and inconvenient, especially when dynamic routing would automatically do this for you.

Dynamic routing protocols are far more scalable, handling thousands of routes across multiple routing devices. They automatically failover to alternate paths with little to no downtime. This is because running a dynamic routing protocol constantly exchange messages about changes in the network, which minimizes your managerial workload.

Interior and Exterior Gateway Protocols

Autonomous System Routers under the same administrative domain



MOD 3- 30

An Autonomous System (AS) is a collection of routing devices under a common administrative domain. Your Internet service provider (ISP) owns their own internal network: they have autonomy over this system, and it is their AS. Each corporation has autonomy over their network: it is their AS.

To route packets inside an AS, each company chooses an interior gateway protocol (IGP). An IGP is simply a routing protocol that runs inside an AS. Examples of IGPs include the Routing Information Protocol (RIP) and Open Shortest Path First (OSPF). Each company chooses the IGP that works best for them.


In the figure AS 1 has three routing devices. They use RIP to exchange network information to discover all paths inside the AS, and then find the best path for each network within the AS. AS 2 has a similar arrangement, but they have chosen to use OSPF.

These companies might be partners and want to share information. They connect to the internet, via an ISP that owns a completely different AS, like AS 3, for example. The two companies wish to route between AS 1 and AS 2. Robust, enterprise-class routing between these Autonomous Systems requires an Exterior Gateway Protocol (EGP). The only EGP currently in use today is the Border Gateway Protocol (BGP).

Distance Vector Routing Protocols

- An early, less sophisticated routing protocol
- Each router only aware of directly connected peers
- Examples: RIPv2, RIPv6

Slow **convergence**
Not scalable (max 15 hops)



MOD 3- 31

The earliest methods devised to exchange IP information between routing information are based on distance vector algorithms. Each router advertises its distance from each network, and the vector or direction packets must travel to get to that network. These algorithms were very simple in their construction and operation.

The distance is measured by number of hops or hop count. In the figure, how many routers must R3 “hop over” to get to network 10.0.3.0/24? The answer is 0 routers – it is directly connected to that network.

What vector or direction must R3 forward packets to reach 10.0.3.0/24? The answer is: exit port 1. It is as if the router says, “to get to any host on destination network 10.0.3.0/24, which is zero hops away, I must forward packets out my local port 1”. Using something like RIP or RIPv2, R3 advertises its distance and vector to R2.

R2 receives this and knows, “If R3 is 0 hops away, and I must hop over R3 to get to that network, then I am one hop away from that route”. To get to destination 10.0.3.0/24, which is one hop away, R2 must forward packets out local port 24, to next-hop router 10.0.2.1. R2 advertises its distance and vector to R1.

R1 receives this and knows, “If R2 is 1 hop away, and I must hop over R2 to get to the destination, then I am 2 hops away”. To get to 10.0.3.0/24, which is 2 hops away, R3 forwards packets out local port 23, to next-hop router 10.0.1.1”.

From this example, you can see that routers using a distance vector protocol are not aware of the entire network topology: they are only aware of what they are directly connected to, which includes their subnets/networks and their distance vector neighboring routers. This simplifies the configuration and route discovery process, but distance vector protocols, like RIP, historically have had two major limitations:

Slow convergence: When a failure occurs, distance vector protocols can take minutes to converge, depending on network size, complexity, and architecture.

Limited scalability: Each router can be no more than 15 hops away from any other router. Understand that you could still have dozens, or even hundreds of routers, as long as it is a relatively flat architecture.


Distance vector protocols include the Routing Information Protocol (RIP), RIP version 2 (RIPv2), and RIP Next Generation (RIPng). RIP and RIPv2 do not support security feature to protect the advertisement messages that are exchanged between routing devices. And because of scalability issues with distance vector protocols, it is not recommended to use them, even in a small network.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Link-State Routing Protocol

- Dijkstra algorithm calculates best paths
- Each router is aware of the entire topology
- Examples: OSPFv2, OSPFv3

Fast **convergence**
Very scalable



MOD 3- 32

Link-state protocols resolve limitations and disadvantages of distance vector protocols. Routers that run link-state protocols possess information about the network topology. Thus, they can independently calculate their routing table using the sophisticated Dijkstra algorithm in order to select the best path for every network destination.

Actual implementations of link-state protocols include Open Shortest Path First version 2 (OSPFv2), which is used for IPv4; OSPFv3, which is used for IPv6. These protocols provide faster convergence times compared with distance vector protocols. Another advantage is scalability.

Many years ago, one disadvantage of protocols like OSPF is that they created more CPU, memory, and bandwidth utilization than protocols like RIP. However, routers have become technologically faster and more capable over the years, and our approach to design more efficient network architectures has greatly reduced this concern. In other words, you shouldn't be concerned about the overhead involved with link-state protocols because of the processing capability that exists in routing devices today, like AOS-CX switches.

The ArubaOS-CX multilayer switches are designed to meet today's network requirements. You can view the OSPF capabilities on an AOS-CX switch with the `show capacities ospfv2` command.

```
Switch-6300# show capacities ospfv2
```

System Capacities: Filter OSPFv2

Capacities Name	Value


Maximum number of OSPFv2 areas configurable in the system	128
Maximum number of OSPFv2 interfaces configurable in the system	128
Maximum number of OSPFv2 interfaces per area in the system	128
Maximum number of OSPFv2 neighbors allowed across all VRFs	128
Maximum number of OSPFv2 processes supported across each VRF	1
Maximum number of routes in OSPFv2 supported across all VRFs	64000

Note. This training focuses on OSPFv2 with a single area design.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

OSPF Introduction

- RFC 2328 defines OSPFv2 for IPv4 routing
- Popular enterprise solution
- Link-state protocol: Topology-aware



MOD 3-33

RFC 2328 defines OSPFv2, which describes how to route IPv4 packets. OSPF does not use TCP or UDP; instead OSPF advertisements are placed directly inside an IP packet. Therefore, it does not have a TCP or UDP port number, it uses IP protocol number 89, as shown in the figure.


OSPF is a very popular enterprise IGP routing solution, due to its hierarchical scalability and security mechanisms. For example, OSPF peers can authenticate packet exchanges.

As a link-state protocol, OSPF routers advertise information about their connected layer 3 interfaces and networks (links) and the cost associated with that interface. Let's see how it works.

Router ID

Unique, 32-bit address for each router

Humans see "Server Switch" → Server Switch
Routers see 10.0.100.0 → **RID: 10.0.100.0**



Determine RID

1. Manually configured → Core-1 (config) # **router ospf 1**
Core-1 (config-ospf-1) # **router-id 10.1.100.1**
2. Highest loopback IPv4 address
3. Highest non-loopback IPv4 address

MOD 3- 34

The OSPF configuration begins with the creation of the OSPF routing process. There can be multiple OSPF processes running on the switches, each having a unique ID which is only locally significant. In the figure, OSPF process ID is set to 1. In addition, all OSPF routers require a unique router identifier (RID). This is a 32-bit IP address which is typically associated with an actual IP address on the routing device. Routers include their ID in all OSPF packets that they generate. Humans might look at the figure and see "Core-1, Core-2 and Sever Switch". The routers might identify themselves as 10.1.100.1, 10.1.100.2 and 10.0.100.0.

You could depend on each router to automatically generate an RID itself, or you can take control of the situation and manually assign RIDs with the `router-id` command. Automatic assignment is tempting, since it requires so little effort. However, many experienced engineers prefer to manually assign RIDs, due to certain documentation, troubleshooting, and management advantages.

AOS-CX uses the following sequence to determine the router ID:

If you manually specify the RID, then that is what the router uses.

If you do not specify the RID, the loopback interface with the highest IP address becomes the RID.

If no loopback interfaces exist, the regular non-loopback, active interfaces with the highest IP

address become the RID. Non-functional Interfaces that are in a down state are not considered.


Let's explore using a loopback interfaces as a router's RID.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Router ID: Loopback Interface

A logical interface, always up

- Enhances stability
- Routable
- Many use cases/advantages



```
Core-1(config)# interface loopback 0
Core-1(config-if)# ip address 10.1.100.1/32
```

```
Core-1(config)# router ospf 1
Core-1(config-ospf-1)# router-id 10.1.100.1
```



MOD 3- 35

A loopback interface is a logical interface that is always is in an up state (unless you manually disable it). This interface enhances stability for processes and protocols that depend on the interface status to work, like OSPF. If you use a physical interface as the router ID, and that link fails, the router ID changes to an address on another alternative active interface, which is sub-optimal. Since loopbacks are always up, there is no danger of instability or wasted resources due to churning RIDs. Therefore, use loopbacks to “rid yourself” of this potential issue.

The IP address associated to a loopback interface is routable, which means that external devices can initiate communication to it. As you gain more education and experience, you will continue to learn about the advantages of loopback interfaces: scalability, troubleshooting, and network management.

In the figure, interface loopback 0 was created and assigned an IP address. Once this has been completed, you would then configure the OSPF routing process, specify the loopback address as the RID. But you just learned that this would happen automatically anyway, so...why configure it?


What if someone on your team created a second loopback interface, with a higher IP address? At some point, when you would reboot the switch or restart the OSPF routing process, OSPF will then begin using the second loopback address instead of the loopback 0 address. This is not necessarily bad, but if you prefer to have consistent, predictable behavior that matches your network documentation and facilitates quick troubleshooting, best practice

is to not only configure the loopback interface, but also manually define that loopback IP address as the OSPF process's RID value as well.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

OSPF General Operation Overview

Phase	Objective	Result
1. Neighboring	Build a neighbor relationship	Neighbor table
2. Topology database	Exchange of routing information	Link State Data Base
3. OSPF routing table	Run SPF to get the best path	OSPF routing table
4. IP routing table	OSPF publishes results to Routing Table (FIB)	Routing table entries



MOD 3- 36

The figure provides an overview of four key phases of OSPF's operations:


1. Form neighbor relationships with directly connected peers to build a neighbor table: a list of directly-connected OSPF peers.
2. Exchange route information to build an OSPF Link State Database: a list of every path in the known topology, which is also called a topological database.
3. Run the Shortest Path First algorithm (Dijkstra's algorithm) on that topological database of every OSPF path to extract the best OSPF paths. This information is placed in the OSPF routing table.
4. Publish the best OSPF paths to the router's IP routing table. The IP routing table contains all routes learned from all routing protocols: OSPF, static, directly connected, and BGP. This is also known as the forwarding information base (FIB).

OSPF General Operation: Build a Neighbor Table

Neighbor table

Objective: Form relationship with OSPF peers

Method: Send hello packets, validate compatibility



MOD 3-37

Phase 1: Build a Neighbor table

First, OSPF neighbors introduce themselves to their directly connected neighbors on interfaces that OSPF enabled. During this process, they validate that they are compatible. In this method the OSPF routers send OSPF hello packets out each OSPF interface.

In the figure, each router is configured similarly (except for their IP address and perhaps interface number). First you create the OSPF process, and then you add appropriate interfaces to that process: SVI 2 in the case of Core-1 and Core-2. You enter the interface context and use the `ip ospf <process-ID> area <area#>` command to associate an interface to an OSPF process.

Thus, Core-1 sends hellos out all OSPF-enabled interfaces, including the one shown here, on subnet 10.11.0.0/30. "Hello. I am Core-1 and I think that we live in Area 11, a normal area, and that we do not need to use secure authentication. We are connected on subnet 10.11.0.0/30".

Core-2 sends similar information in its hello packet. If the criteria match, the OSPF routers agree to be neighbors: they form an OSPF neighbor relationship, which is sometimes referred to as an adjacency. If parameters do not match (typically due to misconfiguration), the routers refuse to form a neighbor relationship, and will not share routes and your network is broken. Let us assume that all routers in the figure successfully formed an adjacency with each directly-connected peer. This will be reflected in each router's neighbor table.

Note: Only a small part of hello packet contents are shown and discussed here, to convey the main idea. There is much more information exchanged in hello packets, about which you will soon learn.

Note. The term adjacency and neighbor refer to the same relationship between two OSPF routers in a broadcast domain. As you will soon learn, just because two routers form an adjacency does not necessarily mean that the two routers will share link/routing information.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY


Hello Messages

- Ensures 2-way peer communication
- Neighbor relationship forms if peers have **matching criteria** →

- Network number and mask
- Area ID and Type
- Timers
- Authentication

Verifying neighbors

```
Core-2# show ip ospf neighbors detail
Neighbor 10.11.100.1, interface address 10.11.0.1
-----
Process ID 11 VRF TABLE-11, in area 0.0.0.11 via interface
vlan110
Neighbor priority is 1, State is FULL
DR is 10.11.0.1, BDR is 10.11.0.2
Options is 0x42
Dead timer due in 00:00:30
Retransmission queue length 0
Time since last state change 00h:02m:39s
```



OSPF routers in the same VLAN/subnet or who are directly-connected will send hello packets to ensure two-way communications. Hellos are also used as a failure detection mechanism. By default, these packets are sent every 10 seconds to the multicast IP address 224.0.0.5. This is the reserved “attention all OSPF routers” multicast address, and this is associated with a multicast layer 2 MAC address of 01:00:5E:00:00:05. Recall that all layer 2 switches flood broadcast and multicast frames out all ports in the broadcast domain, so all directly connected routers exchange hellos (which are the first packets to be sent in the OSPF exchange process).

Remember that another main purpose of hello packets is to build and maintain a neighbor table. Peers only form a neighbor relationship if they are compatibly configured.

They must agree that they are on the same subnet and have the same subnet mask.

They must be in the same area and agree on the area type.

Their timers must match (10-second OSPF hello timer, etc.). Verify the hello interval with the command `show ip ospf neighbor detail`.

They must be configured to use the same authentication type.

The MTU size should match (this is not a requirement, but if it doesn't match, you can experience issues when sharing a large amount of routing information).

The AOS-CX default hello interval is 10 seconds. You can customize the interval from 1 to 65553 seconds. A typical hello interval is 1 second.

```
Switch(config)# interface <interface-id>
```

```
Switch(config-if)# ip ospf hello-interval <seconds>
```


The dead interval is time after which a neighbor is declared dead. Typically, this value is four times the hello interval, so 4 seconds for the dead interval if you use a 1-second hello timer. Note that the neighbor values must match for an adjacency to form. You can use the `show ip ospf neighbors [detail]` command to verify if the adjacencies have successfully formed.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

DRs and BDRs

Objective Centralize route dissemination in a subnet

- 1** DR/BDR election
 - Highest interface priority
 - Highest Router ID
 - No preemption
- 2** OSPF routers share updates with DR/BDR
- 3** DR/BDR shares updates to all routers in subnet
- 4** OSPF routers acknowledge the updates



An OSPF router will not form adjacencies to just any OSPF router. Instead, a client/server design is implemented in OSPF on each broadcast segment. For each multi-access broadcast segment, such as Ethernet, there is a designated router (DR) and a backup designated router (BDR) as well as other OSPF routers, called DROTHERs. As an example, if you have 10 VLANs/subnets in your switched area, you'll have 10 DRs and 10 BDRs. As you will soon learn, if you have point-to-point links (only two OSPF routers), then there is no advantage nor need for DR/BDRs.

The OSPF router with the highest priority becomes the DR for the segment. If there is a tie, the router with the highest router ID (not IP address on the segment) becomes the DR. By default, all routers have a priority of 1 (priorities can range from 0 to 255—it's an 8-bit value). If the DR fails, the BDR is promoted to DR and another router is elected as the BDR. If a router does not become a DR or BDR, it is referred to as a designated router other, or commonly DROTHER, for short. There is no preemption of the DR/BDR role. If a DR/BDR exist and a router with a higher router ID joins the segment, it will not preempt the existing DR/BDR. Election only occurs if there is no DR or BDR on a segment.

When an OSPF router comes up, it forms adjacencies with the DR and the BDR on each multi-access segment to which it is connected; if it is connected to three segments, it will form three sets of adjacencies. Any exchange of routing information is between these DR/BDR routers and the other OSPF neighbors on a segment (and vice versa). An OSPF router talks to a DR using the IP multicast address of 224.0.0.6. The DR and the BDR talk to all OSPF


routers using the 224.0.0.5 multicast IP address.

OSPF routers use link state advertisements (LSAs) to communicate with each other. One type of LSA is a hello, which is used to form neighbor relationships and as a keep-alive function. Hellos are generated every ten seconds. When sharing link information (directly connected routes), links are sent to the DR (224.0.0.6) and the DR disseminates this to everyone else (224.0.0.5) on the segment. On point-to-point links, since no DR/BDR is used, all OSPF packets are addressed to 224.0.0.5.

Note: The DR/BDR feature was specifically designed based on the hardware (or lack thereof) capabilities of routers in the 1980s. Today, this is a moot point and you should not be concerned which of your routers play which role: DR, BDR, or BROTHER.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

OSPF Neighbor States



Neighbor adjacency complete

MOD 3- 40

OSPF uses a finite state machine (FSM) to process the neighbor state transitions between routers, when certain conditions are satisfied. This process can be divided into two main phases:

Establish Neighbor Adjacencies

Synchronize OSPF databases

It is important to clarify that in a broadcast network, database synchronization only happens between the DR/BDR and DROTHER routers. In other words, DROTHER routers only establish neighbor adjacencies and stay in the 2-WAY state (which is considered a stable state) but never synchronize databases with other routers in a DROTHER state: however, they will move into a full adjacency with the DR and BDR, which allows them to share information about links and networks.

As an example, consider that two core switches attempt to become OSPF neighbors over a broadcast network.

Establish Neighbor adjacencies:

Core-1 and Core-2 start with the DOWN neighbor state.

Core-1 transits to the INIT state when it sends the first hello message. This message includes

Router ID = 10.1.100.1; the hello message includes a seen field, which includes the routers has Core-1 received hello messages previously. Since this is the first hello message the field has a NULL value.

Core-2 receives Core-1's hello message and responds, indicating that the message has been seen and values are compatible. Core-2 transits to INIT state.

Core-1 receives the hello message and transit to the 2-WAY state. Core-1 again sends again the hello message to Core-2 but this time it includes both Router IDs in the seen field.

Core-2 receives this message and moves to the 2-WAY state.

Let's see how this plays out in a real-world scenario.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

OSPF Neighbor States: Sample Event Log

```
show events -r -d hpe-routing
```

```
T13:11:39.000761+00:00 sw-aggr1 hpe-routing[7720]: Event2401LOG_INFOAMM1/1AdjChg:  
Nbr rtr ID 10.1.0.3 on IP addr 10.1.2.2( area ID 0.0.0.0): Exchange -> Full
```

```
T13:11:39.000105+00:00 sw-aggr1 hpe-routing[7720]: Event2401LOG_INFOAMM1/1AdjChg:  
Nbr rtr ID 10.1.0.3 on IP addr 10.1.2.2( area ID 0.0.0.0): Exstart -> Exchange
```

```
T13:11:38.999138+00:00 sw-aggr1 hpe-routing[7720]: Event2401LOG_INFOAMM1/1AdjChg:  
Nbr rtr ID 10.1.0.3 on IP addr 10.1.2.2( area ID 0.0.0.0): Two-way -> Exstart
```

```
T13:11:38.999093+00:00 sw-aggr1 hpe-routing[7720]: Event2401LOG_INFOAMM1/1AdjChg:  
Nbr rtr ID 10.1.0.3 on IP addr 10.1.2.2( area ID 0.0.0.0): Init -> Two-way
```

```
T13:11:38.998994+00:00 sw-aggr1 hpe-routing[7720]: Event2401LOG_INFOAMM1/1AdjChg:  
Nbr rtr ID 10.1.0.3 on IP addr 10.1.2.2( area ID 0.0.0.0): Down -> Init
```

<date stamp removed for clarity>




MOD 3-41

The figure shows an example of the neighbor process happening in a real-world environment, time stamped from the bottom, upwards (the newest message is first and the oldest is listed last based on using the `-r` option with the `show events` command).

1. First, sw-aggr1's neighbor, 10.1.0.3, moves from down to an init state: a hello was sent to that neighbor.
2. Next, sw-aggr1 receives a hello from that neighbor, and so the state moves from init to 2-way.
3. If this router or the neighboring router is either a DR or BDR, the two neighbors move to an exstart state: they begin to exchange LSDB. This is covered shortly.
4. The two neighbors will then move to the exchange state to share missing information with each other until each has a matching LSDB.
5. Once this is complete, they are in a full state. You might say that these OSPF peers have formed a full adjacency. This is covered shortly.

OSPF Neighbor States



MOD 3-42

In the example, Core-1 and Core-2 are the only two devices in the network, and so they continue to the database synchronization process.

Synchronize OSPF Databases.

6. Core-1 initiates the synchronization process by sending a database description (DBD) packet. The switch transits to the exstart neighbor state.

7. Core-2 does a similar process: it moves to the exstart state by sending a DBD packet.

The goal of the exstart state is to determine which switch will become the master for the exchange process, based on the highest router ID. Understand that the master role defines only which switch will initiate the database exchange process. This role has nothing to do with DR/BDR: it is only for Link State Database (LSDB) exchange. In this example, Core-2 has a higher router-ID and becomes the master.

8. Core-2 sends another DBD packet as it transitions to the exchange state. Core-2 is now sharing a summary of the contents of its LSDB with Core-1.


9. Core-1 also sends a DBD packet and moves to the exchange state, sharing its LSDB with Core-2.

After several exchanges are performed during the exchange state, each router will have a

summary of the others' entire LSDB. They compare the information received and, in the next step, request the missing information.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

OSPF Neighbor States



```
Core-1# show ip ospf neighbors
```

Neighbor ID	Priority	State	Nbr Address	Interface
10.1.100.2	1	FULL/DR	10.11.0.2	LAG10



MOD 3-43

10. Core-1 and Core-2 request the missing information using link state update (LSU) packets. The other peer answers by sending the requested database information. Both switches transition to the loading neighbor state: this state is the actual routing information exchange.


11. Core-1 and Core-2 move to the full state when there is no more information to be exchanged. Now, both devices have the same LSDB information.

You can verify the OSPF neighbor state using the command `show ip ospf neighbors`.

OSPF Network Types


Point to Point

Two peers on link



Broadcast

Two or more peers on link



Configure

```
Switch(config)# interface <interface-id>  
Switch(config-if)# ip ospf network {broadcast point-to-point}
```



MOD 3- 44

From the protocol perspective two options exist for network types:

Point-to-Point Network: Only two peers are on the link. When an interface is configured to be part of a Point-to-point link, OSPF knows that a single neighbor device is expected on the interface. PPP Serial links (not used too much in networks now) is an example of this network type.

Broadcast Networks: Two or more peers might be on the link. When an interface is configured to be part of a broadcast network, OSPF knows that more than one neighbor might be discovered on the interface.

AOS-CX interfaces default to the broadcast network type, but you can modify this with the `ip ospf network` command on the respective interface.

```
Switch(config)# interface <interface-id>
```

```
Switch(config-if)# ip ospf network {broadcast point-to-point}
```

To verify the type of network that is in use you can use the `show ip ospf interface` command.

Unless you truly have multiple routers in the same broadcast domain, it is recommended to configure switch-to-switch links with the point-to-point network type.

Note. OSPF can also use the non-broadcast multiaccess (NBMA) network type. This network

can support multiple devices (multi-access) but does not support the broadcast capability. Frame Relay is an example of this network type. The use of this type of network is deprecated for modern networks, and so will not be covered in this course.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Validate the Neighbor Table

```
Core-1# show ip ospf neighbors
```

```
VRF : default
```

```
Process : 1
```

```
=====
```

```
Total Number of Neighbors : 1
```

Neighbor ID	Priority	State	Nbr Address	Interface
10.1.100.2	255	FULL/DR	10.11.0.2	vlan2



MOD 3- 45

Show ip ospf neighbor command allows you to verify the neighbors.


A state of full indicates that the router can share link information with the peer. In the figure, the two routing switches are in a broadcast medium and the neighbor is the DR. If this was a point-to-point link, you would want to define the point-to-point link type for the network type to speed up the initial startup of OSPF (convergence), since no DR/BDR election process is necessary on a point-to-point link. If you see 2-way as the state, this indicates that the neighbor is a DROTHER: all the required items match between the two peers, but no routing information is shared. In this instance, routing information is only shared with the DR/BDR peers.

OSPF General Operation: Build a List of Every Link

Build topology database

Objective: Learn every link, router, and interconnection

Method: Send/receive multicast LSAs



MOD 3- 46

Phase 2: Build a topology database, also called a Link State Database (LSDB)

A router shares known subnets with other routers if and only if a neighbor adjacency is fully formed. The objective is to build a topology database: a list of every single link (subnet), every single router, and exactly how those routers and subnets are interconnected. We humans look at a network diagram, in a pictorial format to understand how router and links are interconnected. Routers look at the LSDB (in a numerical format) to understand this. The topology database is referred to as a link state database (LSDB).


The method involves sending special OSPF multicast messages called link state advertisements (LSAs). The Server Switch sends LSAs out all interfaces, "Attention all OSPF routers. I am Server Switch, and I am directly connected to 10.1.1.0/24, 10.1.2.0/24, 10.20.0.0/22". Every other OSPF router receives these LSAs and adds the information to its topological database. Soon, every router has received LSAs from every other router. Thus, each router has a full topology database. Again, this database is essentially the diagram you see in the figure, in a numerical format. There is only one reality from the routing switches' perspective, which is shown in the figure. Therefore, every OSPF router has an identical topological database; and basically, the LSDB tells you who is connected to what.

Note: The figure only shows Server Switch advertising its directly-connected subnets, because that is all that it currently knows. However, once Server Switch receives LSAs from the other routers (as they have just received from Server Switch) it will advertise those routes as well. Routers advertise the entire contents of their topology database to all other routers.

The topology database is also known as the link-state database (LSDB).

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

OSPF LSAs



MOD 3-47

OSPF routers generate different types of link state advertisements (LSAs), each with a different purpose and scope in the OSPF routing domain. Each LSA includes information that describes the portion of the network that a router announces. LSAs are collected to make up the LSDB, which represents the topology of the network. This will be used by OSPF to calculate best path to every destination in the network. When OSPF routers establish adjacencies, they exchange LSDB databases so every router in the network is aware of the entire topology.

In AOS-CX you can verify this information by using the `show ip ospf lsdb` command. This is a powerful troubleshooting command. If your network diagram shows that there are five routers, but you only see three listed, you know that there is a problem.

OSPF General Operation: Build a List of Best Paths

Build OSPF table

Objective: Build the best paths for each destination

Method: Run SPF (Dijkstra) algorithm on topology DB

Two paths in the LSDB to 10.20.0.0/22

- Out LAG10 to next-hop Core-2
- Out 47 to next-hop Server Switch

Run SPF Algorithm


Lowest cost: Out 47 to next-hop **Server Switch**

OSPF Routing TABLE

IP Destination	Next Hop	Interface
10.20.0.0/22	10.1.1.1	47

Server Switch

Server subnets
10.20.0.0/22



MOD 3- 48

Phase 3: Build a route table

Although there is only one topology, each router's position and connectivity within that topology is unique. Thus, each router's objective is to determine the best path to each link from where they sit within the topology. The method is to run the SPF (Dijkstra) algorithm on the topology database or LSDB.

For example, examine the network shown in the figure. Core-1's SPF algorithm analyzes the topology database and sees that Core-2 could reach destination 10.20.0.0/22 by routing packets out port LAG10. And Server Switch can reach the same destination via port 47. The algorithm considers the cost associated with each path (based on bandwidth) and chooses the lowest cost (fastest) path. This best path is added to the list of best paths – the route table.

In the example the best path is to send the traffic directly to Server Switch.


OSPF General Operation: Build a List of Best Paths

Build FIB table

Objective: Publish the best path on the FIB

```
Core-1# show ip route
```

```
10.20.0.0/22, vrf default  
via 10.1.1.1, [110/200], ospf
```




MOD 3- 49

Phase 4: Build a route table

The best path previously calculated by OSPF is published to the routing table (FIB). In this case the best path from Core-1 to the server subnets uses the path through Server Switch.

OSPF Area

- Group of OSPF routers with same LSDB
- Hierarchy improves scalability



MOD 3- 50

OSPF routing process

This section rounds out the rest of the OSPF configuration. Previously in the “OSPF introduction” section, you learned how routers obtain their RIDs and how to manipulate this. Without an RID, the OSPF process on an AOS-CX switch won’t even start. Therefore, using the IP address on a switch’s loopback is the preferred method of assigning an RID. The rest of the section will discuss an introduction to OSPF configuration.

OSPF areas

An area is a group of OSPF routers that share the same LSDB. All router must be part of an area. When you split a large network into separate areas, you reduce the size of the LSDB in each router, lower CPU utilization, and increase overall network stability. This is because each router in an area must only maintain the topology for that area.

For example, SW1 and SW2 in the figure need not learn about the entire topology: they only needs to know about area 10 topology. Any other destination is reached via the area border router(s), called ABRs. The ABRs in the figure are connected to two areas – 10 and 0 – and thus have two LSDBs, one for each area.

Routers SW1 and SW2 are called internal routers: all their interfaces are in a single area. If an internal router must route outside of its area, it simply forwards packets to its Area Border Router (ABR). An ABR is a router connected to two or more areas. SW3 is the ABR for area

10, with two interfaces in area 10 and two interfaces in area 0.


By definition, all areas must connect directly to the special backbone area, called area 0. This is the most important area in the hierarchy, and so should have a redundant design. Communication between two non-backbone areas **MUST** communicate via the backbone area 0. Routers with interfaces in area 0 are called backbone routers. Routers SW4, SW5, and SW6 have all interfaces in a single area, and are therefore called internal routers. Routers SW4, SW5, and SW6 are also backbone routers, however, SW3 is not an internal router since it has interfaces in two different areas.

The area ID is a 32-bit address that can be written in dot-decimal notation (area 0.0.0.0), or in decimal notation (area 0) , as shown in the figure. AOS-CX supports both notations.

This course is only focused on single-area designs, in which all router interfaces are in area 0. You can learn more about hierarchical OSPF solutions in the advanced courses.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Path Selection



The path cost is the “metric” – that which is measured to determine the best path

MOD 3- 51

After all routers have successfully exchanged LSAs and LSUs within an area, they have an identical LSDB. Remember, the LSDB is a list of every link, every router, and how those routers and links are interconnected. It is a list of every path to every destination in the area.

Next routers must run the SPF (Dijkstra’s) algorithm to find the best OSPF paths to each destination subnet. The best path is the one with the lowest cost, and cost is based on bandwidth. Therefore, if there are multiple paths to a single subnet, OSPF chooses the path with the lowest cumulative cost: the fastest path.


Consider the topology shown in the figure. Core-1 has two paths to reach destination subnet 10.20.0.0/22. To determine each path’s cost, simply add the values indicated for each link. Here’s what you get:

Using the Server Switch as the next hop, the cost is $100 = 100$.

Using Core-2 as the next hop, the cost is $100 + 50 = 150$

Obviously, using Server-Switch results in the lowest cost, and so that path is added to the route table.

OSPF Convergence



MOD 3- 52

There are two components to OSPF routing convergence:

Detecting topology changes

Recalculating and selecting the next best OSPF routes

Topology change detection is supported in two ways by OSPF. The first, and quickest, is a failure or change of status on the physical interface on a directly-connected link. The second is a timeout of the OSPF hello timer. An OSPF neighbor is deemed to have failed if the time to wait for a hello packet exceeds the dead timer, which defaults to four times the value of the hello timer. The default hello timer is 10 seconds, so the default dead timer is 40 seconds. The dead interval timer is necessary if the two OSPF routers are not directly connected.

When a change is detected, an LSA is sent to all routers in the OSPF area to signal the topology change. In the example shown, the link between Core-1 and Sever Switch has failed. Server Switch and Core-1 detect this outage: assuming you've configured this as a point-to-point link, their links go from an up state to a down state. Thus, Server Switch and Core-1 originate the topology change LSA. Upon receiving the update, all the OSPF routers update their LSDB and rerun the SPF algorithm to calculate their new best paths to any affected network, such as 10.20.0.0/22. The Server Switch will not change anything since it is still directly connected to the best path. Core-1, on the other hand, will recalculate its best path for 10.20.0.0/22. In this case, it will use the path via Core-2.


Ultimately, Core-2 receives LSA's about the change, but no recalculation is needed from its point of view since Core-2's best path is still it's connection to the Server Switch.

Each router performs route recalculation after a failure has been detected. This causes all routers to recalculate all their routes using the Dijkstra (SPF) algorithm.

Note: Because it is much quicker to detect a downed link than using a dead interval timer, it is highly recommended to use and configure OSPF point-to-point links whenever possible.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Equal Cost Multi-Path (ECMP)




MOD 3- 53

Equal Cost Multi-Path (ECMP) is a mechanism that enables load sharing of two or more equal-cost paths to the same destination network. In the figure there are two equal-cost paths Core-1 can use to reach subnet 10.1.0.0/24. Path 1 is via router with RID 10.0.100.2 and path 2 is via router with RID 10.0.100.3. When multiple paths exist to the same subnet that have the same cost AOS-CX router will perform load sharing using ECMP. ECMP is similar to link aggregation layer 3 based hashing: packets with different source/destination IP address are assigned different paths.

Passive Interfaces

- Router stops sending/receiving OSPF Hello packets *out that interface*
- Router continues to advertise *about that interface* to others
- Useful for endpoint-facing interfaces

```
Switch(config)# interface <interface-id>  
Switch(config-if)# ip ospf passive
```



MOD 3- 54

The OSPF configuration involves enabling the protocol on logical and physical interfaces. Thus, the router can generate LSAs to advertise subnets to other routers. This implies that the router sends periodic Hello messages on all OSPF-enabled interfaces.

In some cases, this is not desired. The most common example is when only hosts exist on that subnet, commonly referred to as a stub link in OSPF. However, the OSPF router would still constantly send OSPF packets on a network where nobody (the user hosts) will process them. Hosts do not respond to OSPF multicast 224.0.0.5 and 224.0.0.6. It is simply wasted processing cycles and bandwidth on the link. Also, if bad actors are on the link, they could learn information about your network, and potentially mount an attack.

The solution is to configure these host-facing router interfaces as passive. When an interface is passive, it stops sending and receiving OSPF packets out that interface. However, it continues to advertise about that interface to other routes. In the example, Switch-1 does not send hello packets out its 10.1.0.0/16 network, but it will tell Switch-2 about that network. Switch-2 has a similar arrangement, as relates to the 10.2.0.0/16 link.

In AOS-CX you can use the `ip ospf passive` command in the interface context to enable a passive interface.

```
Switch(config)# interface <interface-id>
```

```
Switch(config-if)# ip ospf passive
```



Knowledge Check

Self-check on key learning points



Please check your understanding with a brief knowledge check.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Question #1

Given the IP address 172.16.3.42, and a subnet mask of 255.255.255.0, what is the host portion of the address?

The host portion of the address is **42**

a


Knowledge Check

MOD 3- 56

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Question #2

Match the TCP/IP stack layer to an appropriate description.



Knowledge Check ✓




MOD 3- 57

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Question #3

Refer to the diagram. Server-1's ping packet reaches Server-2. However, Server-2's reply packet does not reach Server-1. Which of the options below is a likely cause of this issue? (Select two.)

- a. Server-1's default-gateway is not properly configured.
- b. Server-2's default-gateway is not properly configured.
- c. Core-1 has a route to reach Server-2 but lacks a return path to Server-1.
- d. Core-2 has a route to reach Server-2 but lacks a return path to Server-1.



a

Knowledge Check


MOD 3-58

B,D

Question #4

See the diagram. On router Core-1, your colleague has entered the command **ip route 10.1.20.0/24 10.1.12.2**. Which of the following options might you use to discuss this situation with your colleague?

- a. It appears that Core-1 will be able to route packets to subnet 10.1.20.0, which uses a mask of 255.255.224.0
- b. If devices on subnets 10.1.10.0/24 and 10.1.20.0/24 cannot communicate, it might be because you specified an incorrect next-hop route.
- c. If devices on subnets 10.1.10.0/24 and 10.1.20.0/24 cannot communicate, it might be because you need to configure a return path on Core-2.
- d. It appears that host 10.1.12.2 will be able to reach devices on subnet 10.1.10.0/24.



Knowledge Check

MOD 3-59

Question #5

Which statements below are true about the Open Shortest Path First (OSPF) routing protocol? (Select two)

- a. It is less scalable than RIP.
- b. For simplicity, it is best to let the router choose its own Router ID (RID).
- c. You can leverage loopback interfaces to enhance OSPF stability.
- d. If you do not configure a Router ID (RID), the OSPF router prefers to use the numerically highest IPv4 address assigned to a loopback interface.

a


Knowledge Check

MOD 3- 60

C,D

Question #6

Match an appropriate description to each of the OSPF entities below.



Knowledge Check ✓



MOD 3- 61

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY



Lab Activity

Consider the lab activity.

CONFIDENTIAL - FOR TRAINING

Lab Tasks

Lab 3.01
Basic IP Configuration

Lab 3.02
Edge Switch Management

Lab 3.03
Static Routes

Lab 3.04
Basic OSPF Configuration

**Module 4 –
Network
Redundancy**



MOD 3- 63

The figure provides a brief overview of labs for this module. Please see your lab guide for details. When complete, you will continue with Module 4– Network Redundancy.

Rev.22.41


aruba
a Hewlett Packard
Enterprise company

Network Redundancy

Aruba Campus Access
Fundamentals

Hello and welcome back. This is Module 4: Network Redundancy.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY



Objectives

- Discuss Spanning Tree
- Explain VRRP and VSX



Of course, you must maximize the resiliency of your vital corporate network systems. After completing this module, you will be able to describe the open standards for spanning tree and VRRP, as well as Aruba's VSX feature.

Spanning Tree

We begin with one of the first redundancy protocols, the spanning tree.

CONFIDENTIAL - FOR TRAINING

Redundant Network



MOD 4- 4

Network communications is often mission critical. Layer 2 switch or link failures cannot be tolerated. For example, the left-hand in the figure shows a single path to connect two access switches using Core-1. This design has single points of failure. If Core-1 fails, the network is down: Host A and B cannot communicate.

One common way to mitigate this is by adding a redundant Core switch, as in the figure's right-hand example. In this scenario, if the Core-1 fails, the network remains connected. Hosts A and B can still communicate over the redundant link using Core-2. While this redundant link mitigates a single point of failure issue, it creates a new challenge. The next three sections cover the main issues of layer 2 redundancy:

Layer 2 loops



Broadcast storms

Multiple frame copies

Layer 2 Loop

Redundant links create L2 loops, which cause problems


- Broadcast storms
- Multiple frame copies
- Instability of the MAC address table



MOD 4-5

Connecting layer 2 switches with redundant links creates layer 2 loops. A loop is even created if you connect an Ethernet cable from one port on a switch to another port on that same switch. The figure shows two different ways that result in layer 2 loops. If not properly handled, these loops cause serious problems that can effectively disable a network. These loops can cause broadcast storms and multiple frame copies, which can in turn, cause instability of the MAC address table of the switches and CPU utilization.

Broadcast Storms



Broadcast storm

- Broadcast frames circle the network for eternity
- Waste of Bandwidth and CPU resources
- Network is down



MOD 4-6

You learned in module 3 that switches flood broadcasts, as well as multicasts and unknown destination, out all interfaces in the same VLAN, except the ingress port. This flooding cause problems on a redundant, looped network like the one in the figure.

Let's look at this example more closely:

Suppose that Host A sends a broadcast frame.

Switch Access-1 receives this frame, and floods it out (copy 1).

Core-1 and Core-2 receive this and floods it out their connections to each other (copy 2).

Core-1, Core-2, and Access-2 receive the second copy and flood them out their other ports (copy 3).

Access-1 receives the third copy, and the cycle continues.

This frame circles around the network for all of eternity.

And remember that Access-1 flooded this broadcast to Core-2. So, a copy circles the network in the other direction, forever.


As a natural part of network and endpoint operation, nearly all devices send broadcast frames, often many times per minute. Every broadcast from every device circulates around

the network forever.

Soon, all available bandwidth and CPU cycles are used up processing broadcasts. No resources are available to process normal data communication frames. Eventually, this will consume all your bandwidth and drastically affect the performance of the devices in the network: the network is effectively unusable!

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Multiple Frame Copies



```
Access-1# show mac-address
```

```
MAC age-time      : 300 seconds
```

```
Number of MAC addresses : 2
```

MAC Address	VLAN	Type	Port
90:20:c2:bc:ee:00	20	dynamic	1/1/21
90:20:c2:bc:ee:00	20	dynamic	1/1/22

Multiple frame copies sent by Access-2 create MAC database instability on Access-1

MOD 4-7

You may recall that switches not only flood broadcasts out all ports in the VLAN (except the ingress port), but they do the same thing with unicasts to an unknown destination. If a switch receives a unicast packet to some destination MAC address, and that address has not yet been learned in the MAC address table, then the switch floods the frame out all ports in the VLAN (except ingress port). These can also circulate and waste bandwidth and CPU cycles. These frames also cause other issues.


Suppose that Host B sends a frame to Host A (see figure). Access-2 has not yet learned Host A's MAC address, and so it floods this frame out ports 21 and 22. Thus, Core-1 and Core-2 both switches receive a frame with source MAC 90:....00, inbound on port 21. Both cores send a copy to Access-1 switch, which receives it on port 21 and port 22. Access-1 now believes that it can reach Host B via port 21 and 22.

Access-1 is thus confused: Host B can only exist in one place, but it appears to exist in two places. Basically, Host B's MAC address will constantly switch in the MAC address table from port 21 to 22 and back...and so forth as Host B's frames arrive on Access-1 from Core-1 and Core-2. The multiple frame copies generated by Access-2 creates MAC database instability on Access-1. As you can see, this is a serious problem. You need redundant links for reliability, but this redundancy can bring your network to its knees.

Spanning Tree – Operation Overview

Goal Maintain a redundant loop-free topology

- Elect a root switch
- Connected tree of switches emanates from this root
- Redundant ports in blocking mode



MOD 4-8

It is time to learn about the solution to these challenges: the spanning tree protocol (STP). STP was invented in 1984 by the brilliant Miss Radia Perlman. Legend has it that she created the algorithm for this protocol in a few hours. She was so pleased with her creation that she then wrote a poem about it, and then when home for the day.

The IEEE 802.1D standard version of the STP was developed to build and maintain redundant, yet loop-free networks. With STP, you eliminate single points of failure, while avoiding loops and MAC table instability. STP creates a loop-free topology by automatically disabling redundant links.

The figure shows a highly redundant network, for resiliency and fault-tolerance. These redundant links could cause loops. But once STP is engaged, certain redundant links are automatically disabled: there are no loops, no broadcast storms, no MAC address table instability. These disabled links remain available to provide redundancy as needed.

In STP, logically the redundant paths are disable, but physically they are up. Originally STP referred to the redundant ports (and paths) as being placed in a blocked state. This changed in the STP standards and these ports, today, are referred to be in a discarded state.

To do this, one switch in the Spanning Tree domain is elected as a root bridge or root switch. Many trees grow out of the ground, from its roots, as a single trunk, and then branch out from there. Likewise, the root bridge is the reference from which the spanning tree grows. All other switches are non-root bridges (switches), sometimes knows as designated switches. A loop-


free path “grows” from this root switch out to all non-designated (non-root) switches.

Note: What is now called a switch is a very fast, multiple-port version of what used to be called a bridge (a few decades ago). When discussing STP, you will often see the term bridge. In this context, bridge and switch can be thought of as the same thing. When you see bridge, know that the discussion is actually about a switch. Both the term root bridge and root switch may be used in this course, as they are considered interchangeable.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Spanning Tree protocol: Convergence

- Happens when a failure occurs
- Loop-free topology is always maintained



The STP algorithm runs on all layer-2 switches, and STP will logically disable redundant links to avoid loops (see left side of figure). However, if a failure occurs, STP converges on a new topology of active links, which are used to forward frames (see right side of figure). A loop-free topology is thus maintained.

Note: The word converge can mean something like “to meet at some point”. Convergence, in the context of networking, means that the devices all come to an agreement about a new network topology. All devices converge on a new set of active paths to be used for frame forwarding.

Spanning Tree versions

STP

802.1D, pre-2004

- Original protocol
- Slow convergence
- Obsolete, Not recommended

RSTP

802.1w merged to
802.1D 2004

- Faster convergence
- Costs that fit with modern port speeds
- Default *operation* between AOS-CX switches, when no MSTP region is set up

MSTP

802.1S merged to
802.1Q 2005

- RSTP behavior in multiple instances
- Default *mode* for **AOS-CX** switches (but no MSTP region settings)



MOD 4- 10

This section introduces the components of the STP standard, including the different STP version, including rapid STP (RSTP), bridge identifiers (BIDs), bridge protocol data units (BPDUs), port states, path costs, and port roles.

There are three main versions of the STP protocol, see figure.

Failure detection in the original standard is based on timers. With the IEEE 802.1D default configuration, the root switch (and only the root switch) originates a hello packet every 2 seconds. These hello packets are forwarded on out to the rest of the switches in the domain. If some switch downstream (farther away from the root switch) ceases to receive hello packets for 20 seconds (the default max age timer for STP), an outage is assumed, and all switches begin to converge on a new topology. The duration of this process is defined by the forward delay timer, which is 15 seconds by default. This protocol (802.1d) is considered obsolete and its use is no longer recommended in modern networks.

The rapid spanning tree protocol (RSTP) was developed in 1998 to speed convergence. Instead of only the root switch originating hello packets, all switches originate them. This means that RSTP now has a true keep-alive mechanism that can respond in seconds (or less). No need for old and slow max age and forward delay timers. You will soon learn the details of this new convergence process.

Note: RSTP is backward compatible with the original standard 802.1D. However, to maintain this capability, some benefits of RSTP are lost.

802.1s or the multiple spanning tree protocol (MSTP) was developed to improve the performance of STP for implementing multiple loop-free topologies or instances to load balance the traffic across all links. This protocol helps to create optimal paths and increase bandwidth. MSTP runs on AOS-CX switches by default. However, if no setting is configured then this protocol behaves like RSTP. You will understand why this happens later in this module.

Note:


AOS-CX 6300 Spanning-Tree is globally enabled by default

AOS-CX 8325 Spanning-Tree is globally disabled by default

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Rapid Spanning Tree: Key Components

- Bridge Identifier (BID)
- Bridge Protocol Data Unit (BPDU)
- Path Cost
- Port States
- Port Roles



MOD 4- 11


Let's explore key elements that the RSTP algorithm uses to decide which ports shall be enabled for traffic forwarding, and which ones will not. This includes:

- Bridge Identifier (BIDs)
- Bridge Protocol Data Unit (BPDUs)
- RSTP port states
- RSTP port costs
- RSTP port roles

Bridge ID

Lowest Bridge ID wins the root election

- Unique identifier for each RSTP bridge
- Composed of a Priority and MAC address
- Value expressed in multiples of 4096
- AOS-CX use by default a priority of 32,768




MOD 4- 12

STP assigns each switch a unique identifier called Bridge ID (BID). This identifier is composed of a 2-byte priority and 6-byte MAC address. The priority defaults to 32,768, so by default, all switches have the same priority. However, each switch has a unique MAC address, and so each BID will be unique. The switch with the lowest ID wins the election to be the root switch.

So, which switch would be the root in the figure? Remember, you have to look at both the combined priority and MAC address of each switch, which becomes its BID. In the figure, the Core-1 switch is elected as the root bridge. Obviously, an administrator changed the priority on Core-1, since on AOS-CX switches, the priority defaults to 32,768.

Bridge Protocol Data Unit

- All switches exchange BPDUs every 2 seconds
- Acts as a keep-alive mechanism



MOD 4- 13

All switch participants in the STP algorithm exchange control messages called bridge protocol data units (BPDU), sometimes referred to as hello frames.

In the original 802.1D standard, BPDUs were generated only by the root switch, and then had to trickle down to other switches. This led to the need for the long max age and forward delay timers, which greatly slowed convergence.

With RSTP, all switches originate BPDUs with its current information, every 2 seconds (the default hello-time period). Thus, if a port stops receiving BPDUs for three consecutive hello timers, the switch quickly knows that it has lost connectivity to its neighbor. It ages out the protocol information and begins to converge.

Because each switch originates BPDUs, this becomes a true keep-alive mechanism. Failure detection will take no longer than 6 seconds.

Port States

STP (802.1D) Port State	RSTP (802.1w) Port state	Port Included in Active Topology	Port learns MAC addresses
Disabled	Discarding	No	No
Blocking			
Listening			
Learning	Learning	Yes	Yes
Forwarding	Forwarding		



MOD 4- 14

During original tree establishment and any ensuing convergence, switches must decide which ports will forward data and which ports must be disabled to prevent layer 2 Loops. STP uses port states to transit from a blocking/disabled port to a forwarding port state. The table in the figure summarize port states and their specific task. You see the states that were used on 802.1D as compared to 802.1w port states.

To transition from blocking to forwarding, the original 802.1D standard takes 30 seconds: 15 seconds are spent in a listening state, and 15 in the learning state. One of the reasons RSTP is more efficient than 802.1D is because the port transitions quickly from discarding to learning to forwarding.


Note: Blocking and listening state are deprecated in RSTP.

Note: In RSTP the learning state is transitory and is only used during re-convergence upon a topology change. Discarding or forwarding states are stable. Also, a discarding port is not physically disabled. It can exchange STP and LLDP BPDUs but blocks all other traffic.

Path Cost

Link Speed	Default Port Cost in ArubaOS-CX
100 Mbps	200,000
1 Gbps	20,000
10 Gbps	2,000
100 Gbps	200
1 Tbps	20

Direct Path	Indirect Path
0 + 20,000 = 20,000	0 + 2,000 + 2,000 = 4,000



BPDUs advertise the path cost

MOD 4-15

RSTP may have several possible paths to get from the root switch to some non-root switch. It chooses the best path based on accumulated cost, which is based on link speed. The figure shows the AOS-CX default port costs.

Consider the example shown in the figure. Intuitively, you might think that Access-1's best path to the root switch is the direct path. However, this is a 1 Gbps link, which has a cost of 20,000.

The indirect path from Access-1 via Core-2 to the root switch includes two 10 Gbps links, which only have a cost of 2,000 each. $2,000 + 2,000 = 4,000$. Therefore, since the cost is less than 20,000, it becomes the preferred path. (Strictly speaking the root switch will add its cost to itself in this case the cost is equal to 0).

Thus, Access-1 prefers this path to reach the root (since it is the best and least-cost path), and the other redundant path is disabled to avoid loops.


Port roles - Designated and Root

Designated Port

- Closest to the root switch
- Role for all root switch ports
- Port state = Forwarding
- Downlink port, away from root

Root Port

- Another switch is closer to the root
- Best path to the root switch
- Port state = Forwarding
- Uplink port, toward root



MOD 4- 16

Switches that use 802.1w RSTP do not have a complete view of the topology. They build and maintain the loop-free topology by exchanging BPDUs, which indicate how close they are to the root switch. BPDUs help switches to calculate the correct port role for each of their ports.

All ports on the root switch are designated ports (see figure). The root switch is like the boss of the domain. It does not block its ports: only non-designated ports must worry about that.

Looking at the non-root switches, what is the difference between the root port (RP) and designated port (DP)? They both forward frames. The RP is simply the port that the local switch uses to the root switch: the best path toward the root switch. The DP is designated to accept traffic from downstream root ports. Based on this topology, ports marked as DP or RP will forward traffic and other ports (non-designated) are placed in a discarded state.


Port Roles - Alternate and Backup

Alternate Port

- Not closest to root switch
- No loop to itself
- Becomes Root Port if the active fails
- Port state = Discarding

Backup Port

- Not closest to root switch
- Creates a loop to itself
- Becomes the Designated Port if the active fails
- Port state = **Discarding**



MOD 4- 17

To speed up convergence, RSTP introduces two new port states (see the figure).

Alternate ports:

- Port is not the closest to the root on this link.
- This link does not offer the switch its best path to the root.
- Port is not connected to the same switch and has no Loop to itself.
- This port becomes the root port if the active port fails.

As an example, consider the following situation. Access-1 port 1/1/21 can become an alternate port since it fulfills the requisites to it.

Backup ports:

- Port is not the closest to the root in this link.
- This link does not offer the switch its best path to the root.
- Port is connected to the same switch. The switch has a looped connection to itself or it has a loop created by a layer 1 hub.
- The backup port becomes designated port if the existing designated port fails.

A loop could be caused if a layer 1 hub is placed on the topology. As an example, take in consideration the network in the figure. In this situation ports 1/1/26 and 1/1/27 on Access-1 are in a loop situation because a hub was introduced to the topology. If RTSP was running on Access-1, one of these ports is marked as the DP port and the other as the backup port. Likewise, since neither of these ports have a path to the root, they will not become an


alternate port.

Note: Introducing layer 1 hubs into the topology must be avoided.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Operation: Building a Loop-Free Topology

- Elect Root switch
- Select Root Port on non-root switches
- Select Designated Port
- Other ports move to Discard state



This section pulls together all that you have learned and shows you how the entire process works. The Rapid STP algorithm converges on a loop-free topology by performing the following steps:


- Elect the root switch.
- Select the root port on all non-root switches or non-designated bridges.
- Select the designated port on each switch-to-switch links.
- All remaining ports are placed in a discarded state.

RP and DP ports will be placed in a forwarding state while other ports (alternate and backup) move to a discarding state to prevent loops. Let's look at this process in more depth.

Root Switch Election

- Based on BID
- The lowest BID wins
- Priority value helps define Root Bridge
- AOS-CX default priority = 32,768

Device	Priority	MAC
Core-1	4096	00:....04
Core-2	8192	00:....03
Access-1	32768	00:....02
Access-2	32768	00:....01



For both the 802.1D and 802.1w standards, root switch election is based on the BID: the lowest value wins. Recall that the two-part BID is composed of a priority, which defaults to 32,768 for all switches, and a globally unique MAC address.

You can rely on this default, rather random behavior, but it is not recommended. Some small, low-powered switch on the edge of your network might win the election. This makes for a less stable tree, with sub-optimal paths and poor resiliency. Consider the topology shown in the figure. If the priority is set to the default, then Access-2 will become the root switch since it has the lowest MAC address.

You want to ensure that one of the high-powered, more centrally-located core switches wins the election. You have a more robust, resilient, and optimally pathed tree structure. To do this, you simply lower the priority value of the preferred switches, lower than the 32,768 default. In this example, Core-1 becomes the root switch simply by setting its priority to 4096. However, if Core-1 fails, then Access-2 will become the root switch, which would not be optimal. Therefore, it is important that Core-2 becomes the second-best option. In this case Core-2 is setup with a priority of 8192.

Recall that the priority value must be configured in increments of 4096. You might lower the value of Core-1 to 4096, and the value of Core-2 to 8192.


Non-Root Switches: Select Root ports

Root Port

- Another switch on link is closer to root
- Best path to root for non-root switch

Selection criteria

- Lowest path cost to root bridge
- Lowest sender BID
- Lowest port priority
- Lowest port ID



Non-root switches or non-designated bridges must select the best path to the root switch, by selecting its root port: the port connected to this best path. Root port selection criteria are shown in the figure. For example, Core-2 will receive BPDUs on ports 1, 2, 43 and 44. Core-2 must decide which port is the best one:

- Lowest path cost to the root bridge. Assuming that all links are the same speed then, indirect paths on port 1 and 2 are discarded simply because the cost is higher than direct paths. Ports 43 and 44 have the same cost, the next criteria must be used.
- Lowest sender Bridge ID. BPDUs received on ports 43 and 44 are generated by the same device (Core-1), this criteria does not help to break the tie.
- Lowest port priority. The sender (Core-1) will include a port priority on the advertisement. The port priority comes in range between 0-240 and in AOS-CX the default value is set to 128, a bridge will consider a lower number the winner. In this case assuming that default values are set this criteria does not help to break the tie.
- Lowest port ID. The final decision is taken on the port ID, where the lowest the value is the winner. Core-1 lowest port number is 43, so the port on Core-2 that is connected to this port becomes the root port (in this case is Core-2's port 43).

Access-1 will receive BPDUs on ports 21, 22. Access-1 must decide which port is the best one:

Lowest root bridge ID. All devices agree that the root bridge is Core-1; this criteria does not help to select the best path.

Lowest path cost to the root bridge. Assuming that all links are the same speed then, indirect paths on port 22 is discarded simply because the cost is higher than the direct path. Ports 21 becomes the root port.

And Access-2 will receive BPDUs on ports 21, 22. Access-2 must decide which port is the best one: A similar decision will take as Access-1, port 21 the direct path offers a best path so it will become the root port.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY


Select Designated Ports and Alternate Ports

Designated Port

- Closest to root switch
- All root switch ports are designated
- Use same criteria as root switch

Alternate Port

- All ports left that are not RP or DP
- Port must not create a loop



The criteria to select a designated port (DP) is the same as root port.

- Lowest root bridge ID.
- Lowest path cost to the root bridge.
- Lowest sender bridge ID.
- Lowest port priority.
- Lowest port ID.

The following statements are a shortcut to determine quickly a designated port on a link.

All ports in the root switch are always designated ports.

In a switch-to-switch link where previously a root port was elected on one side, the other side must always be a designated port.

See the figure, Core-1 is the root switch and its ports will always be the closest to the root bridge; therefore ports 1, 2, 43 and 44 on Core-1 become designated ports.

On Core-2, let's focus on ports 1, 2 and 44.

Port 1. Core-2 will evaluate if this port is closest to the root, this process is done by comparing the BPDU received from Access-1 on this port and the one that it sends out. This process follow the same rules as the root port, so let's start the analysis:

- Lowest root bridge ID. All devices in the topology agree that Core-1 is the root, this

criteria does not break the tie.

- Lowest path cost to the root bridge. Access-1 and Core-2 advertises the same cost to reach the root switch (both are 1 hop away).
- Lowest sender bridge ID. This criteria breaks the tie since Core-2 has a lower BID than Access-1. Therefore, port 1 in this link becomes the DP.

Port 2. Core-2 will evaluate if this port is closest to the root, this process is done by comparing the BPDU received from Access-2 on this port and the one that it sends out.

- Lowest root bridge ID. All devices in the topology agree that Core-1 is the root, this criteria does not break the tie.
- Lowest path cost to the root bridge. Access-2 and Core-2 advertises the same cost to reach the root switch (both are 1 hop away).
- Lowest sender bridge ID. This criteria breaks the tie since Core-2 has a lower BID than Access-2. Port 2 in this link becomes the DP.

Port 44. This port cannot be designated since the port on the other side of the link is closest to the root.

For Access-1 let's focus on port 1 and 22:

Port1. This port is the only RSTP speaker in the link, this port becomes the DP.

Port22. This port cannot be designated since the port on the other side of the link is closest to the root.

For Access-2 let's focus on port 1 and 22:

Port1. This port is the only RSTP speaker in the link, this port becomes the DP.

Port22. This port cannot be designated since the port on the other side of the link is closest to the root.

In this example, the topology does not include any layer 1 hub which implies that there will not exist a backup port only alternate ports. Simply all ports that were not elected as DP or RP ports will become alternate ports and their state will be marked as discarding. In the figure you can see that Access-1's port 22, Access-2's port 22 and Core-2's port 44 become alternate ports.


Topology Change Mechanism: Begin

- Caused when non-edge ports move to new state
- BPDU TC bit is set
- Switches flush their MAC address table

1 Link Failure

2 Access-1 Port 22 moves to Root port

3 Access-1 sends TC BPDUs out port 22



MOD 4- 22

In RSTP, a topology change occurs when non-edge ports move to a different state, or when BPDUs are no longer received. A switch detects this topology change. The switch actively informs other switches of the topology change (TC). It sets the TC bit in BPDUs and transmits them. It flushes the MAC address table entries associated with all non-edge ports.

Other switches receive a BPDU with the TC bit set from a neighbor. The switch clears the MAC address table entries learned on all its ports, except the ingress port for the TC BPDU. These switches in turn send BPDUs with TC set.


Note. In the original 802.1D standard, any switch could send TC BPDUs to notify others, but the instruction to clear the MAC address table had to always come from the root switch. This two-step process took more time to complete.

To understand how a topology change is handled by Rapid STP consider the following scenario:

1. A link failure occurs between Core-1 and Access-1.
2. As soon as Access-1 detects that its root port is down, it moves its Alternate port 22 to be the new Root Port. Port 22 immediately moves to a forwarding state. The rest of the ports on Access-1 are unaffected.
3. Access-1 generates a TC BPDU and sends it out port 22.

Topology Change Mechanism: End

- 4 Core-2 flushes MAC address table
- 5 Core-2 generates a TC BPDU
- 6 Core-1 receives TC-BPDU, flushes MAC table



MOD 4- 23

4. Core-2 flushes its MAC address tables.
5. Core-2 sends a TC-BPDU.
6. Core-1 receives the TC-BPDU and flushes its MAC address table.

Note: Core-1 also detects the link failure and does the same process as Access-1, generating a TC and advertise it to its neighbor, Core-2.

Note: Edge port link state changes will not trigger a TCN.

Edge Ports

- Not participant in Spanning Tree Algorithm
- Cannot cause a loop
- Fast transit to Forwarding state


```
Switch(config)# interface 1/1/1
Switch(config-if)# spanning-tree port-type admin-edge
```

Admin-edge moves to forwarding, doesn't participate in STP unless BPDU received

Admin network option (default)

- Switch waits for 3 seconds for any BPDU
- If nothing is received, then move to forwarding mode

```
Switch(config)# interface 1/1/1
Switch(config-if)# spanning-tree port-type admin-network
```



MOD 4- 24

Edge ports connect to endpoint, and so they do not expect to receive BPDUs. They need not participate in the STP algorithm, because they cannot cause a loop: they are leaf on the spanning tree, so to speak. Since they cannot cause loops, they can quickly transition to a forwarding state, with no intermediate steps.

If BPDUs are received on an edge port, then the port will act as a normal STP port and participate in the algorithm to prevent layer 2 loops.

There are two options for edge ports: admin-edge and admin-network, admin-network is the default. You can manually configure edge ports, as shown here.


```
Switch(config)# interface 1/1/1
```

```
Switch(config-if)# spanning-tree port-type admin-edge
```

An alternative to the admin-edge option is the AOS-CX admin-network option. With this option, the port looks for BPDUs for the first 3 seconds after the link is up. If no BPDUs are received, the port becomes an edge port and immediately forwards frames. If BPDUs are detected, the port becomes as a non-edge port. It participates in normal STP operation.

Multiple Spanning Tree – 802.1s

- Maintain low CPU while load sharing
- Optimization based on number of possible loop-free topologies



This section introduces you to another part of today's 802.1D standard, multiple STP (MSTP).

You can mitigate Single STP Topology issues using MSTP. MSTP enhances RSTP by building a new topology or instance of the spanning tree protocol. Instances then can carry different VLAN traffic: you get load sharing and optimal paths for different groups (instances) of VLANs.

The figure shows two topologies on the same physical network. In the left-hand example, VLANs 1-20 all share a single STP instance. Core-1 is the root, and the tree grows from there. VLANs 1-20 share these forwarding links, with the unused links available for redundancy.


On the right, VLANs 21-40 share a different STP instance. Core-2 is the root, and the topology emanates from that root. The solid links are shared by VLANs 21-40, and the other links are available for backup. The dormant links for VLANs 1-20 are the active links for VLANs 21-40, and vice versa. Nice!

Common Spanning Tree

- Defines a unique topology tree
- BPDUs are untagged
- VLANs are not considered

Drawbacks

- No load balancing
- Sub-optimal pathing
- Poor link utilization



MOD 4-26


Common spanning tree (CST) defines one instance of STP for the entire network, regardless of the number of VLANs. This is the original scenario that you just explored. All BPDUs use untagged VLAN 1, and so VLANs are not considered.

There is no load balancing, so you get the sub-optimal pathing, poor link utilization, and potentially slower performance. You will now learn about these issues and how to resolve them with MSTP.

MSTP: Regions

Group of switches with common configuration

MST instance name • MST revision number • VLAN mapping



- AOS-CX switches run MSTP by default
- Regions are different by default
- Instance 0 is created by default

```
Region config:
MST config ID      :Switch MAC
MST config revision :0
MST config Digest. :6723...ea99
```

Devices in different regions run RSTP and build CST Topology (valid for all VLANs)

MOD 4- 27

MSTP introduces the concept of a region: a group of switches that have been configured with the same MSTP configuration. The items that must match between the switches in the region include the following:

- Region name
- Region revision number
- VLAN-to-instance mappings

The above information is hashed and is sent in BPDUs. Neighboring MSTP switches create the hash and compare their hash (digest) with the incoming hash to determine if they match: i.e., they are in the same region. which is mapped or associated with a group of VLANs.

Note. A hash algorithm is a mathematical function that is applied to an input (x). A hash function is interesting because from the output (y) you cannot derive or infer what the input was. Thus, it is referred to as a one-way function. Another characteristic of a hash is that if the same input (x) has been entered, the result will be always the same. MSTP takes advantage of this characteristic.

AOS-CX switches run MSTP by default, when STP is enabled. At the factory default settings, each switch has the same revision number (0) and all VLANs are in instance 0. However, the region name (config ID) defaults to the switch's system MAC address. Thus, each AOS-CX

switch operates in its own region. When this happens, the topology behaves like (CST) creating a single topology or instance common for all VLANs. RSTP then exchange messages between the individual devices, as you have learned in previous sections of this module.

Instance 0 or the internal spanning tree (IST) is used as both: a regular instance in MST and the creation of the CST in a multi-region deployment for backward compatibility with RSTP speakers, for this reason instance 0 is known as CIST (common internal spanning tree).

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

MSTP: Region Configuration

Common configuration for all switches


```
(config)# spanning-tree config-name CXF
(config)# spanning-tree config-revision 1
(config)# spanning-tree instance 1 vlan 1-20
(config)# spanning-tree instance 2 vlan 21-40
```

Core-1 configuration

```
(config)# spanning-tree priority 1
(config)# spanning-tree instance 1 priority 1
(config)# spanning-tree instance 2 priority 2
```

Core-2 configuration

```
(config)# spanning-tree priority 2
(config)# spanning-tree instance 1 priority 2
(config)# spanning-tree instance 2 priority 1
```



In the figure, notice that instance 1 and 2 are created. Along with the default instance 0, you have a total of 3 instances. The common configuration creates a single region where you can create multiple instances for optimal pathing and load balancing. Instance 1 is mapped VLANs 1-20 and uses Core-1 as the root switch. Instance 2 is mapped VLANs 21-40 and uses Core-2 as the root switch. Instance 0 has all other VLANs, including the untagged (native VLAN, which is 1 in this example).

In AOS-CX, use the following steps to configure MSTP:

1. Configure the STP mode as MSTP. The command `spanning-tree mode` sets the mode to either MSTP or Rapid Per VLAN Spanning Tree (RPVST).

```
Switch(config)# spanning-tree mode mstp
```

Note. To change the spanning-tree mode, you must first disable spanning-tree with the command `no spanning-tree` in the global context.

2. Set up the MST region name. The `spanning-tree config-name` command sets the configuration name for the MST region in which the switch resides. All switches within an MST region must have identical configuration names, note that the name that you provide is case sensitive.

```
Switch(config)# spanning-tree config-name <Config-name>
```

3. Set up the configuration revision number. Use the revision number to differentiate between region configurations. This is valuable, for example, when changing configuration settings within a region, and you want to track the configuration version you use. It is also valuable when you create a new region from a subset of switches in a current region and you want to maintain the same region name. All switches within an MST region must have identical revision numbers.

```
Switch(config)# spanning-tree config-revision <revision-number>
```

4. Set up the MSTP instance and map the VLAN IDs to it. The spanning-tree instance command creates a new instance with VLAN mapped to an existing instance. Each instance must have at least one VLAN mapped to it. When VLANs are mapped to an instance, they are automatically unmapped from the instance they were mapped before. Any MSTP instance can have all the VLANs configured on the switch.

```
Switch(config)# spanning-tree instance <instance-id> vlan <vlan-ids>
```

5. Optionally, set the switch priority for an instance. The same priority values can be used as in RSTP, the range to set up is 0-15.

```
Switch(config)# spanning-tree instance <instance-id> priority <priority-value>
```

6. Enable spanning tree globally.

```
Switch(config)# spanning-tree
```


BPDUs Protection


Problem

Unauthorized switch connects
Potential loops, root election, security issues

Solution

BPDUs protection

Protected port sees BDPUs
Port is disabled, sends SNMP trap



MOD 4- 29

There are two additional security features that most enterprise switching products support, but that is not part of the STP protocol standard. These two features operate within the AOS-CX switch and don't share information with other switches. How a vendor implements them is vendor-independent. These features are BPDUs Guard and loop protection, though other vendors or products might use a different name for them.

BPDUs protection

BPDUs protection (sometimes referred to as BPDUs guard) is used to detect switch that are not supposed to be part of the network. To solve issue the AOS-CX switches support a feature called BPDUs protection, which is designed to protect the active STP topology by preventing spoofed or incidental BDPUs frames from entering the STP domain. In a typical implementation, BPDUs protection would be applied to edge ports connected to end user devices that do not run STP. If STP BDPUs frames are received on a protected port, this feature disables that port and alerts the network manager using an SNMP trap.

BPDUs protection is disabled by default and is enabled on the edge layer 2 or LAG interfaces:

```
switch(config)# interface <interface-ID>
```

```
switch(config-if)# spanning-tree bpdu-guard
```

```
switch(config-if)# exit
```

```
switch(config)# spanning-tree bpduguard timeout <seconds>
```

By default, the port is disabled when a BPDU is received on the port and an SNMP trap is generated. Optionally you can specify a timeout for the offense, where the port is re-enabled automatically after the timer expires; however, if another BPDU is seen on the port, it would be disabled again. Note that the timeout is defined globally for all ports.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Loop Protection

Problem

Edge ports are unprotected
Could create loops


Solution

Loop protection

Send loop protection packets (default: 5s)

Options when loop detected:

- Only trigger a SNMP notification
- Disable transmission on port
- Disable transmission and reception



MOD 4- 30

Loop protection

So far, the focus of the STP section has been on preventing loops on ports that are intentionally redundant—by design. However, there are situations when a loop can be created on ports that were not intended to be connected. This most commonly occurs on switch ports that face users: edge ports. For example, a user connects two Ethernet jacks on the same switch together.

The loop protection feature protects edge switches from downstream loops by transmitting loop protocol packets out of the ports. The enabled ports automatically send loop protection packets at specified intervals (5 seconds by default). When a loop-protection-enabled port receives a loop protocol packet, the port knows there is a downstream loop, and so takes certain remedial actions.

The actions supported on AOS-CX are:

Do not disable. Only a notification is triggered and reported using SNMP.

Transmission disabled. The port that transmitted the loop detection packet is disabled. This is the default action.

Transmission and reception disabled.

In AOS-CX this feature can be enabled on a layer 2 interface or LAG. The AOS-CX configuration looks like this:

```
switch(config)# interface <interface-ID>
```

```
switch(config-if)# loop-protect
```

```
switch(config-if)# loop-protect action {do-not-disable | tx-disabled | tx-rx-disabled}
```

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

VRRP




The second redundancy solution you'll explore is the open standard Virtual Router Redundancy Protocol (VRRP).

CONFIDENTIAL - FOR TRAINING

Need for Layer 3 Redundancy

Add second DG, however

- Hosts can only have one DG
- You must manually change DG if primary fails
- For DHCP, hosts must disconnect/reconnect




MOD 4- 32

In the previous module you learned about the benefits a default gateway (DG) for endpoints. Recall that an endpoint may only have one DG, and a single DG means a single point of failure. In the figure, if Core-1 fails, PC-1 and all other hosts using it for the DG are now isolated. What can you do?

You could add another DG for redundancy. This seems like an easy solution...but maybe not. Each endpoint's DG is either configured manually or obtained via DHCP. When Core-1 fails, you must either manually reconfigure each host with a new DG, or reconfigure the DHCP scope, and then ask all end users to disconnect from the network and reconnect, or power cycle their PC, or teach them how to trigger a DHCP release and renew action. (In a Windows command prompt for example, use `ipconfig /release`, then `ipconfig /renew`). Whatever method you use, it is not very elegant or scalable, and is disruptive for end users. There must be a better way, right?

First Hop Redundancy Protocols

Objective Device level redundancy for default gateway



MOD 4- 33

The real solution for this challenge can be resolved from the network side, instead from the endpoint side, using some kind of first hop routing protocol (FHRP). This adds resiliency for endpoints by using a coordinated gateway solution: no change to endpoint IP configuration, no DHCP modifications, and no end user disruption. It is automatic!


Note. The term resiliency refers to the ability of the network to adapt to changes and failures.

An FHRP solution creates a single coordinated gateway from two or more physical routers. The two physical routers present themselves to endpoints as a single device, with a single Virtual IP (VIP) address. This VIP acts as the endpoint's DG.

Normally, the primary routing device serves the DG role, forwarding traffic for endpoints. The secondary unit monitors the primary device state. If the primary fails, the secondary device takes over. It takes on the primary role and VIP and forwards endpoint traffic. From the endpoint perspective the Virtual IP address is always available. There is no DG address change, and there is no disruption for end users.

Virtual Router Redundancy Protocol


Objective Standard FHRP protocol that mitigates DG failure outages



RFC 5798 defines the Virtual Router Redundancy Protocol (VRRP): a standard FHRP that enables two or more routing devices to provide gateway redundancy for endpoints.

VRRP uses an active-standby architecture. Only one gateway actively forwards traffic sent to the VIP address. This primary forwarding device is called the active, while the non-forwarding device is the standby.

VRRP Instances



MOD 4- 35

AOS-CX allows you to deploy multiple instances of VRRP, often to balance the load for VLANs, as shown in the figure. Each instance has a unique virtual router ID (VRID) number, as shown in the figure. VRRP instance 1 serves VLAN 10, while instance 2 serves VLAN 20.


Switch Core-1 is the Master or active forwarder for VLAN 10, with Core-2 as the standby. Meanwhile, Core-2 is the Master for VLAN 20, with Core-1 as the standby. This gives you a nice load balancing capability.

Note: On AOS-CX, you can define multiple instances on the same VLAN.

Master Election

Master election based on priority

- 1- 255, highest priority wins
- ArubaOS-CX default priority = **100**
- Same priority - highest IP address is the winner




MOD 4- 36

VRRP members exchange multicast advertisement messages to elect the master DG, using address 224.0.0.18, IP protocol number 112. To control the master election, you set a priority value from 1 - 255, where highest priority wins. If both devices have the same priority, then the gateway with the highest IP address wins the election. In AOS-CX the default priority value is set to 100.

Virtual IP Address

- Endpoints forward their traffic to this IP
- Master in turn handles traffic directed to the VIP
- A MAC address is also created
- 5 first octets are **always the same**
- Last octet match the **GroupID number**

Virtual MAC address:
00:00:5e:00:01:XX



MOD 4- 37


The Virtual IP (VIP) address is the result of the DG coordination. You assign a unique, real IP address to each individual physical gateway, as normal. The VIP address must also be unique. In the figure, you assigned 10.0.10.1 to the Master, 10.0.10.2 to the Backup, and 10.0.10.3 is used as the VIP.

Each VRRP instance has a unique VRRP address. 10.0.10.3 might be the VIP address for Instance 10, used for VLAN 10. meanwhile, 10.0.20.3 might be the VIP address for Instance 20, used for VLAN 20.

You must ensure that this VIP address is configured as the endpoint DG. Thus, endpoints forward their traffic to the VIP. The VRRP master receives these packets and routes them. If the master fails, the standby unit takes over. The devices do not learn about the router's physical IP addresses, 10.0.10.1 and 10.0.10.2, in the example.

A virtual MAC address (vMAC) is automatically assigned to the VIP. As defined in the standard, this address shall be 00:00:5e:00:01:XX, where XX = the VRID. In the figure, the VRID = 10, and so the vMAC is 00:00:5e:00:01:0A (Hex A = 10 in decimal). Thus, when endpoints ARP for their DG address of 10.0.10.3, they will learn the vMAC address, and add it to their ARP table. This simplifies operations, since in case there is a switchover from master to standby, the edge devices can still use the same IP and MAC address for the default gateway.

VRRP Failover Operation



MOD 4- 38


The image shows a scenario where two Core switches run VRRP. Core-1 has a higher priority, and so is the VRRP master. Core-2 is the VRRP standby, constantly monitoring Core-1 status via a keepalive mechanism. Endpoints forward their traffic to 10.0.10.3, which is services by the VRRP master.

Then the master fails. The standby stops receiving keepalive messages, and so knows that the master is down. The former standby takes over as the new VRRP master and begins to forward traffic for VIP 10.0.10.3.


VRRP Preemption

Objective Determine role of Original Master when it is available again

Preemption Enabled



Preemption Disabled



Preemption is enabled by **default**

MOD 4- 39

If you continue to review the previous figure and the failover operation, you can note that that Core-1 failed, and so Core-2 took over as the new master. What happens when Core-1 comes back online? This depends on how you configure VRRP preemption.

If preemption is enabled, then Core-1 will reassume its original master role and Core-2 reverts back to its standby role. This is the AOS-CX default setting. It is nice because you know that under normal operating conditions, when all devices are up, the same router always acts as the master. This can be especially important if you are using multiple VRRP instances. If a master fails, the remaining device(s) will carry the load for all endpoints, while the former master, once again operational, remains unused.

If preemption is disabled, then Core-1 will not resume its original master role. Core-2 remains the master, and Core—1 takes on the standby role. You must manually disable preemption with the command **no preempt** if you want to use this mode. With preemption disabled you lose the benefits described above. Some administrators might choose to disable preemption if they are worried about the (very brief) time lag that might occur during the preemption process: while the routers switch back over to new operational states, potentially during the middle of a busy day. Due to the high-performance nature of AOS-CX devices, this is rarely a concern.

VRRP Configuration

Master configuration

```
router vrrp enable
```

Enable VRRP globally

```
interface vlan 11
```

```
    vrrp 11 address-family ipv4
```

Create VRRP instance

```
    address 10.1.11.1 primary
```

Set VIP

```
    priority 110
```

Adjust the priority

```
    no shutdown
```

Backup configuration

```
router vrrp enable
```

```
interface vlan 11
```

```
    vrrp 11 address-family ipv4
```

```
    address 10.1.11.1 primary
```

```
    no shutdown
```

No need to set priority on Backup,
default is 100



MOD 4- 40

You can configure VRRP using commands presented in the figure.





Introduction to VSX

The last part of this module explores Aruba's switch virtualization technologies: Virtual Switching eXtension (VSX) and Virtual Switch Framework (VSF). This section provides a very brief introduction to VSX. VSX is covered in depth in the Aruba professional-level course.

CONFIDENTIAL & PROPRIETARY

VRRP vs Switch Virtualization



MOD 4- 42

Although VRRP and STP are supported protocols, Aruba has developed a better solution. Look at the figure. The left example implements VRRP and STP which is active/standby solution. This means that only one of the switches is actively processing the traffic, the solution is sub-optimal. Of course, you can enable multiple MSTP instances and configure multiple VRRP groups so one switch is a VIP master and the other is a VIP backup for a given VLAN. This however allows load-sharing only if the traffic originates from different VLANs – per VLAN load-sharing.

What if you want to use both switches and utilize all of the uplinks even if the traffic comes from the clients connected to the same VLAN? The solution to the problem is virtualization that allows truly active/active traffic processing. You can combine two or more switches (depending on the technology) to create one virtual switch. You can now use link aggregation instead of STP, which will use all the uplinks and allows load-sharing even if the traffic originates from the same VLAN. This is because when using LACP, the multiple switches that comprise the virtual switch send the same system identifier, and thus the edge devices think they are connected to a single switch. For IP routing both switches share the same IP address, and both are capable of layer 3 forwarding, you get a truly active/active L2 and L3 traffic processing.

AOS-CX Switch Virtualization: VSX vs VSF

Agg/Core Layer

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

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

Ease of deployment is important – large number of access switches



MOD 4- 43


Aruba's VSX and VSF are similar to each other, and even to earlier virtualization techniques like VRRP, where multiple devices are perceived as a single device. This improves resiliency while reduces configuration and management overhead. But there are distinct differences between VSX and VSF, as summarized in the figure.

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

Configuring VSF and backplane stacking is a plug-and-play affair - a switch, at default settings, listens to discover another VSF switch on some designated ports. This ease of deployment is especially important at the access layer : you may have dozens, or hundreds of access switches. VSF accommodates up to 10 nodes, while a VSX configuration has exactly two switches: primary and secondary.

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 like MAC and ARP addresses, and more.

VSX Benefits – MC-LAG



Control Plane

- Dual control plane for better resiliency
- Unified management (synchronized configuration and easy troubleshooting)
- Independently software upgradable with 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 4- 44

Here is an example of an Aruba-recommended topology in which you deploy VSX at the aggregation layer.

While you can deploy VSX at the core if the topology requires, 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. At the same time, VSX permits unified management on the features that administrators 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 (MCLAG), distributed, or VSX LAG. These VSX LAGs operate at layer 2 because the access layer switches do not route, and the aggregation layer VSX pair acts as the default gateway for the access VLANs. The layer 2 LAGs eliminate the need for Multiple Spanning Tree Protocol (MSTP) or Rapid Per-VLAN Spanning Tree (RPVST) 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.



Knowledge Check

Self-check on key learning points



Please check your understanding with a brief knowledge check.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Question #1

Which of the statements below accurately describe spanning-tree protocols? (Select three.)

- a. They mitigate Layer 2 loops, broadcast storms, and multiple frame copies.
- b. Switches elect a root bridge, from which the spanning-tree emanates.
- c. All switch ports forward traffic, to improve resiliency.
- d. When a link failure occurs, switches converge on a new topology.

a


Knowledge Check

MOD 4- 46

A,B,D

Question #2

Match the descriptions to the correct STP version.



Knowledge Check ✓




MOD 4- 47

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Question #3

See the figure. Assuming default settings and normal operation, which interface will switch Access-1 use as its root port, and why does it choose that interface?

- a. Port 1/1/1, because it is directly connected to the root
- b. Port 1/1/1, because it has a lower port number
- c. Port 1/1/2, because it has a higher port number
- d. Port 1/1/2, because it has a lower cost



a

Knowledge Check

MOD 4- 48

D

Question #4

Your colleague is confused about why a certain switch in your network was elected as the root switch. Which of the statements below would help your colleague understand this aspect of spanning-tree protocols? (Select two.)

- a. Root switch election is based on the lowest Bridge ID (BID) and the highest priority
- b. Root switch election is based on the lowest Bridge ID (BID), which you can override by configuring a priority value on your preferred root switch.
- c. The Bridge ID (BID) is derived from an IP address on the device.
- d. It is a best practice to manually configure a priority value on both your preferred root switch and backup root switch.

a

Knowledge Check

MOD 4- 49

B,D

Question #5

Which statements below are true about the Virtual Router Redundancy Protocol (VRRP)? (Select two.)

- a. It provides better first-hop redundancy as compared to configuring hosts with two Default Gateways (DG).
- b. With VRRP, two routing devices actively forward traffic.
- c. You can define VRRP instances to improve load sharing for VRRP deployments.

a

Knowledge Check

MOD 4- 50

A,C

Question #6

Which statements below provide accurate comparisons of Aruba's Virtual Switching eXtension (VSX) versus Aruba's Virtual Switching Framework (VSF)? (Select three.)

- a. Both solutions allow multiple devices to be perceived as a single device.
- b. You typically run VSX on aggregation and core switches.
- c. VSF can be easier to configure.
- d. Like VSF, VSX maintains separate control and management planes.

a

Knowledge Check

MOD 4- 51

A,B,C

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY



Lab Activity

Consider the lab activity.

CONFIDENTIAL - FOR TRAINING

Lab Tasks

4.01
Spanning-Tree

4.02
VRRP

4.03
Introduction to VSX

Module 5 - VSF



MOD 4- 53

The figure provides a brief review of labs for this module. Please see your lab guide for details. When complete, you will continue with Module 5: VSF.


Rev.22.41

aruba
a Hewlett Packard
Enterprise company

VSF

Aruba Campus Access Fundamentals

Hello and welcome back. This is Module 5 is all about Virtual Switch Framework (VSF).



Objectives

- Describe VSF
- Configure VSF
- Describe Auto-VSF
- Explain VSF MAD

One great way to optimize network resiliency and ease network management tasks is by using the Virtual Switch Framework (VSF) feature on access layer ArubaOS-CX switches.

After completing this module, you will be able to describe VSF technology, operation, and use cases. You will be able to configure VSF, and describe aspects like auto-VSF and VSF Multi-Active Detection (MAD).

VSF Introduction

We begin with an introduction to VSF.

CONFIDENTIAL - FOR TRAINING

Introduction to Stacking Technologies


Objective A group of switches managed as a single virtual device

Features

- Centralized control and management plane
- Independent data planes

Benefits

- Ease of management
- Network simplification



MOD 5- 4

Stacking technologies allows you to manage a group of switches as a single device , commonly referred to as a virtual switch. Control and management plane functions are centralized in one group member, but each member in the stack runs its own independent data plane. The tight relationship between the control and data planes is maintained; it just happens on an inter-switch basis.


Stacking benefits include:

- **Ease of management:** You no longer need to connect to, configure, and manage each individual switch. You simply configure the primary switch. That configuration is then automatically distributed to other virtual switch members. This simplifies network setup, operation, and maintenance.
- **Network simplification:** Since multiple devices share a common control plane, routing protocols and STP are no longer needed inside the stacking group. Connected devices perceive the group as a single device.

Distributed Data Plane and Distributed Link Aggregation


LAG to a switch

- If SW1 fails, AP1 is lost
- Switches must run STP



LAG to a stack

- If SW1 fails, AP1 is OK
- No need for STP



Stacking technologies maintain the data plane distributed across all the members: this implies that each physical device individually creates and populates forwarding tables like ARP and MAC. These tables are then shared across all members using the control plane.

Mobility gateways (formerly referred to as controllers), firewalls, APs, and other devices can benefit from stacking with LAG enable switches, since they perceive a LAG connection to a single device. However physical links terminate in different stack members (referred to as distributed link aggregation). If SW1 fails in the example shown in the figure, the traffic can still use other LAG links.

LAG and Stacking features enable the network to fully use all available links at the same time. Notice that Spanning Tree is not needed because the stack operates from a single control plane. Aruba highly recommends this implementation since it provides resiliency and is easier to manage one virtual switch than multiple physical switches.

Aruba VSF Stacking Solution and Platforms

VSF on ArubaOS-CX switches

- 6200 and 6300 series only
- Enabled by default
- Switch interconnect is Ethernet
- Same product series required




MOD 5- 6

Aruba Virtual Switching Framework (VSF) defines a virtual switch, which is composed by individual physical switches interconnected using Ethernet links. All member switches share a single control plane. Devices connected to the VSF perceive a single device. As you recently learned, this virtual switch behavior has significant benefits related to simplified management and improved connectivity.

In ArubaOS-CX, VSF is supported on the 6200 , 6300M, and 6300F access switches. You can configure a maximum of 10 members in the stack. This feature is enabled by default.

The VSF feature is also available in the Aruba AOS switching family, including the 5400 and 2930 series, but the feature is disabled by default on those switches. You cannot form a VSF stack with switches from different model families. This means that VSF can only be formed using any combination of 6200 or 6300 ArubaOS-CX switching series; i.e., you cannot have both 6200 and 6300 switches in the same stack.

VSF Member Roles and Links



MOD 5- 7

Aruba VSF creates a single control plane that runs on a single VSF member. This is the primary member and always uses member ID 1. This device assumes the conductor role in the stack.

Note: A factory-default Aruba 6200 or 6300 switch boots up as VSF enabled switch with member ID 1. This implies that the switch behaves as the primary member.

A secondary member provides for high availability in case primary failure. You can choose any member in the stack to take the secondary role, but it must be explicitly defined. Just configure any member ID except 1, which is reserved for the primary switch.

The other devices in the VSF are members, which only perform data plane functions and cannot assume the primary/conductor role.

VSF switches are interconnected using uplink ports.

VSF Open Virtual Switch Database

Objective Track/sync state and configuration data for the stack



You previously learned about the Current State Database (CSDB) when discussing checkpoints in module 2. As the most important element of the ArubaOS-CX architecture, it is contacted by all control and management protocols. In a VSF stack this database runs on the primary/conductor switch.

When VSF runs on a group of switches the Open Virtual Switch Database (OVSDB) is also created. This new database runs on the conductor switch and contains state and configuration data for the VSF stack itself. The conductor switch synchronizes OVSDB content with the secondary switch to ensure that it can quickly take over the active role without interruption.

The OVSDB database includes six tables:

VSF member table

VSF link table

System table: Includes the number of members, MAD status, fragment status, and topology type.


Subsystem table: Boot time for each member

Interface table

Topology table

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

VSF Topologies



MOD 5-9


Of course, VSF members must be physically connected to form the switching stack, utilizing one of two topologies: Daisy chain or Ring.

- Daisy chain. As the name implies, this topology interconnects VSF members with a single chain of Ethernet connections. From the figure you can easily see that a switch or link failure causes the stack to be split. This means that part of the stack is unable to provide endpoint connectivity.
- Ring topology. This topology is recommended since it offers a backup path in case of a switch or link failure.

Note: Full mesh topology is not supported with VSF.

VSF Requirements

- ✓ All members must run the same AOS-CX version
- ✓ All switches must be the same series 6200 or 6300. Can combine models i.e. 6200-25G and 6200-48G; 6300M and 6300F
- ✓ VSF Link is using a regular 10, 25, or 50 Gbps port
- ✓ Maximum 8(6200 series) or 10 (6300 series) members are supported



MOD 5- 10

VSF has some requirements that you must meet:

- All members must run the same AOS-CX version of software.
- All switches must be the same series 6200 or 6300. You can combine models in a series, for example, 6200-25G and 6200-48G.
- VSF Link uses regular Ethernet ports.
- It is recommended to use the uplink ports which support 10, 25 or 50 Gbps speeds.
- A maximum of 8 members supported on 6200 stack and 10 members on 6300 stack.

Note: Firmware is automatically synchronized within the VSF stack by the conductor. Joining switches will automatically upgrade or downgrade to switch's firmware running on the conductor switch.

VSF Configuration



Now that you understand some of the basics of VSF let's talk about how to set it up.

CONFIDENTIAL - FOR TRAINING

VSF member ID and Port numbers

Access-6300# **show interface brief**

Port	Native VLAN	Mode	Type	Enabled	Status	Speed (Mb/s)
1/1/1	1	access	SFP+DAC1	yes	up	10000
1/1/2	1	access	SFP+DAC1	yes	up	10000
...						
2/1/1	1	access	SFP+DAC1	yes	up	10000
2/1/2	1	access	SFP+DAC1	yes	up	10000
...						
3/1/1	1	access	SFP+DAC1	yes	up	10000
3/1/2	1	access	SFP+DAC1	yes	up	10000



MOD 5-12

When the VSF stack forms, all physical devices use a single control plane. This means that all switch interfaces are available for configuration, using the standard ArubaOS-CX Member/Slot/Port notation, as shown in the figure. Use the command `show interface brief` to see all available ports for all switches within the stack.


Note: Since all 6200 and 6300 switch series models are fixed-port switches, the slot number is always 1.

VSF configuration

Consider the scenario where three 6300 switches are assigned to the same VSF stack, shown in the figure. You see the port configuration for all three switches. The link command specifies which physical port or ports are used to connect the switches. If you use a ring topology, then you would have two link interfaces 1 and 2. After Access-2 reboots, you add the `vsf secondary-member 2` command, which will cause it to renumber its member ID from 1 to 2. The switch will have to reboot a second time. The same would be required of Access-3, where its member number would be changed to 3. Note that since there is already a primary and secondary member, the third switch will be considered a normal member of the stack.

Once you have set up VSF, use the `show vsf` command to determine member status and role. As you can see setting up VSF is a very simple process; you'll be doing this in the lab later in the module.

VSF Configuration Example



```
Access-1(config)# vsf member 1
Access-1(config-vsf)# link 1 1/1/27
Access-1(config)# vsf secondary-member 2
```

```
Access-2(config)# vsf member 1
Access-2(config-vsf)# link 1 1/1/27
Access-2(config)# vsf renumber-to 2
Switch reboots
```

```
Access-1# show vsf
```

Id	Mac Address	type	Status
1	88:3a:30:97:c7:00	JL668A	Master
2	88:3a:30:97:1a:80	JL668A	Standby

MOD 5-13

Consider the scenario where a couple of 6300 switches are assigned to the same VSF stack, as shown in the figure. You see the port configuration for switch Access-1, and for Access-2 in the figure. After Access-2 reboots, you add the vsf secondary-member 2, after which, again, the switch reboots.

Use show vsf to determine member status and role.

VSF Pre-Provisioning

- Prepare VSF Link and member configuration
- Based on switch model
- VSF-connected switch boots with proper config

Example: VSF Provisioning for a 6300F 24-port switch

```
Switch-6300(config-vsff)# vsf member 4
Switch-6300(config-vsff)# type jl668a
Switch-6300(config-vsff)# link 1 1/1/25
Switch-6300(config-vsff)# link 2 1/1/26
```

P/N	Model
JL658A	6300M 24SFP+/4SFP56
JL659A	6300M 48SR PoE CLS 6/4SFP56
JL660A	6300M 24SR PoE CLS 6/4SFP56
JL661A	6300M 48G PoE CLS 4/4SFP56
JL662A	6300M 24G PoE CLS 4/4SFP56
JL663A	6300M 48G/4SFP56
JL664A	6300M 24G/4SFP56
JL665A	6300F 48G PoE CLS 4/4SFP56
JL666A	6300F 24G PoE CLS 4/4SFP56
JL667A	6300F 48G/4SFP56
JL668A	6300F 24G/4SFP56

Check Aruba Networks website for complete list of products and P/N

MOD 5- 14

ArubaOS-CX supports member pre-provisioning. This enables you to prepare the VSF Link and member configuration for a specific 6200 or 6300 switch model before the switch is connected to the stack. This configuration is performed from the primary/conductor switch. Then, when the new switch boots up, it can automatically join the stack and obtain its proper configuration.

The figure shows a VSF pre-provisioning example for a 6300F 24-port switch done from a conductor switch. The table in figure indicates some available options for the switch model number. The type command, followed by a space and a "?", will show the available model options. The list will change when new products or models are released. A firmware update may be required to support these new models.

Auto VSF




As you can see, setting up VSF manually is a fairly simple process. Let's see if we can make VSF deployments even easier with the auto-VSF feature.

CONFIDENTIAL - INTERNAL

Auto VSF

- Mechanism to automatically form a stack
- Reduces the number of user intervention touch points
- Must be running AOS-CX 10.07 or higher



MOD 5- 16

The VSF auto-stacking feature provides a mechanism to automatically form a stack when the stack members are physically connected in a desired topology. This reduces the number of user intervention touch points to form a VSF stack.

A manual stack formation procedure generally requires you to explicitly log in to each of the switch, configure the links, renumber the member ID, and then make the physical connection to form a VSF stack of desired size and topology. This is error prone since there are multiple touch-points involved in the whole work-flow for each member.

The auto-stacking feature eases this problem by reducing the number of touch-points involved to simple physical connections of the links. A new factory default switch can be added into an existing stack by physically connecting it to a given member of the stack. The new switch is automatically assigned a member ID, then it reboots. After reboot, the newly added member joins the stack.

There are two major components to the auto-stacking solution:

Peer discovery: Initiated from the conductor using one of the methods of designating a conductor switch.

Auto-join eligibility: Determined by the configuration state of each stack member. A switch with a factory default configuration is always auto-join eligible.


Peer Discovery & Auto-join Eligibility

Peer discovery

- Initiated from the conductor
- Uni-directional process
- Higher VSF port sends discovery packet
- Peer replies with MAC, part number, auto-join eligibility

Auto-join eligibility

- Factory default switch is auto-join eligible
- If connected to the existing stack it reboots and joins the stack
- Non-factory default switch can be forced to join the stack using `vsf force-auto-join`



Force auto-join will work only if the switch does not have any pre-existing VSF configurations

MOD 5- 17

Peer discovery

Auto-stacking peer discovery is a unidirectional process. It starts on the conductor of the existing VSF with the VSF link containing the higher denomination VSF port. The conductor sends a VSF peer discovery protocol packet on that port. The peer receives the packet, determines if it is valid, and sends a response with information including its auto-join eligibility, MAC address, and part (model) type. If the peer is auto-join eligible, the VSF member and link configurations are automatically added to the running-configuration of the active conductor.

Auto-join eligibility

A switch in its factory defaults configuration state is considered to be auto-join eligible. If the auto-join eligible switch is connected to existing stack, it will automatically reboot and join the stack. Once it moves out of factory default configuration state, it is not considered as auto-join eligible and cannot automatically join an existing stack. However, an administrator can still manually configure the links, renumber the device to make it part of a stack.

Force auto-join support

Only a switch with a factory default configuration is considered to be auto-join eligible. In order to support factory express deployments where the administrator wants to add a switch which is in its non-factory default configuration, the force auto-join configuration support is

provided. Use the `vsf force-auto-join` command to force the switch to join the stack automatically. Once the user sets force auto-join in the switch configuration, the switch will be considered as auto-join eligible and will join the stack even though the switch does not have the factory default configuration.

Note: Force auto-join will work only if the switch does not have any pre-existing VSF configurations. You would need to either remove the VSF configuration or reset the switch back to its factory default state if you wanted it to join a stack.


CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Forming a VSF Stack Using auto-stacking Command

1. Install the switches and physically connect them on the reserved auto-stacking interfaces.
2. On the conductor switch use the `vsf start-auto-stacking` command. The links and secondary member will be automatically configured.

```
switch(config)# vsf start-auto-stacking
```

3. Since the switches are already physically connected,, each switch in the stack reboots automatically and join the stack one after another automatically.



MOD 5- 18

Remember that you need to meet these requirements for auto-stacking to work:

- All switches must be in the factory default configuration.
- All stack members must be connected in a desired topology on the reserved VSF link ports.

The figure illustrates the steps to use auto-stacking. It's fairly simple since the only configuration command needed is done on the conductor: `vsf start-auto-stacking`.

Reserved Interfaces for Auto-Stacking

Platform	SKU part number	Ports reserved
6300	JL658A	25 and 26
	JL660A	
	JL662A	
	JL664A	
	JL666A	
	JL668A	
6300	JL659A	49 and 50
	JL661A	
	JL663A	
	JL665A	
	JL667A	
	JL762A	
6200	JL724A	25 and 26
	JL725A	
6200	JL726A	49 and 50
	JL727A	
	JL728A	



MOD 5- 19

Based on the product type of a switch, the following two interfaces are reserved for the auto-stacking process:

- 24-port switch models: 25 and 26
- 48-port switch models: 49 and 50

Users can physically connect the switch to an existing stack on one of these reserved auto-stacking interfaces.

For auto-stacking, it is required that the existing stack conductor's VSF link port have a higher value. For example, if you connect member 2 to member 1 (existing VSF stack conductor) , you must connect interface 26 of member 1 to interface 25 of member 2. Otherwise, the member 2 will not join in auto-stacking.

VSF MAD



Don't get angry about potential system failures! You've got VSF MAD.

CONFIDENTIAL - FOR TRAINING


VSF Link Failure

Split brain condition

- Primary and Secondary members are isolated, but still up
- Both fragments think they are in control
- External devices exchange control messages with both fragments

NOT
GOOD!

Unexpected
network
behavior



Solution:

Disable all ports on SW2 - the fragment with the original Secondary



MOD 5- 21

VSF Link failures could cause a fragmented stack. The figure shows a fragmented VSF stack, due to link failures. Originally, SW1 was primary and SW2 was secondary. After fragmentation, there is no direct connection between the two switches. SW1 continues to think that it is primary. SW2 no longer hears from SW1, and so now thinks that it is the primary. This is known as split-brain condition: both fragments continue to function although there is no communication between them.

A split-brain can cause unstable, unexpected network behavior. A packet received in one fragment and destined for the other fragment is discarded. Even worse, both segments use the same IP address and the same routing information: external devices could start populating duplicated data, which can cause very strange network behavior.


Note. Split-brain situation could also occur if a VSF member fails in a daisy chain topology. Aruba always recommends deploying a stack using a ring topology.

The best way to solve a split-brain situation is to disable the ports on one of the segments: the segment that contained the original secondary switch.

Multi-Active Detection – MAD

Split detection mitigates split-brain

- Standby member verifies if Primary is up
- All ports of all members in non-primary fragment are shutdown



MOD 5- 22

You can use Multi-Active Detection (MAD) to avoid split-brain situations. Thus, if a VSF link failure occurs, the segmented switches that do not include the primary switch verify primary member status. If the original primary is up, then all members in the fragment that does not include the primary member will disable all their ports, with the exception of the VSL link ports.

VSF uses the following mechanism to detect and verify the status of the primary member:

- Management interface split detection. This method requires you to connect out-of-band Management (OOBM) interfaces to primary and secondary stack members. These interfaces must be in the same layer 2 broadcast domain (VLAN); thus, if you have more than two switches in a stack, you'll need to connect them to an intermediate switch that is not participating in the stack. This network is used to identify active stack fragments. Each member broadcast Split Detection Protocol (SDP) packets to identify stack fragments that are currently operational.

MAD is disabled by default. To enable MAD on the OOBM port, use the `vsf split-detect mgmt` command.



Knowledge Check

Self-check on key learning points



Its time for a knowledge check.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Question #1

Which of the following statements accurately describe Aruba's Virtual Switch Framework (VSF)? (Select four.)

- a. VSF defines a virtual switch, composed of individual switches interconnected via a proprietary backplane link.
- b. You can configure a maximum of 10 members in a VSF stack if using 6300
- c. VSF Primary and Secondary roles can fill the Conductor role, but a VSF Member cannot.
- d. VSF Conductor and Standby maintain synchronized state and configuration data by leveraging an Open Virtual Switch Database (OVSDB).
- e. VSF supports daisy chain, ring, and full-mesh topologies.
- f. You can leverage VSF pre-provisioning to facilitate adding new members to a VSF stack.

a

Knowledge Check

MOD 5- 24

B,C,D,F

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Question #2

Which VSF feature provides a mechanism to automatically form a stack?

- a. VSF pre-provisioning
- b. VSF priority configuration
- c. VSF LAG
- d. Auto VSF

a

Knowledge Check

MOD 5- 25

D

Question #3

VSF Multi-Active Detection (MAD) mitigates a split-brain condition by shutting down all ports of all members in a non-primary fragment.

- True
- False

a

Knowledge Check

MOD 5- 26

T



Lab Activity

Consider the lab activity.

CONFIDENTIAL - FOR TRAINING

Lab Tasks

Lab 05.01: VSF

Task 5.1

- Prepare the topology

Task 5.2

- Create a VSF Stack using Automatic Join of New Member

Module 6 – Intro to Aruba Solutions



MOD 5- 28

The figure provides a brief review of lab tasks. Please see your lab guide for details. When complete, you will continue with Module 6 – Introduction to Aruba Solutions.

Rev. 22.41




Introduction to Aruba Solutions

Aruba Campus Access Fundamentals

Hello everyone! This is Module 6, Introduction to Aruba Solutions.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY



Objectives

- Explain ESP
- Describe Aruba switching products
- Describe the Aruba WLAN portfolio
- Explain Aruba Central and Aruba ClearPass

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


- Explain Aruba Edge Services Platform (ESP) - ESP is cloud-based architecture for networking.
- Describe various Aruba switching products.
- Describe various Aruba Wireless LAN (WLAN) products.
- Understand the Aruba Central management platform.
- Describe the Aruba ClearPass product- ClearPass is an authentication server platform for employees and guests.

Aruba ESP

Let's start by introducing the Aruba Edge Services Platform (ESP).

CONFIDENTIAL - FOR TRAINING

Aruba ESP Architecture



MOD 6- 4


The Aruba Edge Services Platform (ESP) is based on four core attributes:

- Unified Infrastructure - means you can manage different types of devices in different locations using a single interface. It eliminates silos and streamlines operations.
- Zero Trust Networking – this means that every single device that plugs into the network gets authenticated and authorized. You are not inherently trusted and zero trust security includes role-based security to minimize the need for manual configuration.
- Artificial intelligence operations - gives the network a 6th sense, resolving issues before they impact the business or at least resolve them faster
- Aruba Central - powers ESP, which is a software as a service(SaaS) Management solution, ready to monitor and manage any size company. Aruba Central's flexibility gives you greater choice and control over how you consume and pay for technology to accomplish their goals.

Reference Architecture: Small Campus

Design parameters

- CPPM for User Roles
- VLANs on access switches
- VLANs terminate on core switches
- SSIDs mapped to VLANs
- Core switches enforce VLAN policy for wired and wireless



The figure shows an example network of many of the different components that you will typically see in a small campus network.

Aruba Switch Portfolio



Let's move on to Aruba CX switching portfolio.

CONFIDENTIAL - FOR TRN

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 switch software licensing or subscription costs

MOD 6-7


The Aruba CX switching advantages:

1. A single operating system across your entire environment spanning from edge to data center.
2. Aruba Central is a management tool. It is a unified infrastructure, single pane-of-glass to manage and monitor your Aruba wired and wireless devices. Central supports a multi-editor mode for advanced configurations, which allows you to orchestrate common configuration changes across one or multiple CX switches.
3. Aruba Network Analytics Engine (NAE) is a tool to help you with monitoring and 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.
4. 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.
5. Dynamic segmentation is a way to unify network access and network security regardless of whether the user is using wired or WiFi connection.
6. The top-shelf, Gen7 ASIC is a microprocessor specialized in processing network traffic. It

offers programmable, data center class performance with no bottlenecks and with no switch software licensing or subscriptions.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Aruba CX Product Portfolio



MOD 6- 8

The Aruba CX portfolio includes modular access, aggregation, and core switches. At the access and aggregation layers you get power over Ethernet (PoE), multi-gigabit speeds, secure segmentation, and stacking capabilities. At the campus core and data center, you get high availability, bandwidth, and speed: with 40 and 100 GbE scalability.

Here is an overview and product positioning for the Aruba CX switching product portfolio.

Aruba CX 6000 and 6100

- Overview: Ideal for smaller networks of around 100 – 500 devices with 10 GbE uplinks.
- Positioning: branch offices, mid-market enterprise, and SMBs looking for a reliable, cost-sensitive entry level access.

Aruba CX 6200

- Overview: These fixed-port access switches support 8-member stacking with built-in 1/10GbE 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 WIFI 6, plans for Internet of Things (IoT), and looking for

automation to help with rollouts. Supports bandwidth intensive applications, plus investment protection with future-ready smart rate multi-gigabit 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-gigabit and high power 60W PoE port capability.
- 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 1/10/25GbE, 1/10/40 GbE, and 40/100GbE connectivity with 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. The Aruba Fabric Composer (ACF) product 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 25.6Tbps, 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, aggregation, ToR, 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 10/40/100GbE 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 and buffers) is required.

Aruba CX 10000 Series Switch

Overview: Provides advanced services to the data center such as distributed stateful firewall, encryption services, and stateful network address translation

(NAT), zero-trust networking (ZTN).


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

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


CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Aruba Central for AOS-CX Management

Fast, error-free management via Aruba Central



The image shows the Aruba Central web interface in 'Simple' mode. A callout box labeled 'MultiEdit' with a toggle switch pointing to the 'Off' position is shown above the interface. The interface displays a configuration page for a switch, with a left-hand navigation menu and a main content area showing various configuration sections like 'Properties', 'Interfaces', and 'Routing'.



The image shows the Aruba Central web interface in 'MultiEdit' mode. A callout box labeled 'MultiEdit' with a toggle switch pointing to the 'On' position is shown above the interface. The interface displays a configuration page for a switch, with a left-hand navigation menu and a main content area showing a diff view of configuration changes. The diff view shows a table with columns for 'Before', 'After', and 'Diff', highlighting changes in configuration commands.

Simple, yet powerful UI-based config.
Fast, routine changes – no advanced skills required


Multi-Edit for advanced config options.
Diff view: Ensure changes match intent

Toggle between modes with one click. Built-in validation ensures changes are conformant.

MOD 6- 9

Aruba Central features an even simpler, more powerful UI, to ease routine changes to interfaces, ports, VLANs, and more on the Aruba OS-CX switches. You can easily toggle to MultiEdit mode for advanced config settings, with assistances like command completion, syntax highlighting, and built-in validation of network policies. You have improved flexibility over how to implement configuration changes, where you can flip back and forth using the UI or MultiEdit mode depending on your needs and preferences.

Network Automation with AOS-CX



MOD 6- 10

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 and troubleshooting.

NAE uses python-based agents, developed by Aruba and the Aruba community, to automate network monitoring and troubleshooting operations. Agents monitor the state, settings, and statistics for system components in the current state database (CSDB), for example, 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 switch's web UI's analytics dashboard. You can also automate the NAE through the REST API using tools such as Postman, Ansible, or another 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 CX switches using the Central configuration UI. Custom NAE scripts can be pushed via template configurations as well.

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 6- 12

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

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

Dynamic Segmentation: Operation Overview



MOD 6-13

With dynamic segmentation, the switch authenticates the user by interacting with an authentication, authorization, and accounting (AAA) server, like ClearPass. ClearPass profiles each device, deems appropriate endpoints as acceptable, and sends accept messages to the switch. The messages include tunnel information, the assigned VLAN, and the secondary role. The switch uses this information to establish the tunnel to the gateway (mobility controller or MC for short). The user's traffic is tunneled to the gateway where the gateway applies the appropriate policies to it, like firewall policies. In the figure, a user was profiled with the Finance role, and so their port is configured for VLAN 15, and their traffic is securely tunneled to the gateway. The gateway applies centralized policy enforcement to limit finance traffic as appropriate. Printers and cameras are similarly treated, as shown in the figure.

ClearPass can also indicate if device traffic should be locally forwarded by the switch, directly toward its ultimate destination, without tunneling to the controller. You sacrifice some of that centralized PEF enforcement, but you accommodate devices that are uber-sensitive to delay and jitter.



Aruba WLAN Portfolio

MOD 1- 14

Let's move on to the Aruba wireless LAN (WLAN) portfolio.


CONFIDENTIAL - FOR TRAINING

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 6-15

Aruba WLAN provides 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, Aruba delivers critical security and reliability to mobile users, IoT, and latency-sensitive applications, even in Wi-Fi crowded areas. Enterprises with 5G strategies also benefit from Aruba Wi-Fi 6, serving as an onramp to 5G (in place of small cell/DAS). The different AP models represent different levels of performance and capacity as well as different operational environment: indoor vs outdoor.

Aruba AP 500 Series



AP 504 Wi-Fi 6 AP

- 2 RP-SMA connectors for external dual band antennas

- 2x2 MIMO up to 1.2Gbps for 5GHz, 574Mbps for 2.4GHz
- 10/100/1000Base-T and MDI/MDX
- PoE-PD: 802.3af/at (class 3 or 4)
- 802.3az Energy Efficient Ethernet (EEE)

AP 505 Wi-Fi 6 AP

- 2 Integrated dual-band downtilt omni
- Optimized for horizontal ceiling mount



MOD 6- 16

The figure summarizes key specifications for the Aruba 500 series APs. Please see the Aruba 500 Series Indoor Access Point data sheet for more details:

<https://www.arubanetworks.com/products/wireless/access-points/indoor-access-points/500-series>.

Aruba AP 530 Series



AP 534 Wi-Fi 6 AP

- 4 RP-SMA connectors for external dual band antennas

- AP type: Indoor, dual radio, 5GHz and 2.4GHz 802.11ax 4x4 MIMO
- 5GHz radio: 4 spatial stream MIMO, up to 2.4Gbps data rate
- 2.4GHz radio: 4 spatial stream MIMO, up to 1,147Mbps data rate
- POE-PD: 48Vdc (nominal) 802.3af/at/bt POE (class 3 or higher)
- 802.3az Energy Efficient Ethernet (EEE)

AP 535 Wi-Fi 6 AP

- 4 Integrated dual-band downtilt omni
- Optimized for horizontal ceiling mount



MOD 6-17

The figure summarizes key specifications for the Aruba 530 series APs. Please see the Aruba 530 Series Indoor Access Point data sheet for more details:

<https://www.arubanetworks.com/products/wireless/access-points/indoor-access-points/530-series/>.

Aruba AP 550 Series



AP 555

- Indoor, dual/tri-radio, 5 and 2.4GHz 802.11ax 4x4 MIMO
- 5GHz dual-radio: 8 spatial streams, up to 4.8Gbps data rate
- 5GHz tri-radio: 4 spatial streams, up to 2.4Gbps data rate
- 2.4GHz: 4 spatial streams, up to 1,150Mbps data rate
- E0, E1: HPE SmartRate port (RJ-45, max speed 5Gbps)
- 100/1000/2500/5000BASE-T and MDI/MDX
- POE-PD: 48Vdc (nominal) 802.3af/at/bt POE (class 3 or higher)
- 802.3az Energy Efficient Ethernet (EEE)



MOD 6- 18

The figure summarizes key specifications for the Aruba 550 series APs. Please see the Aruba 550 Series Indoor Access Point data sheet for more details:

<https://www.arubanetworks.com/products/wireless/access-points/indoor-access-points/550-series/>.

Aruba AP 600 Series



AP 650

- Very high-capacity Wi-Fi 6E indoor access point
- 7.8 Gbps max aggregate throughput
- High availability with configurable 5 Gbps dual Ethernet ports

AP 630

- High-capacity Wi-Fi 6E access point
- 3.9 Gbps max aggregate throughput
- Up to seven 160 MHz channels in 6 GHz

AP 610

- Entry-level Wi-Fi 6E indoor access point
- 3.6 Gbps max aggregate throughput
- 2.5 Gbps Smart Rate Ethernet port to minimize wired bottlenecks

Tri-band: supports 5 and 6GHz bands (6 and 5, or 2.4 and 5GHz)



MOD 6-19

AP 650

- Comprehensive tri-band coverage across 2.4GHz, 5GHz, and 6GHz to deliver 7.8 Gbps maximum aggregate data rate
- 4x4 MIMO radios to deliver peak performance and increased capacity using MU-MIMO and OFDMA (uplink and downlink for both)
- Up to seven 160MHz channels in 6GHz support low-latency, bandwidth-hungry applications like high-definition video and augmented reality/virtual reality
- applications
- Unique ultra tri-band filtering enables 5GHz and 6GHz to operate without restrictions or interference
- High availability with configurable 5 Gbps dual Ethernet ports for hitless failover of ethernet and power

AP 630

- Unlocks the 6 GHz band to more than double the available capacity
- Comprehensive tri-band coverage across 2.4 GHz, 5 GHz, and 6 GHz to deliver 3.9 Gbps maximum aggregate data rate
- Up to seven 160 MHz channels in 6 GHz support low-latency, bandwidth-hungry applications like high-definition video and augmented reality/virtual reality applications
- Unique ultra tri-band (UTB) filtering enables 5 GHz and 6 GHz to operate without restrictions or interference
- High availability with 2.5 Gbps dual ethernet ports for hitless failover of Ethernet and power

AP 610

- Unlocks the 6 GHz band to more than double the available capacity in wireless networks
- Comprehensive dual radio/tri-band coverage across 2.4 GHz, 5 GHz, and 6 GHz to deliver up to 3.6 Gbps combined peak data rate
- Compact and low power consumption: virtually unrestricted operation from an 802.3af (class 3) PoE source
- Up to seven 160 MHz channels in 6 GHz to support low-latency, bandwidth-hungry applications like high-definition video, and augmented reality/virtual reality applications
- 2.5 Gbps Smart Rate Ethernet port to minimize wired bottlenecks
- Built-in GPS receivers and fine time measurement (802.11mc) to enable APs to automatically locate themselves and serve as reference points for accurate indoor location measurements

Aruba AP 300 Series



AP 340

- High-performance Wi-Fi 5 campus access point
- Support for dual 4x4 5GHz Wi-Fi
- Scales up to multi-gig Ethernet



AP 303

- Low-cost Wi-Fi 5 campus access point
- Ideal for medium-density enterprise environments



MOD 6-20

AP 340

- Designed with an integrated, 802.3bz compliant, HPE SmartRate multi-gigabit Ethernet port to eliminate wired bottlenecks
- Maximum data rates of 2,166 Mbps in the 5 GHz band and 800 Mbps in the 2.4 GHz band (for an aggregate peak data rate of 3.0 Gbps in dual-band and 4.3 Gbps in dual-5GHz)
- Includes integrated Bluetooth Low Energy (BLE) radio, for advanced location and indoor wayfinding
- Participates in Aruba's Dynamic Segmentation solution


AP 303

- Up to 1.2 Gbps aggregate peak data rate
- Daisy-chain another AP using the secondary PoE-out port
- Integrated Zigbee and Bluetooth 5 radio for IoT connectivity
- Resolve sticky client issues with MU-MIMO-aware ClientMatch

- Participates in Aruba's Dynamic Segmentation solution

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Aruba Remote/Hospitality APs



AP 303H

- High-performance Wi-Fi 5 access point
- Flexible wall or desk mount options
- Ideal for hospitality and branch offices



AP 505H

- High-performance, cost-effective Wi-Fi 6 access point
- Support for PoE Out on wired ports
- Ideal for hospitality, branch, and teleworker deployments



MOD 6- 21

AP 303H

- Cost-effective Wi-Fi 5 (Wave 2) Access Point. A flexible, secure solution ideal for hospitality, branch offices, medical clinics, and remote workstations.
- Unified dual-radio 802.11ac Wave 2 2x2:2 hospitality and branch AP with internal antennas, three local Gigabit Ethernet ports, PoE out and USB host interface
- Supports wall-box and desk mount deployments

AP 505H

- High-end dual radio Wi-Fi 6 Hospitality AP with 1+4 Ethernet ports, PSE, USB
- High performance, cost-effective, medium-density, 2X2 radios
- 505H High-end dual radio, four Ethernet ports, 2X PoE-outdoor, 500 Mbps
- 503H Mid-range radio, two Ethernet ports, PoE Indoor, 100 Mbps cryptographic VPN performance

Aruba Gateways



9000 Series

Small/midsize branch
4 Gbps firewall throughput



7005/7008

Small branch w/ 4-8 ports and
optional PoE
2 Gbps firewall throughput



7010/7024

Midsize/large branch w/16-24 ports
and PoE
4 Gbps firewall throughput



7030

Large branch/small campus
8 Gbps firewall throughput



Mobility Controller (VA)

Branch and campus
8 Gbps firewall throughput

When to use gateways:

- ✓ VLAN 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 6- 22

Aruba Gateways

The main reason you need gateways in a large-scale deployment is to simplify network design and provide scalability. The gateway is responsible for traffic termination, so all wireless traffic is being tunneled from the APs to the gateway, which then passes through the firewall function on the gateway and from there it is forwarded on its way to the traffic's final destination. This eliminates excessive span of VLANs. Aruba offers both gateways and controllers to support the needs of larger enterprises who require additional scalability, 24x7 operations, and greater security. Gateways are optional: in other words, you can run the APs in standalone mode: but if you need scalability and simplicity, gateways are the solution of choice.

You should use a gateway deployment when any of the following is required:

- Simplified VLAN design
- SD-WAN
- Security: tunneled WLANs, segmentation/guest WLANs, network encryption
- More scalable and robust mobility: layer 2 and /or layer 3 roaming across domains
- Dynamic Radius Proxy

- Scale: More than 5,000 clients per site or more than 500 APs per site

Note: In AOS 8, there were two different terms for the hardware shown in the figure: gateways and mobility controllers (MCs). Gateways provided SD-WAN connectivity and MCs terminated the APs. Starting in AOS 10, only a single term is used for these devices: gateways.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Aruba 7200 Campus Gateway



- ✓ Gateway for large enterprises
- ✓ Up to 32k concurrent clients
- ✓ 100 Gbps maximum wired throughput
- ✓ 40G ports
- ✓ 1 RU



MOD 6- 23

Aruba 7200 Series gateways enhance WLANs for high-performance and high-density enterprise requirements, ideally suited for large campuses. The 7200 Series can be deployed using zero touch provisioning (ZTP) to simplify deployment. The 7200 Series provides connectivity for up to 32,728 concurrent users or client devices and over 2 million active firewall sessions. With up to 80 virtual CPUs and 100Gbps of maximum throughput to perform policy enforcement firewall (PEF) features, the 7200 series are in a class of their own: they are ideal for the most demanding enterprise, college, and large public venue requirements.

Aruba 9200 Series Campus Gateway



- ✓ Next-generation gateway for medium and large enterprises
- ✓ Increase in storage and memory
- ✓ 40 Gbps maximum wired throughput
- ✓ Capacity depends on license
- ✓ 1 expansion slot




MOD 6- 24

The Aruba 9200 Series Campus Gateway delivers enhanced Wi-Fi scalability and security to meet demands at the network edge. Engineered for mission-critical environments, the 9200 Series delivers high-performance, seamless roaming, 24x7 reliability, and always-on connectivity: even during upgrades. Unlike other edge appliances, the 9200 Series provides consistent policy enforcement across wired and wireless and dynamic segmentation for any user or device, even IoT.

The base unit provides 20 Gbps of throughput, 64K devices, 500K IP addresses, 32K GRE tunnels, and 32K IPsec tunnels. Applying a hardware capacity license will increase the performance and scalability of the 9200. With the Silver license the maximum throughput is 30 Gbps and 128K devices, 1M IP addresses, 64K GRE tunnels, and 64K IPsec tunnels. And finally, with the Gold license, you can achieve 40 Gbps throughput and 256K devices, 2M IP addresses, 128K GRE tunnels, and 128K IPsec tunnels.

Running Aruba APs and Gateways with AOS10



Evolution to unified operating system

MOD 6- 25

Aruba offers APs and gateways running AOS 10.x. You can now deploy and manage your WLANs from a single and unified cloud-based network management system: Aruba Central. This combination of AOS 10.x-based devices and Aruba Central is powerful! It simplifies network administration with automated workflows, end-to-end visibility, AI-powered insights, and analytics to enhance and optimize the wireless experience for users.

Let's explore AOS 10 a bit more.

ArubaOS 10: Overview

Powered by **Aruba Central**



Analytics



Configuration
Management



Control Plane
Services




Monitoring and
Reporting



Security




Location



Large Campus



Small/Medium Campus



Branch



Hybrid
Workplace




Microbranch

MOD 6-26

AOS 10.x provides a single architecture for customer WLAN access deployments. It integrates access networks, cloud service offerings like Aruba Edge Services Platform (ESP), which provides a unified architecture for branch and campus, wired and wireless, with a no sub-WLAN architectures such as AOS6/8 and Instant 4/6/8: only AOS 10 is needed.

ArubaOS 10: Architecture



AOS 10.x is comprised of three layers.


- A required Infrastructure Layer of on-premise WLAN APs and switches
- A required central services layer for services, applications, and NMS/CMRT
- An optional gateway Layer of on-premise security gateways

The advantages of this architecture are that it:

- Scales beyond the Mobility Conductor 8.x domain of 10k APs or Instant Clusters of 2K devices
- Enables cloud services including AI-powered insights, workflow automation, and robust security
- Unifies WLAN, wired LAN and WAN management and visibility with consistent UI and workflows

Let's zoom in on the optional gateway layer.

ArubaOS 10: Architecture



MOD 6-28

AOS 10 provides additional scalability and security for enterprise customers. Customers should consider a gateway if they have one or more of the following requirements:

- Scale higher than 5,000 clients or 500 APs (Note that this is due to management complexity and not a hard technical limit)
- Security: Tunnelled WLANs, segmentation with guest WLANs, network encryption over their LAN, or user-based tunneling (UBT)
- Greater mobility: scalable and robust L2/L3 roaming
- MultiZone support to allow different gateways to work with the same AP infrastructure in a secure, segmented manner
- Dynamic RADIUS Proxy (DRP): Ensures RADIUS traffic is sent to the correct server. Useful if corporate employee traffic should use a local RADIUS server, and guest traffic should use a centralized captive portal server.

Note: The gateway is an optional component. Aruba Virtual Gateways are not currently supported for wireless use cases.




Introduction to Aruba Central

Let's take a look at Aruba Central.

CONFIDENTIAL - FOR TR

Aruba Central



Aruba Central is a cloud-based platform for managing Aruba-based wireless, WAN, and wired networks. Aruba Central supports group configuration to provision and manage similar devices with less work. Continuous monitoring and AI-based analytics provide real-time visibility and insight. Machine Learning leverages a growing pool of network data and deep experience. It is very secure, with HTTPS connections, certificate-based and cloud authentication and policies.

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.

Note: Central runs part of the control plane for some network devices and solutions, using the Aruba SD-Branch solution and Aruba APs and gateways running AOS10.

Cloud-Managed

Simple



- No management hardware or software to maintain
- Easy network deployment

Reliable



- Guaranteed uptime
- Local management if cloud connectivity fails

Intelligent



- Capabilities to ensure performance & visibility
- Analytics to improve business & user experience

Flexible



- Easy to scale network
- Ongoing innovation
- Subscription options
- Anytime, anywhere access



MOD 6- 31

Central is simple, with no hardware or software requirements on the customer's end. It's in the cloud, with all the resources needed for performance and visibility. Central is reliable, with guaranteed uptime, and local management in case your connectivity to Aruba Central experiences issues or failures.

Plus, you get intelligent features to improve performance and visibility, with analytics to improve the business and user experience. And flexibility? It's easy to scale up as needed, to accommodate growing systems and ongoing innovations, with simple subscription options and anytime, anywhere access.

Flexible and Cost-Effective Cloud Networking

HARDWARE

Limited lifetime warranty

APs

IAPs
AOS10




Switches

Aruba-OS
Aruba CX



Gateways

SD-Branch
AOS10



CLOUD MANAGEMENT

License per managed device
Transferrable between devices

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



MOD 6- 32

Central supports management, monitoring, and configuration for the Aruba IAPs, Aruba AOS and AOS-CX switches, SD-Branch gateways, and AOS10 APs and gateways. And since it's a cloud-based management system, it only requires a simple 1, 3, 5, 7, or 10-year subscription license, which includes tech support, with more additional services available.

Aruba COP

Delivering cloud-like scale and efficiency



Controllers,
Conductors, &
Controller-based APs

Enterprise Scale

Supports up to 25K
network devices



Instant APs

Built-in Redundancy

No loss of operations if
an appliance fails



AOS-S and AOS-CX
Switches

Programmability

APIs and webhooks to
automate workflows with
3rd-party tools

MOD 6-33

We want to provide organizations with the choice and flexibility to deploy and manage their networks in ways that best align with their technical requirements. And that's why we offer Aruba Central On-Premises (COP) to complement our Aruba Central cloud offering.

Aruba COP delivers most of the same services and capabilities of Aruba Central. It has the same user interface, same code, and same user experience and workflows: it's just powered by purpose-built server appliances that are deployed in a customer's data center environment. In this way, networking teams still gain a cloud-like management experience for better operational agility and efficiency, while also satisfying security or compliance mandates that require on-premises infrastructure.

Like AirWave, Aruba COP is designed to simplify the deployment and maintenance of wireless and wired networks by unifying operations under a single pane of glass, with support for Aruba access points : both gateway-based and instant APs, as well as mobility controllers and/or conductors and both portfolios of Aruba switches (AOS and AOS-CX). But an added advantage of Aruba COP is that it delivers much higher scalability: supporting up to 25K devices in seven-node server cluster.

COP has limits in terms of number of managed devices depending on the number of servers hosting COP. With three-node deployment you can manage up to 8,000 devices, with a five-node deployment supports up to 16,000 devices, and a seven-node deployment supports up to 25,000 devices. In addition to this, there are limits for AI Insights that can only be used in a five-node and seven-node deployment. A five-node deployment support up to 160,000 clients

and a seven-node up to 300,000 clients.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY



Introduction to Aruba ClearPass

Aruba ClearPass helps to ease security operations and improve your security stance.

CONFIDENTIAL - FOR TR

Aruba ClearPass

Authentication server that handles all NAC functions

Services

- Policy Engine
- 802.1X
- MAC Auth
- Guest
- TACACS+
- Profiling
- Context Database
- 100++ RADIUS dictionaries

IT Tools

- Policy Simulation
- Access Tracking
- Template-based policy creation
- LDAP Browser
- Per Session Logs
- Advanced Reporting
- AirGroup

Security Exchange (3rd Party Integration)

- API's
- Syslog Feeds
- Extensions
- Ingress Events



MOD 6- 35

Aruba ClearPass is an authentication server that can easily handle all network access control (NAC) functions. It is a decision-making point when clients connect to the network and go through authentication and authorization process. One of the great things about ClearPass is that we include many features within the base software set. Aruba has had TACACS and RADIUS servers for over 10 years, so they're tried and tested. The policy engine, basic guest, and device registration tools work without the need for dedicated appliances or extra licenses.

The visibility dashboard, policy creation templates profiling, and troubleshooting services have been there since we've been shipping the product. The exchange interoperability components (API and syslog messaging) are included. Customers are free to use built-in integration templates or create their own workflows. And, if they want, customers can try any of the three expandable applications because we include 25 enterprise licenses with the base platform. ClearPass is available as an appliance or a virtual machine (CVM) installation.

ClearPass Applications

- ✓ Automated workflows
- ✓ Rules by user role and device types

Onboard



OnGuard



MOD 6- 36

ClearPass Onboard is the easiest way to support a bring-your-own-device (BYOD) environment. Based on its own internal certificate authority (CA), Onboard provides an intuitive and easy to use portal for your users to securely configure their devices to connect to your company's network. To enforce the concept that every device should be granted access on its own unique merits, ClearPass Onboard assigns a unique identity to each device that goes through the Onboard process. For auditing capabilities, this identity includes information about the client device as well as the user that put the device on the network.

Posture checks and client system health validation are necessary for most networks. Using an agent, either deployed through a web browser to the client or as a permanent software applet, ClearPass OnGuard can provide security status validation on your wired, wireless, and VPN networks. OnGuard makes it easy to enforce the organization's device compliance policies.



Knowledge Check

Self-check on key learning points



Please check your understanding with a brief knowledge check.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Question #1

Which of the options below are the main aspects of Aruba's Edge Services Platform (ESP) architecture? (Select all that apply.)

- a. Artificial Intelligence Operations (AIOps)
- b. Zero Trust Networking
- c. Aruba Central
- d. Unified Infrastructure



Knowledge Check

MOD 6- 38

A,B,C,D

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Question #2

Which Aruba feature solves the challenge of tedious switch port configurations and securing devices that lack support for 802.1X/EAP authentication?

- a. IPsec tunneling
- b. Dynamic Segmentation
- c. Analytics dashboard
- d. Third-party integration



Knowledge Check

MOD 6- 39

B

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Question #3

Which Aruba Access Points (APs) support three RF bands: 2.4GHz, 5GHz, and 6GHz?

- a. All Aruba APs
- b. All 500 and 600 series APs
- c. All APs but the older 300 series and hospitality APs
- d. Only the 600 series APs
- e. No Aruba APs support all three bands



Knowledge Check

MOD 6-40

D

Question #4

Which statements below accurately describe Aruba Gateways? (Select two.)

- a. They can simplify VLAN design by acting as a tunnel termination point for Generic Routing Encapsulation (GRE) tunnels.
- b. They are required components of all AOS-CX validated design deployments.
- c. Running Aruba APs with optional Gateways on AOS10 enables you to manage WLANs from Aruba Central.

a

Knowledge Check

MOD 6- 41

A,C

Question #5

Aruba Central enables you to manage APs, switches, and gateways at all locations from a single portal, supported only as a cloud-based solution.

- a. True
- b. False



Knowledge Check

MOD 6- 42

F

Question #6

Aruba ClearPass is an authentication server that handles all Network Access Control (NAC) functions, with extensive third-party integration features.

- a. True
- b. False



Knowledge Check

MOD 6- 43

T

Rev.22.41




Central for device management

Aruba Campus Access
Fundamentals

Hello and welcome back. This is Module 7: Central for device management.

CONFIDENTIAL - FOR TRAINING USES ONLY



Objectives

- Perform device onboarding
- Create Central Groups
- Use UI config mode
- Use template config mode
- Describe Central licensing

Aruba Central provides centralized and cloud-based management platform for Aruba's ESP.

After completing this module, you will be able to perform Aruba Central device onboarding and create and use Central groups to ease administration.

You will be able to use user interface (UI) and template configuration mode and describe Central licensing.

Device Onboarding

We begin with device onboarding.

CONFIDENTIAL - FOR TRAINING

Device Onboard Workflow



MOD 7- 4

The onboard device process starts at the device inventory. First, you must add any device that you want Aruba Central to manage to the device inventory list. Once you've added the devices, you must assign each device a subscription. This enables Aruba Central to manage the device. Finally, you configure your devices.

Adding Devices to Central Inventory

Zero Touch Provisioning

Manually Adding Devices

Using Cloud Activation Key

Using Install Manager




MOD 7- 5

There are four ways to have devices communicate with Central. Zero-touch is the best way as it simplifies installations. All devices receive a cloud activation key that you can use to direct devices to Central. Alternatively, you can also manually add a device to Central.

You can download the Aruba app to your smartphone to use the zero-touch option. This app allows you to simply take a picture of the device's serial number (SN) and MAC address. Then Central will automatically update Activate with the SN and MAC. If you have access to an Activate account, then you can add in Activate rules to assist in the zero-touch provisioning (ZTP) process. If not, you can call Aruba support and they can assist you in setting up the appropriate Activate rule or rules for your devices.

Aruba Activate



Zero Touch Provisioning

Firmware repository




MOD 7- 6

Before starting the Zero Touch Provisioning (ZTP) process, you need to understand what Aruba Activate is: it is a key component for provisioning Aruba devices. Aruba Activate is a cloud-based service designed to enable customers to efficiently manage their Aruba device inventory and facilitate the provisioning of their Aruba devices. Use of Aruba Activate is free to customers.

Aruba Activate provides two main services:

- Zero-touch provisioning system for all Aruba networking devices.
- Firmware repository for Aruba networking devices.

New Customer



The figure shows an example of how a new customer receives the Aruba Central key and joins the Central network using ZTP and Aruba Activate.

Here are the steps for the process:


1. First, the new customer, in this case MyCorp Inc., sends Aruba a purchase order (PO) for 100 devices (Aruba APs, for example). Aruba delivers the 100 devices (APs) as well as the Aruba Central key to MyCorp.Inc.
2. Next, Aruba updates Activate with the PO information including the customer's name, the subscription key (X in the figure) and the 100 APs they purchased.
3. Aruba uses Activate to dynamically provision APs and gateways to point them to Central. This is done by Aruba creating a rule in Activate that directs these devices to contact Central when they boot up at their factory default settings.
4. Now John Customer, a MyCorp.Inc employee, registers for Aruba Central using his company email. Then John adds the Central subscription key that Aruba sent with the purchase confirmation.
5. The Central account communicates with Activate and gives the customer's name to Activate.

6. Then Activate informs the Central account about all 100 APs that belong to this account. The APs will show up in MyCorp.Inc's Central inventory under My Devices.

As you can see from this process, Aruba has automated most of the process of getting the newly purchased devices into Aruba Central.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Zero Touch Provisioning: Device Functionality



ZTP: Device functionality

AOS10 AP

You might choose to assign an AP to an AOS10 group with a firmware policy. Therefore, after communication with Aruba Central, the AP automatically downloads the AOS10 firmware and reboots. Once the upgrade completes, the AP again communicates with Aruba Central, which pushes the AP configuration.

Switches and Gateways

Aruba switches and gateways in factory default mode will do the same. They boot up, get an IP address, and attempt to communicate with Activate. If successful, Activate directs the switch to a Central account where the switch it gets its configuration.

HPE GreenLake Management Task Overview

View, manage, and onboard all devices in your account



HPE GreenLake

Onboard
devices

Add manually: CSV file or enter S/N and MAC per device

Onboard w/Cloud Activation Key: not supported by AOS-CX and AOS10 APs



Aruba Central

Assign app

Assign devices to an HPE GreenLake application: Aruba Central



MOD 7-9


HPE GreenLake is a cloud management and monitoring platform from HPE. It hosts and provides not only Aruba Central services but many other HPE services as well.

HPE GreenLake's device management feature let's you easily onboard, view, and manage all the devices in your account.

One of the first things you do is onboard your various network devices. You can add them manually, using a .csv file or manually enter each device's serial number and MAC address. You can also onboard devices using an Aruba cloud activation key. This method is not supported by AOS-CX switches or AOS10 APs.

After onboarding devices to HPE GreenLake, you must assign them to the correct HPE GreenLake application like Aruba Central. You will learn to perform all these activities in the lab. For now, let's look at some items related to using that cloud activation key.

Aruba Cloud-based Management Flow



Getting Aruba devices show up in Central involves the following steps:


1. First the customer registers the devices in HPE GreenLake as described earlier.
2. HPE GreenLake communicates information about the devices to Aruba Activate as those devices are assigned to Aruba Central application.
3. Now that Aruba Activate is aware of the devices' serial number and MAC address, it will accept ZTP requests coming from these devices.
4. Aruba Activate upgrades device firmware, if necessary, and directs it to the appropriate Central cluster.
5. The devices establish an HTTPS session to Aruba Central cluster and now appear in the Central web interface.
6. The customer can now start configuring and monitoring the devices via Central.

Aruba Central Licensing

Now you will explore Aruba Central licensing.

CONFIDENTIAL - FOR TRAINING

Aruba Central 2.5.5 Licensing



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

MOD 7-12

Aruba Central supports foundation and advanced license keys. These keys are for APs, switches, SD-Branch gateways, and gateways running AOS10. 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 figure.

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: Switch foundation license for 6200 switch series and foundation license for 6400 switch series; and Advanced Gateway licenses for the 7000 gateway family, and Advanced license for 7200 gateway family.

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

Subscription Management

Activities & Tasks

- Get Started**
Not sure where to start? Let's walk you through it! These videos should help you get started.
[Watch Videos](#)
- App Catalog**
Find and deploy applications for your organization's internal use.
[Explore Catalog](#)
- Onboard Devices**
Add and assign devices in your device inventory.
[Manage Devices](#)
- Account Details**
Edit your company account's general information.
[Edit Account](#)
- User Profile**
Manage your user profile.
[Manage User Profile](#)
- Invite Users**
Send a sign-up link to users so you can collaborate with your team.
[Send an Invite](#)
- User Roles**
Manage built-in roles or create custom roles.
[Manage Roles](#)
- Assign User Access**
Assign a custom role or built-in roles to your team.
[Assign Roles](#)
- Subscriptions**
Manage your device and service subscriptions.
[Add Subscription](#)

MOD 7- 13

You view and manage HPE GreenLake subscription keys at the Subscription Management page (HPE GreenLake > Activities & Tasks > Subscriptions). The Subscriptions page shows your current subscriptions and allows you to add more as needed. To access it, click the Subscriptions tile as shown in the figure.

Subscription Management

Device Subscriptions

Manage and add device subscription keys.

Add Device Subscription

Search Subscription Keys



31 Subscription Key(s)

Key	Tier	Quantity	Open Seats	Expiration Date	
AAONEB9YCGS86Y8R	FOUNDATION_70XX	10	10	Jul 29, 2030	...
EAD6EFLXNQWTXZ7Z	FOUNDATION_70XX	1	1	Jul 29, 2030	...
E68GLVGZRMGTGCD4	FOUNDATION_SWITCH_6100	1	1	Jul 29, 2030	...
ALM4VXDQOXX7ANT	BASIC_AP	0	0	Jul 29, 2030	...
EOJDMB9UK8TLSSIF	FOUNDATION_SWITCH_6200	3	1	Jul 29, 2030	...
EXVHCAVNG2XKK6ER	FOUNDATION_SWITCH_6300	3	2	Jul 29, 2030	...

MOD 7- 14

From here (see figure) you can see a list of current subscriptions and add more subscriptions. Just click the Add Device Subscription button.

Subscription Assignment

Activities & Tasks

**Get Started**

Not sure where to start? Let's walk you through it! These videos should help you get started.

[Watch Videos](#)

**App Catalog**

Find and deploy applications for your organization's internal use.

[Explore Catalog](#)

**Onboard Devices**

Add and assign devices in your device inventory.

[Manage Devices](#)

**Account Details**

Edit your company account's general information.

[Edit Account](#)

**User Profile**

Manage your user profile.

[Manage User Profile](#)

**Invite Users**

Send a sign-up link to users so you can collaborate with your team.

[Send an Invite](#)



MOD 7- 15

After adding Subscription keys, you must assign them to Aruba devices. From the HPE GreenLake management page, select Manage Devices, which will take you to the page shown in the figure. Then click the Onboard Devices tile.

License Assignment

Device Management

Overview

Devices

Auto-Subscribe

Devices

View and manage all devices in your inventory.

Add Devices

Devices Require
Assignments
0



Devices Require
Subscriptions
1



Assigned &
Subscribed
10



Total Devices
11



Search by Serial, Model, or MAC Address



Remove Assignment

Apply Subscription

1 of 1 Device(s) selected

☒	Serial Number	Model	MAC Address	Application	Subscription Tier
☒	CNDBJSS0FX	JX945A	C8:B5:AD:C7:85:E8	Aruba Central Internal (Central_Internal)	--



MOD 7- 16

You are then taken to the window shown in the figure. Complete the license assignment here (HPE GreenLake > Activities & Tasks > Manage Devices). Choose devices to which the subscription shall be assigned and click Apply Subscription. To update license assignments, select the device, click Detach Subscription and assign a new subscription to the device. To avoid the manual process, you can select Auto-Subscribe.

For now, let's assume you click Apply Subscription.

Manual License Assignment

Apply Subscriptions To Devices

Apply subscription keys to your selected devices.

Access Points

Model JX945A

0 out of 1 devices subscribed

Apply Subscriptions

Finish



MOD 7- 17

You are taken to the Apply Subscriptions to Devices page, shown in the figure. For each selected device type, click Apply Subscriptions.

Manual License Assignment

Apply Subscriptions

Select a subscription tier and subscription keys to assign to **1 Access Points Model JX945A**

Select Subscription Tier

ADVANCED_AP

ADVANCED_AP
FOUNDATION_AP

☒	Subscription Key	Tier	Available	Expiration Date
☒	EXNKGXBAT3GWRNNC	ADVANCED_AP	996	Dec 30, 2030

1 subscription selected (0 more devices require subscriptions)

Clear Selection

Cancel

Apply Subscriptions

Then, choose the subscription type (Subscription tier) to be applied to the device and choose the Subscription key you want to use for this device. Click Apply Subscriptions.

Auto-Subscribe

Devices

Manage your devices and their subscriptions.

Inventory

Auto-Subscribe

Auto-Subscribe

Configure auto-subscribe for each supported device type.

Add

Access Points

Subscription Tier

Select Subscription Tier

Foundation AP (10 seats left)

Advance AP (9 seats left)

Advance Gateway (30 seats left)

Cancel

Save

Delete

Edit

MOD 7- 19

With Auto-Subscribe you can assign licenses to all new devices automatically when they connect to Central.



Central Groups, Sites and Labels

Aruba Central groups, sites, and labels can ease workflows and save you time.

CONFIDENTIAL - FOR TR

What are Central Groups and Sites

Groups

- Used for management, configuration, monitor, filter and reports
- A device can belong to only one group at a time

Sites

- Used for monitor, filter and reports
- Should follow physical location
- A device can belong to one site at a time



MOD 7- 21

Groups

Groups in Aruba Central are a primary 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, or venue. In Aruba Central, you use sites as a primary navigation element. So, 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 place them in the same group.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Aruba Central group



MOD 7- 22

A group's main purpose is to simplify your device configuration. All devices in the same group will receive the same configuration. However, you can also use groups to monitor devices. 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 based on the group with which they are associated. 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 figure shows an AP group named AP-Building-1. In this scenario, these APs must support three SSIDs: employee, voice, and guest. Therefore, 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 securely 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 Switch-Building-1 for the Aruba switches' configuration in the building.

Group Types

UI Group

Configure with a web browser

Template Group

Configure with CLI-based templates



MOD 7- 23

Group types include:

UI group

You configure UI group device members using a web interface (GUI) within Central, 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.

Template group

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.

Default

The default group is a system-defined container that consists of devices running factory default configurations. When a new device with a 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 administrator 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 7- 24

A device's 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.

Creating a Group Persona



MOD 7- 25

You can create a persona when you create a group. Persona and architecture can be set at the group level. All devices within a group inherit the same persona from the group settings.

While creating a group, you can mark 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 points tab.

Group Personas For Devices

APs



Campus/Branch:

WLAN functionality

ArubaOS 10

ArubaOS 8 (including IAP-VPN)

Microbranch:

WLAN, SD-WAN-lite

ArubaOS 10

Gateways



Mobility:

Overlay WLAN, wired LAN

ArubaOS 10

Branch:

Aruba InstantOS SD-Branch
(LAN+WAN)

ArubaOS 10

ArubaOS 8

VPN Concentrator:

VPN concentrator

Microbranch: ArubaOS 10

SD-Branch: ArubaOS 8 or ArubaOS 10

MOD 7- 26

Access Points can have the following persona:

- Campus/Branch—in this persona, the AP provides WLAN functionality. This persona applies to both ArubaOS 10 and legacy ArubaOS 8 (including IAP-VPN) architectures.
- Microbranch—in this persona, the AP provides SD-WAN-lite functionality in addition to the WLAN AP functions. This persona applies only to ArubaOS 10 architecture.

Gateways can have the following persona:

- Mobility—in this persona, gateways provide Overlay WLAN and (or) wired (UBT) LAN functionalities. This persona applies only to the ArubaOS 10 architecture.
- Branch—in this persona, gateways provide Aruba InstantOS SD-Branch (LAN + WAN) functionality. This persona applies to both ArubaOS 10 and ArubaOS 8 architectures, with ArubaOS 10 being a superset of ArubaOS 8.
- VPN Concentrator—in this persona, gateways provide VPN concentrator functionality for ArubaOS 10 Microbranch or ArubaOS 8 or ArubaOS 10 SD-Branch deployments.

The two architectures for creating groups are highlighted here:

- ArubaOS 8: Instant AP-based deployment, including Aruba InstantOS 6.x or Aruba

InstantOS 8.x (IAP, IAP-VPN), or Aruba InstantOS 8.x SD-Branch deployments.

- ArubaOS 10: ArubaOS 10 deployment, including AP-only underlay, AP+gateway overlay, Microbranch, or AOS 10.x SD-Branch deployments.


You can apply a device persona to both ArubaOS 10 and ArubaOS 8 deployments. All of the ArubaOS 8 personas apply to ArubaOS 10 deployments as well. The ArubaOS 10 deployment supports a couple of additional personas that do not apply to ArubaOS 8 deployments. The persona workflow differs based on the deployment type.

Note: A group cannot manage ArubaOS 8 and ArubaOS 10 devices. The ArubaOS version is selected as part of the group creation process. All the devices in that group must comply with the group persona.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Groups: Template and UI-based Configuration

Define a preferred configuration method during group creation



MOD 7- 27

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, (for example, 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 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 "Unprovisioned" group to one of your defined groups
After move, device reloads, inherits group config (loses current config)

MOD 7- 28

Create the Group

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

You can mix/match UI and Templates in a single group: for example, a UI configuration for access points and gateways, and use a CLI template configuration for the 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 the group between UI and Template groups after creating the group.

Define the Group

You then define the group by choosing 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 ArubaOS8 at the time of the creation, the only way to manage ArubaOS10 devices is by creating a new group with the AOS10 architecture.

Add Devices to the Group

After you define groups, you move devices into those groups as members – perhaps a devices from the unprovisioned group to one of your defined groups. Be careful: after a device is moved to a new group, it reloads and inherits the group configuration. This means that the device will erase the 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 7- 29


A site refers to a physical location where a set of devices are installed; for example, campus, branch, or venue.

Sites are more than just a tool to filter views, place devices on a map, and generate reports. You can use sites for several Aruba Central services such as live upgrades, AirMatch, install manager, network health dashboard, User Experience Insight (UXI), etc. As a best practice, you should assign every device to an Aruba Central site. Devices can only belong to one site.

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 or postal code is not.

Note: VisualRF, which will be covered later in this course, 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.

Using Open Location Code



CREATE NEW SITE

SITE NAME
New Building

STREET ADDRESS
QMMP+MWQ

CITY
Roseville, California, USA

United States ▼

California ▼

ZIP/POSTAL CODE


Add

MOD 7- 30

The Open Location Code (OLC) is a geocode system for identifying an area anywhere on the earth. Location codes created by the OLC system are referred to as plus codes. Aruba Central supports plus codes when creating new sites. Simply use plus code as street address. The site will be automatically positioned on the map in Central. A good example of why OLC is useful would be a campus that has multiple buildings. The street name and house number for each building is the same because it is the same campus. Therefore, entering an address when creating site would not be precise. Plus codes allow you to specify the exact location when setting up the site.

You will learn how to add, modify, and manage sites during the lab activities. Keep in mind that many Aruba Central services use the site information for planning and managing the devices and their activities. Therefore, be sure that all devices are assigned to the right site to allow those services to run correctly.

Sites and Network Health MAP



MOD 7- 31

Sites with proper addresses will appear on the network health map. The map provides a pictorial view of the various network sites. The sites are color-coded with red or green. Red indicates potential issues, and green indicates no issues. You can click the site on the map to view details on any issues at this site. You can also zoom into a site and get to the street level view.

Labels

**Aruba
Central**



Filter devices for monitor, manage, report purposes

Help determine device ownership, departments, functions

Label – City-ABC

Label – Building-1

Label – BLD-1-FLR-1

Label – BLD-1-FLR-2

Label – BLD-2

Label – Conference rooms

City ABC

Building-1

AP Building-1 Floor 1



AP Building-1 Floor 2



AP Building-2



MOD 7- 32

Labels are tags that are used to filter devices for monitoring, managing, and reporting purposes. A device can have multiple labels, which can help you determine device ownership, departments, and functions.

Suppose you have several buildings in a city, all with their own devices. Therefore, you could create a label of BLD-1-FLR-1, and all the devices in building 1, floor 1 would associate to that label. Select the label to see only the devices associated with it and unselect the label to remove the filter.

You might create other labels as shown in the figure: for other buildings, floors, cities, and perhaps a conference room label for even more granularity. The noted AP here can be associated with four labels: conference rooms, city ABC, building 1, and BLD-1-FLR-1.



Knowledge Check

Self-check on key learning points



Its time for a knowledge check.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Question #1

Choose the statements below that accurately describe Aruba Central device onboarding. (Select three)

- a. Zero Touch Provisioning (ZTP) is an onboarding method that relies on the cloud-based Aruba Activate service.
- b. Manually adding devices to Aruba Central is often considered the easiest method, because it does not rely on other services.
- c. You first onboard devices to HPE GreenLake, and then assign them to the Aruba Central app.
- d. Before a customer can configure and monitor devices in Aruba Central, those devices must be onboarded.

a

Knowledge Check

MOD 7- 34

A,C,D

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Question #2

Your manager asks you about Aruba Central licensing. Which of the options below would help your manager understand this aspect of your deployment? (Select five)

- a. Foundation keys are available for gateways, APs, and switches.
- b. Licenses are available in 2-year increments.
- c. You need one license per managed device.
- d. You can purchase additional Advanced licenses to add functionality and features.
- e. You must add subscription keys to Aruba Central, and then assign them to network devices.
- f. The auto-subscribe feature allows you to assign licenses to all new devices automatically.

a

Knowledge Check

MOD 7- 35

A,C,D,E,F

Question #3

Which of the options below offer true comparisons of Aruba Central Groups vs Aruba Central Sites? (Select three)

- a. Both can help ease monitor, filter, and report functions.
- b. Devices can only belong to one site, but can belong to multiple groups.
- c. Groups come in two types: a UI group and a template group.
- d. Only groups include a device persona

a

Knowledge Check


MOD 7- 36

A,C,D

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Question #4

Match the task involved in group management with the order in which you should perform that task.



Knowledge Check ✓



MOD 7- 37

True

Question #5

Labels help you filter devices for Aruba Central monitor, manage, and report purposes. You can assign multiple labels to a single device.

- True
- False

a

Knowledge Check

MOD 7- 38

True



Lab Activity

Consider the lab activity.

CONFIDENTIAL - FOR TRAINING

Lab Tasks

Lab 7.01
Managing Aggregation Switches with Aruba Central

Lab 7.02
Managing Edge Switches with Aruba Central

**Module 8 – AP
Onboarding and
Device Profiles**



MOD 7- 40

The figure provides a brief review of lab tasks. Please see your lab guide for details. When complete, you will continue with Module 8 – AP Onboarding and Device Profiles.


Rev.22.41

aruba
a Hewlett Packard
Enterprise company

Device Profiles & AP Onboarding

Aruba Campus Access
Fundamentals

Hello and welcome back. This is Module 8: AP Onboarding & Device Profiles.



Objectives

- Use device profiling
- Describe LLDP and MAC profiling
- Connect AP to Aruba Central
- Perform initial AP setup

To avail yourself of Aruba Central management features, you must, of course, onboard your devices. You'll find device profiling to be a vital part of your Aruba switch operations.

After completing this module, you will be able to use device profiling to help simplify and secure the deployment. This includes both MAC and LLDP-based profiling. You will also be able to connect your APs to Aruba Central and perform the initial AP setup.



Device Profiling

We begin with device profiling.

CONFIDENTIAL - FOR TRAINING

Profiling Introduction

Network adjusts settings dynamically, based on device type



MOD 8- 4


Profiling and endpoint classification allows the network infrastructure to adjust settings and security policies dynamically, based on device type/model. For example, when an Aruba infrastructure device detects a wireless or wired VoIP phone, it can optimize the connection with the correct quality of service (QoS) settings for voice traffic. Aruba-based networks can also automatically adapt to various IoT device types.

Assigning roles based on identity might work if the company manages all the devices, placing them in correct groups. But it can be simpler and more reliable to let the Aruba solution automatically categorize devices and assign roles based on those categories.

And here's another use case for profiling: sometimes identity-based authentication does not provide the granular control you need. This is because users often connect to the network on multiple devices. If you only use user's identity to assign roles, the same user always gets the same rights, regardless of device type.

But smartphones tend to be less secure against some exploits and employees tend to use them for less demanding work applications, like browsing the internet or checking email. A more limited access profile makes sense when the user connects using a smartphone versus their company laptop. Therefore, by combining identity-based and endpoint category-based criteria, a user gets unlimited access from their laptop, but only internet access from their smartphone.

Device Provisioning



MOD 8- 5

When you enable profiling for device provisioning, the network devices can dynamically assign roles, like Aruba switches, APs, and gateways, can dynamically assign roles to the user devices. This section discusses the components of device provisioning:

- LLDP, CDP, and/or MAC address groups
- Roles
- Device profiles

A role is a collection or group of settings that can be assigned to a switch port: a VLAN, PoE priority, QoS trust mode (DSCP and 802.1 prioritizations), security policy, bandwidth restrictions, custom captive (web) portal page, etc. Devices can now be plug-n-play with no specific port configuration on an Aruba switch. For example, the switch, instead, learns about the device that is connecting and dynamically assigns the role to associated switch port.

You will learn more about roles later in this training. For now, let's explore provisioning operation.

Configuring roles

```
switch(config)# port-access role AP
switch(config-pa-role)# ?
<omitted for clarity>
description          Description 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.
private-vlan           Configure Private-VLAN.
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.
```




MOD 8-6

The benefits of assigning roles to users and/or devices are well known within the Aruba wireless world. Roles have simply been adapted from the Aruba wireless products to the Aruba switching portfolio. Just as with Aruba wireless, roles simplify configuration your configuration. Roles group access and security policies in a single configuration that can be referenced by many devices and/or user types. Look at the role configuration example in the figure.

Notice that you have many options for polices that will be associated to the connected user/device. The most common setting is the VLAN identifier that will be assigned to the user.

LLDP vs MAC-based


LLDP based



Sys-Desc : ArubaOS (MODEL: 305)
SysName : AP-1
ChassisId : fc:7f:f1:ce:95:44
...

LLDP message

MAC based



OUI → **fc:7f:f1**:ce:95:44

MOD 8- 7

Recall that roles enable predefined configurations to be applied to the user's access, based on user's identity as well as their connected device type. But when you connect a device to an Aruba switch, how does it determine the device type? The switch supports LLDP, CDP, and local MAC matching in order to profile the device. Other Aruba devices might use other protocol information to assist in the profiling process. For example, Aruba gateways can examine user DHCP and HTTP traffic to assist in profiling process, and ClearPass can look at over a dozen protocols and mechanisms to assist with the profiling process.

LLDP or CDP messages tell the switch what type of device it is. LLDP is an open-stand protocol while CDP (Cisco Discovery Protocol) is typically used by Cisco devices: both protocols share information, via layer 2, about the device's capabilities. LLDP and CDP messages contain information about the device product, model, software or firmware version, and many other items. The switch can also glean information by focusing on the connected device's MAC address: the first three bytes of every MAC address is the Organizationally Unique Identifier (OUI), which is the manufacturer of the device NIC.

Connected devices are identified using corresponding protocol packets. When the protocol information on the interface ages, the profile or role is revoked from the interface. When a device of a configured type is connected to an interface, the switch automatically applies the corresponding role.

If the switch enforces 802.1X authentication port security, Aruba APs support 802.1X in supplicant mode so they can function as 802.1X authentication clients just like PCs. The ⁴⁵¹

authentication server can also pass back the appropriate role that is to be applied. And if the authentication server is ClearPass, ClearPass has extensive profiling capabilities that can be leveraged in determining the role that it should pass to the switch, which, in turn, the switch will apply to the user's device's access.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

LLDP Device Group

```
Switch(config)# port-access lldp-group ArubaAP-LLDP
Switch(config-lldp-group)# match
    sys-desc    Configure LLDP system description.
    sysname     Configure LLDP system name.
    vendor-oui  Configure LLDP vendor OUI.
Switch(config-lldp-group)# ignore sysname AP10
Switch(config-lldp-group)# match sys-desc ArubaOS
```

```
Switch(config-lldp-group)# show running-config current-context
port-access lldp-group ArubaAP-LLDP
    seq 10 ignore sysname AP10
    seq 20 match sys-desc ArubaOS
```



MOD 8- 8

LLDP device groups are typically used when you are profiling capabilities, but are not using an authentication server, like ClearPass, to do the profiling and to dynamically pass back a respective user role. The figure shows how to create a device group using LLDP. You give the group a name, and then specify the group's matching criteria, which is used to determine group membership. LLDP device groups allow you to match on the criteria highlighted in the figure. Multiple TLVs can be matched or ignored simultaneously; use seq numbers to control the order of matching rules.

MAC Device Group

```
Switch(config)# mac-group ArubaAP-MAC
Switch(config-mac-group)# match
    mac          Configure MAC address
    mac-mask     Configure MAC mask
    mac-oui      Configure MAC OUI
Switch(config-mac-group)# match mac-oui f0:5c:19
Switch(config-mac-group)# match mac 20:4c:03:d4:3b:e3

Switch(config-mac-group)# show running-config current-context
mac-group ArubaAP-MAC
    seq 10 match mac-oui f0:5c:19
    seq 20 match mac 20:4c:03:d4:3b:e3
```

MAC based profiling requires MAC authentication enabled on a port

MOD 8-9

With MAC device groups, you can match or ignore specific MAC addresses as well as part of the MAC address, like the MAC OUI (leading 24 bits) or a range of MAC addresses.

You have enabled the switch to determine group membership. Now, what does the switch do with that information? You must configure a device profile to ensure that switch ports are properly configured for group members.

Device Profile

LLDP group

- ignore sysname AP10
- match sys-desc ArubaOS


MAC group

- match mac-oui f0:5c:19
- match mac 20:4c:03:d4:3b:e3

Role AP

- vlan trunk 10,20
- vlan trunk native 10
- poe-priority high

```
Switch(config)# port-access device-profile Access-Point
Switch(config-device-profile)# enable
Switch(config-device-profile)# associate mac-group ArubaAP-MAC
Switch(config-device-profile)# associate lldp-group ArubaAP-LLDP
Switch(config-device-profile)# associate role AP
```



MOD 8-10

The device profile associates a device group and the user role that should be applied to devices that match. Recall that the role assigns the user's access policies, like the VLAN and / or security policy applied to the user's access.

In the figure, a device profile named Access-Point has been created and is associated with the two groups you just created with the role named AP. Therefore, the switch knows which role to assign to a port when a connected device matches a device group. If you configure a match-criteria that matches across multiple device groups, then the priority considered is LLDP, and then local MAC matching. That is, LLDP has precedence over local MAC matching.

Profiling Verification

```
Switch(config)# show port-access device-profile interface all
Port 1/1/2, Neighbor-Mac 20:4c:03:d4:3b:e3
  Profile Name:           : Access-Point
  LLDP Group:             :
  CDP Group:              :
  MAC Group:              : ArubaAP-MAC
  Role:                   : AP
  State:                  : applied
  Failure Reason:         :
```



MOD 8- 11

Use the show port-access device-profile interface all command to verify profiling.

Connecting APs to Aruba Central



Let's move on to connecting APs to Aruba Central.

CONFIDENTIAL - FOR TRAINING

AOS 8 to AOS10 Comparison

AOS 8	AOS 10
IAPs, CAP, RAP	AP, Microbranch AP
Mobility Conductor	No Mobility conductor
Mobility Controllers	Gateways
SD-Branch Gateways	SD-Branch & Mobility
Local or Cloud Mgmt	Cloud Management



• MOD 8-13

Aruba AOS 10 simplifies access point deployment methods with just two AP options: AP and microbranch AP. AOS10 does not include an IAP cluster or IAP virtual controller concept. Additionally, AOS10 does not include a mobility conductor (formerly known as mobility master) since all those features are now in Aruba Central: basically, you get a simpler network architecture utilizing AOS 10.

Note: In AOS 10, Aruba has renamed controllers to gateways, and merged the firmware code for SD-Branch and mobility implementations.

AP to Aruba Central Communication



MOD 8- 14

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

You will learn how to manage APs and groups during the lab activities.

AP System Configuration


The screenshot displays the Aruba Central web interface for configuring Access Points. The left navigation pane shows the 'campus-wifi-ui' context selected under the 'Devices' section. The main content area shows the 'System' tab for 'Access Points'. The configuration fields include:

- Set Country code for group : US - United States
- Timezone : Eastern-Time UTC-05
- The selected country observes Daylight Savings Time
- NTP Server : 10.254.1.21
- DHCP Option 82 XML :
- Login Session Timeout : 60
- Console Access : ☒
- LED Display : ☒

MOD 8- 15

In Aruba Central, select the appropriate context in the left navigation pane. Then navigate to Devices > Access Points. Click the System tab to see essential system configurations: AP country code, time zone, preferred band, NTP server, and much more.

Creating Floorplan



View floor plans associated with APs, clients, rogues, buildings, floors

Locate associated clients and rogue devices


MOD 8- 16

In Aruba Central, select the appropriate context in the left navigation pane. Then select Overview in the left navigation pane. Click the Floor Plan tab to plan sites, create and manage floor plans, and provision access points. You'll need to import a picture of a floorplan, which can be something as simple as a GIF or BMP file, for example, or as extensive as an Auto-Cad file. In the lab you will see the wizard-based process that is used to create a floorplan.

You can only access the floor plan dashboard from a site context. Here's what you can do as it relates to using floorplans:

- Pictorial navigation: view floor plans associated with APs, associated clients, rogues, buildings, and floors.
- Accurately locate associated client devices using RF data from those devices.
- Accurately locate rogue devices (as classified by RAPIDS) using RF data from your devices.

Adding AP To Floorplan



MOD 8-17

The floorplan shows a real-time picture of the radio environment, as well as allowing you to plan the wireless coverage for new sites (referred to as a greenfield installation). To really understand your wireless network, you must know the location of devices and users and the health of your RF environment. With integrated mapping and location data, there's a wealth of information at your fingertips.

The floorplan feature uses sophisticated RF fingerprinting to accurately display coverage patterns and calculate the location of every wireless device in range. There is no need for dedicated RF sensors or a costly additional location appliance: Aruba Central gathers all required information from your existing APs. First, though, you must place your APs on the floorplan based on where they are physically installed. This can be done when you initially create a floorplan or after the fact.



Knowledge Check

Self-check on key learning points



Its time for a knowledge check.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Question #1

What statements below are true about device profiling? (Select three)

- a. Profiling and endpoint classification allows the network infrastructure to adjust settings and security policies dynamically, based on device type.
- b. You can use profiling with device provisioning, so when a device connects, switches are automatically configured with the correct VLAN, PoE priority, security policies, and more.
- c. End systems can only be profiled based on LLDP information.
- d. A switch knows which role to assign to a port when a connected device matches a device group, which you define.

a

Knowledge Check

MOD 8-19

A,B,D



Lab Activity

Consider the lab activity.

CONFIDENTIAL - FOR TRAINING

Lab Tasks

Lab 8.01: Onboarding APs

Task 8.1

- Verify AP wired onboarding

Task 8.2

- Prepare Aruba Central Groups

Task 8.3

- Move APs to the campus-wifi-ui group

Task 8.4

- Configure the Site Floorplan with APs

Module 9 – WLAN Fundamentals



MOD 8- 21

The figure provides a brief review of lab tasks. Please see your lab guide for details. When complete, you will continue with Module 9 – WLAN Fundamentals.


Rev.22.41

aruba
a Hewlett Packard
Enterprise company

WLAN Fundamentals

Aruba Campus Access
Fundamentals

Hello and welcome back. This is Module 9: WLAN fundamentals.



Objectives

- Explain RF bands & channels
- Describe RF basics and WLAN mobility
- Explain 802.11 frames
- Describe aspects of RF design
- Configure WLANs

Reliable, high-performance WLANs are key to successful deployments. After completing this module, you will be able to explain radio frequency (RF) bands and channels and describe RF basics and WLAN mobility. You will be able to explain 801.11 frames and describe key aspects of good RF design.


Finally, you will be able to configure WLANs.

RF Bands & Channels

We begin with a discussion of RF bands and channels.

CONFIDENTIAL - FOR TRAINING

WLAN Organizations



MOD 9- 4

Several organizations are responsible for wireless local area network (WLAN) standards. These regulatory bodies place restrictions on RF technology and how you use it in a WLAN environment.

The IETF is an international Internet standards organization, responsible for many internet standards, as documented in so-called Request for Comments (RFCs). There are several Wi-Fi standards related to WLAN control and provisioning.

Each area of the world is governed by local RF regulations. The FCC regulates RF spectrum in the United States, ETSI regulates Europe, and so on. These bodies dictate RF power limits and usable WLAN frequencies. These organizations are largely concerned with minimizing RF interference.


Within these limitations, the IEEE created the 802.11 standards. 802.11 channels must fall within FCC unlicensed ranges. Therefore, the FCC regulates power and channels, and IEEE standardizes how to use those channels: layer 1 physical layer transmission methods and layer 2 frame types and protocols.

The Wi-Fi alliance is just that – an alliance of WLAN hardware and software manufacturers like Aruba. They provide interoperability standards to assist in ensuring different vendors' APs and clients can work together. This often centers on security and quality of service (QoS) mechanisms.

For now, let's focus on those IEEE standards.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

IEEE 802.11 Standards



MOD 9-5

Early IEEE 802.11 standards like the original 802.11, 802.11b, a, and g standards were great when they were introduced, with data rates up to 54Mbps, and support in both the 2.4 and 5 GHz frequency bands. However, 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 being a typical speed. It accomplished these speed by increasing the channel width, using the advanced 64-Quadrature Amplitude Modulation (QAM), and techniques like multiple input/multiple output (MIMO).

802.11ac only operates in the 5 GHz range with theoretical data rates up to 6.93 Gbps, using very high throughput (VHT) methods, with higher channel width and more complex modulation techniques. 802.11ac is sometimes referred to as Wi-Fi 5.


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

The 802.11ax extension was introduced in 2020, with up to 7.8Gbps, 1024-QAM with modulation and coding scheme (MCS) 10/11. There were two important features introduced with this standard: it operates in the 6 GHz RF Spectrum and supports 1 Gigabit channel width. This version of 802.11ax has been referred to as Wi-Fi 6E.

Let's look at the Wi-Fi Alliance standards.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

2.4 GHz Band and Channels



Signal spread over 20+Mhz
Adjacent channels are not usable in same physical RF area




MOD 9-6

One of the things that Wi-Fi Alliance defines is how to use RF for Wi-Fi. 802.11 channels in the 2.4 GHz spectrum are 20 MHz wide. Clients and APs use this frequency range to transmit and receive data.

The 2.4 GHz band allows for three non-overlapping channels in some regulatory domains such as the United States. Other regulatory domains, such as Japan, allow the use of channel 14, providing four non-overlapping channels.

A WLAN signal uses spread spectrum technology so adjacent channels overlap. When an AP broadcasting on Channel 1 can hear an AP broadcasting on Channel 2, both radios will experience adjacent channel interference or ACI. ACI is like talking over top of another person when having a conversation. Some or all the transmission may be corrupted due to channel overlap. Due to the nature of channel overlap, only the use of channels 1, 6, 11 is recommended. Depending on regulatory domain, channel 14 may also be considered for use.


5 GHz U-NII Band and Channels



802.11a, n, ac, and ax all make use of the 5 GHz spectrum. Given that 5 GHz is a higher frequency, higher data rates can be achieved. Additionally, the presence of more channels allows for increased channel bandwidth. 5 GHz allows for the use of 24 non-overlapping channels.

Highlighted in orange, you will notice at lower power levels, 5 GHz channels do overlap. However, as power increases, so does beamwidth, eliminating adjacent channel interference (ACI) in the 5 GHz spectrum. Design considerations concerning low powered APs in 5 GHz are beyond the scope of this course.


Channel Bonding



As you learned, 802.11 channels are 20 MHz wide, and devices use this frequency range to transfer data. Channel Bonding is the process of combining the bandwidth of two or more adjacent channels: for example, channel 149 bonded with channel 153 becomes channel 149+.

When channels are bonded, devices on this now wider channel have twice as many frequencies to modulate their signal across, allowing for increased throughput. Channel bonding comes with a cost however: for every bonded channel, the noise floor of the channel will increase by 3dB (decibel). As we will learn, a 3dB increase means double in the decibel calculations.

5GHz Allowed Channels in FCC Countries




Bonding channels can aid in providing more bandwidth to a channel but becomes a point of diminishing return as fewer channels are available and the noise floor increases. In most enterprise WLANs, 20 or 40 MHz channel widths are optimal. 80 and 160 MHz channels are less practical due to lack of available bandwidth and increased noise floor. Notice the difference in available channels as channel bandwidth increases.

Also, IEEE's 802.11h dynamic frequency selection (DFS) places restrictions on which channels you can use in the 5 GHz range when co-located with satellite and/or radar signals.

6 GHz – Massive Increase in Capacity


- More than 2x the existing amount of unlicensed spectrum
- Up to seven 160 MHz wide channels
- 6 GHz is greenfield - no need for backwards compatibility



Today, spatial re-use of wide channels like 80s and 160s is not realistic. This is because we are limited with only about 500 MHz of frequency allocation, depending on the country. With 6 GHz, we get a massive increase, where some countries are getting more than two times the existing amount of capacity.

Notice in the figure that some of the 6 GHz channels in the graphic have a different shade of color. You will learn about these special channels later in the module.

Usable 6 GHz Channels in Americas & Europe/CEPT



Throughout the Americas, regulators are opening the entire 1,200 MHz for unlicensed use. To put this another way, we have more 80 MHz channels in this model than there are 40 MHz in the current 5 GHz band. This means we will default to 80 or 160MHz-wide channels in the 6 GHz band. For the first time, it will be practical to have wide channels and spatial re-use for enterprise deployments.

In other regulatory domains, or the European model, only the first 500 MHz portion shall be opened for unlicensed use. With that said, the global organization, Dynamic Spectrum Alliance, published an article in 2022 urging regulators to consider opening the entire band because the importance of permitting unlicensed use of the full 1,200 MHz, since this is critical for spatial re-use.

Initial product focus is on the low power indoor (LPI) device class where access points are limited to indoor operation only. LPI has access to all 1,200 MHz of the frequency, has no DFS requirement, has no spectrum database requirement, and has a massive amount of additional spectrum and capacity.

Note: Technically, channel 233 would be the 59th channel, and you may find references to it in other places, but power will be very low and likely realistically unusable due to band edge limitations.




RF Basics

To design, deployment, and troubleshooting a wireless solution, you really should know the basics of Radio Frequency (RF) energy.

CONFIDENTIAL - FOR TRAINING

Milliwatts and dBm



RF Power is a logarithmic calculation

Distance(m)	Received Signal Strength	
	mW	dBm
0	1	0
0	100	20
1	.0098911	-20
10	.0000989	-40
20	.0000247	-46
100	.0000010	-60
1000 (1km)	.0000000099	-80

dBm is much easier to work with than mW



MOD 9- 13

One scale used to measure RF Power is milliwatts.

This chart shows a signal sent by an AP, decreasing in power as it reaches the client. The AP transmit power is 100 milliwatts (mW), but that energy is spread out over the AP's entire coverage area. The portion of that signal to hit the client's antenna is quite small. In fact, most of this power is lost within the first linear meter. The table at the right shows how drastically the power drops as client-to-AP distance increases.

These tiny numbers get cumbersome to speak out loud (too many zeros). This is because milliwatts is a linear scale, used to measure RF energy, which is a logarithmic phenomenon.

Thankfully, we have another scale to use, decibels as compared to 1mW (dBm). So, 1mw = 0dBm, and 100mW = 20dBm. So, instead of saying, "my signal strength is .00000001 mW", you can simply say "-80dBm.


Let's get some perspective on these numbers. How do you use them in the real world?

dBm and mW Calculations

0 dBm = 1 mW

+3 dBm = 2X power / -3 dBm = 1/2 power

+10 dBm = 10X power / -10 dBm = 1/10th power



dBm	mW
+20	100
+17	50
+14	25
+10	10
0	1
-10	.1
-30	.001
-50	.000001
-70	.0000001
-90	.000000001


MOD 9- 14

When deploying a WLAN, it is key to ensure that clients will always receive a good signal, otherwise the connection may be poor, which leads to low data rates and disconnects. The amount of signal power the client receives is its known as received signal strength indicator (RSSI). An RSSI value of -62 dBm or so is great for most users, while receiving a higher value is considered an excellent signal. However, a lower power can lead to underperformance. Many wireless administrators use -75 dBm as a minimum value to properly communicate. Anything lower than that is largely unusable.

Remember, dBm is decibels relative to one milliwatt: 0 dBm equals one milliwatt. When dealing decibel representations, you must know the rule of tens and threes. A 3 dBm increase is equal to double the power. A 10 dBm increase is equal to 10 times the power. Likewise, a 3dB decrease is half power, and a 10 dB decrease is one tenth of the power.

Calculating Effective Isotropic Radiated Power


$$20\text{dBm} - 3\text{dB cable and connector loss} + 10\text{dBi antenna gain} = 27\text{dBm EIRP}$$



MOD 9-15

Effective or equivalent isotropic radiated power (EIRP) simply means how much power the antenna effectively radiates out into free space. So perhaps the radio generates a 100mW (20dBm) signal that is transmitted into the external cable, which, in turn, is connected to an external antenna. In this example, it is a long cable, and therefore there is 3dB of loss as the power travels through the wire. Also, this particular antenna has a fairly squashed pattern and therefore has a 10dBi of gain. So, $20\text{dBm} - 3\text{dB cable loss} = 17$, plus the 10dBi gain of the antenna resulting in a 27dBm EIRP. As you can see in this example, the antenna provides amplification for the signal, and thus the gain in dBm.

Distance, Performance, and MCS




MOD 9- 16

RF Transmit power, or EIRP, determines the AP's RF coverage area. The transmit power also helps the client to choose the best modulation and coding scheme (MCS) for its transmission. A modulation scheme is a way of changing (modulating) a transmitted waveform to represent data: binary 1's and 0's.


As the client-to-AP distance increases, the signal weakens, waveforms degrade, and the error rates increase. MCS is a key factor in determining data rate and therefore throughput. The more complex the MCS, the higher the data rate. However, as client-to-AP distance increases, the signal weakens, thus distorting these very complex waveforms. Error rates increase. And so, the client must automatically drop down to a lower MCS. The more clearly a client can hear an AP the more complexity it can add to its modulation and encoding of data.

MCS vs Distance.



The figure relates possible MCS encoding schemes versus client-to-AP distance. If a client is very far away, it may have to drop down to MCS 0, which is a very simple modulation scheme called binary phase shift keying (BPSK) and supports only low data rates. As a client gets closer to the AP, it can auto-rate-shift up to a more complex MCS 9 encoding scheme called 256-quadrature amplitude modulation (256-QAM): a far more intricate scheme that supports much higher data rates.

Minimum Signal Strength Requirements by Data Rate




The graph in the figure shows what signal strengths are required to achieve a certain MCS. You can see how higher MCS rates such as 256-QAM require high signal strengths and lower noise floors. As the channel bandwidth is increased, so is the noise floor. Notice how the signal strength requirements to achieve a specific MCS increase with channel bandwidth.

MCS Solutions include:

- Quadrature Amplitude Modulation (QAM)
- Binary Phase Shift Keying (BPSK)
- Quadrature Phase Shift Keying (QPSK)


Signals and Noises



Imagine you with a friend at a rock concert. The music is playing, and you are right up front. You yell at your friend, "These guys rock!". Your friend says, "What?" "I said, these guys rock!" Your friend says, "What?" Never mind.

You were yelling to your friend, standing right next to you: the signal was very loud. But the noise in the room was also loud. There was very little difference between the signal (your loud voice) and the noise. This is an important RF concept.

Signal to Noise Ratio (SNR)




MOD 9- 20

The ambient noise or noise floor is the sum of all undesired transmissions from devices that are either too far away for the signal to be intelligible or nearby electronic devices that create interference. The signal-to-noise ratio (SNR) is the difference between the ambient noise floor and the received signal strength. You express it as a difference, not a ratio.


In this example, the noise floor measures -90dBm. The receiver's signal strength is -65dBm. Therefore, $-90 \text{ minus } -65$ give us an SNR of 25dBm. The higher your SNR, the better your Modulation Coding Scheme (MCS) will be, and the higher the communication data rate. Clients often require a minimum SNR of 15 dBm. However, SNRs in excess of 30dBm provide the best data rates. Most people agree that you want at least 25 dBm or better for the SNR.

Antenna Types

Omni Directional




Directional



MOD 9- 21

Aruba offers APs in two main types: integrated antennas or external antennas. External antenna APs can be customized to provide extra passive gain, increased receive sensitivity, and provide directionality for an AP's RF transmission. Most integrated antennas are omni-directional, meaning they will transmit approximately equal power in all directions. This class will focus on omni-directional radiation patterns.


Horizontal Plane (azimuth)



MOD 9- 22

AP radiation patterns show the horizontal and vertical coverage areas. H-plane or horizontal view, also known as the azimuth shows the radiated pattern coming from the sides of the Access Point. If you have an AP, place it on the desk, white dome facing up. Draw a circle around the AP. That is your horizontal plane.

Vertical Plane (Elevation)




MOD 9- 23

The E-Plane or elevation plane, also known as the vertical plane, shows radiated patterns coming from the front of the AP. Many directional antennas sacrifice horizontal gain in favor of vertical gain to achieve connectivity at longer distances.

Passive E-Plane Antenna Gain


Side view – vertical (e-plane) coverage



Passive gain, omni-directional antennas, however, sacrifice vertical coverage in favor of horizontal coverage. Passive gain antennas can be used when applying coverage to a single floor instead of through a floor. In the figure you see the side view of the affect passive gain has on vertical coverage.

Passive H-Plane Antenna Gain

Top view – horizontal (h-plane) coverage




MOD 9- 25


In the figure, look at the top view of the passive gain omni-directional antenna and how it will provide more horizontal coverage.

AP/Antenna Mounting Options

Ceiling Mount



Side Mount



Floor Mount



MOD 9- 26


APs can be mounted in one of three orientations. The placement of your APs will vary on a situational basis. A typical drop ceiling is as close to a perfect mounting place as you can get. Cables are usually easy to run above a drop ceiling and the porous nature of the material provides attenuation of the signal coming from the back of the AP.

If a ceiling mount is not available, a side mount can be considered. Reasons to use a side mount include the ceiling being too high, forcing the AP to be too far from clients. Also, aesthetics play a key factor in AP placement as well as cable location. If the ceiling is not available due to height, cabling restrictions or aesthetics, consider a side mount on walls or columns. Additionally, special mounts can be used to orientate the AP like a ceiling mount while not attaching to the ceiling.

In very high-density deployments, a floor mount may be considered. Mounting APs under the floor or seats in a stadium for example, creates pico cells. Pico cells are very small service areas by design.

Note: Most Aruba APs are meant to be mounted on the ceiling, with the exceptions of the desktop and hospitality series, which are meant to be side-mounted. If you need to mount an AP, like a hospitality AP on the ceiling, you will negatively affect its RF coverage. Therefore, you need to carefully consider where you will be placing your APs to determine the correct model and whether you can use internal versus external antennas.

SISO: Legacy 802.11a,b,g Devices



MOD 9- 27

This section introduces you to how APs send and clients receive RF transmissions.

Legacy 802.11 networks use single input single output (SISO) radio technologies where only one antenna is used for transmitting or receiving at a time. SISO does not make full use of a radio with multiple antennas. The transmitter sends the signal out of a single antenna. While the receiver will hear the transmission on both antennas, it will only use one of those signals.

Remember, SISO is largely a legacy operating procedure. The benefits of beamforming and MIMO will be covered shortly.

MIMO - Multiple Input Multiple Output

Both antennas used for transmitting signals




Both antennas receive multiple signals



MOD 9- 28

802.11a and 802.11ac use MIMO antenna technologies to create more than one stream of data transmitted out of different antennas at the same time. These streams of data are known as spatial streams and provide significantly faster throughput compared to a single spatial stream technology.

Multiuser - Multiple Input Multiple Output (MU-MIMO)




MOD 9- 29

Multi-user MIMO (MU-MIMO) uses beamforming technology to send traffic to multiple clients in the same transmission. The AP will use a channel sounding procedure to which the clients will respond. The AP uses this information to build a steering matrix. The AP uses the steering matrix to manipulate its transmission, coordinating in phase spatial streams for each client when the frames intended for them egress the radio.

Due to the nature of MU-MIMO, multi-user clients physically near each other do not receive the benefits of beamforming as their transmissions will interfere with one another.

Beamforming and MIMO




MOD 9- 30

Beamforming technology allows an AP to locate client's antennas and modulate its spatial streams in phase. When spatial streams are modulated in phase, two or more copies of the same signal are generated from the AP's radio. The AP times the transmission from each antenna, so the spatial streams reach the receiving antennas at the same time, in phase, increasing the received signal strength at the radio.

Remember, signal strength and noise floor are key factors in determining data rate. Beamforming allows for increased signal strength at the client radio, without adding more power from the AP, which can increase the noise floor.

When deciding on an AP model, supported spatial streams may be a key factor: different products have different capabilities. The figure shows an example of the Aruba 515 AP. Note the difference between MU spatial streams and MU stations. Multi-User stations is number of clients an AP can serve in one down stream with the MU-MIMO transmission.

Total Spatial Streams



AP-515 **4** x **4** : **4** : 4 : 3

4 signals at the same time
to **single** client



4 x 4 Client




MOD 9- 31

The first parameter (e.g. 4x4) indicates number of transmit and receive chains (typically antennas). The second number describes maximum number of signals that can be communicated to a single client at the same time.

Multi User Spatial Streams

AP-515 4 x 4 : 4 : **4** : 3



4 signals at the same time
to **multiple** clients

Q


MOD 9- 32

Next we have a maximum number of Multi-User Spatial Streams. This indicates how many signals can be exchanged with multiple users at the same time.

Supported Multi User Stations

AP-515 4 x 4 : 4 : 4 : 3

Maximum 3 MU clients



MOD 9- 33

Finally, last number defines maximum number of multi-user clients that AP can communicate with simultaneously.


WLAN Mobility



One of the many great advantages of Wi-Fi is end user mobility.

CONFIDENTIAL - FOR TRAINING

WLAN Logical Configuration - BSS



There are two types of configuration involved with a WLAN:

- Basic Service Set (BSS)
- Extended Service Set (ESS)


A BSS includes an AP radio and all associated clients in a coverage area. An AP with two or more radios therefore has more than one BSS. When you define an SSID (the name of the WLAN), it is assigned a different 48-bit MAC address on each of the AP radios, which is used to identify the SSID and the radio it is advertised on. 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 on 802.11 frames, and the AP understands, based on the received radio and the BSSID MAC address, with which SSID the 802.11 frame is associated.

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

All BSSIDs are based on the AP's 2.4 , 5, and 6 GHz radio base MAC address. The first SSID you define on AP1 might end with aa:aa:aa, as shown in the figure. The next SSID you define on that AP might end with aa:aa:ab, the next ends with aa:aa:ac, and so on. The 5 GHz radio's SSIDs would be based on its MAC address, like bb:bb:bb for the first SSID, bb:bb:bc for the second MAC address, and so on.

WLAN Logical Configuration - ESS

Extended Service Set (ESS): One SSID Broadcast on Multiple Radios



In the figure you see an ESS architecture applied to multiple APs. Note that the WLAN advertises the same SSID across multiple APs. However, each AP radio will have its own unique BSSID MAC address for that SSID. This allows the APs to understand when traffic is sent to one of their own radios versus a different AP's radio.

Wireless Device Mobility

Stationary Device



Somewhat Mobile Devices



Highly Mobile Devices



MOD 9- 37

Different devices have different functions on the network as well as different roaming capabilities.

- Stationary devices, such as printers and smart TVs almost never move when in use on the WLAN, therefore, they are not expected to roam from one AP to another while in use.
- Somewhat mobile devices such as handheld scanners, laptops and some IoT devices may move when in use and therefore may have to roam from one AP to another.
- Highly mobile devices such as tablets, phones and some IoT devices are more likely to require smooth roaming as they are moving while accessing the network.

These client definitions are generic and subject to change based on the actual function of the device. For example, an iPad in a lobby that is bolted to a kiosk to register visitors is a stationary device, even though it is a tablet.

802.11 Client Roaming



SSID – Corp1
BSSID=aa:aa:aa:aa:aa:aa



SSID – Corp1
BSSID=bb:bb:bb:bb:bb:bb



SSID – Corp1
BSSID=cc:cc:cc:cc:cc:cc



MOD 9- 38

When an 802.11 client roams from one access point to another, it changes the destination MAC address on its 802.11 frames to the BSSID of the new radio. The SSID however, will remain the same. It is important to remember that clients ultimately make the decision when and where to roam.

802.11 Frames

What did the Wi-Fi information say to 802.11? I've been framed!!

This section introduces 802.11 frames and their types and how they are used in AP communications, like a client connecting (associating) to an AP.

Types of 802.11 Frames

Management frames

Used to manage wireless network

Control frames

Used to control access to the medium

Data frames

Used to carry client data



MOD 9-40


All 802.11 frames fall under one of the three types: management, control, or data.

- 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 user data.

802.11 Management Frames

Joining the AP

1. Beacons
2. Probe Request
3. Probe Reply
4. Authentication
5. Association



MOD 9- 41

The figure has an overview of how a client connects to an AP. Client WiFi devices can automatically scan the respective RF spectrum to see the available SSIDs. They do this by looking for beacons that APs broadcast. APs transmit a beacon every 100ms by default. These beacons contain basic information about the AP and the services offered. Among other things, the following information is included in each beacon:

- Timestamp: To help synchronize devices
- Beacon interval: the time between beacons
- A list of available SSIDs: You can control which supported SSIDs are advertised in beacons and which are not. You will typically allow your APs to broadcast guest SSIDs and disallow them from advertising corporate and employee SSIDs.
- Supported data rates: speed of the client and AP communication.

Beacons allow end stations to passively discover available SSIDs. When you view the list of available Wi-Fi networks on your smartphone or laptop, you are seeing the SSIDs that have been broadcast by APs in the area; this is perfect for guests to easily discover and access your guest network.

But what about those corporate SSIDs that are not broadcast in the beacon for all to hear?

Your employee devices are manually configured with the appropriate SSID, so those stations need not passively wait to hear a beacon. They actively transmit a probe request frame to connect to the employee network.

When you try connecting to the WLAN, your device will send a probe request to the AP. The AP will then reply with a probe response, confirming the network information learned from the beacon. At this time, negotiation of the authentication protocol is done, and the association can begin.

Note: It is best to limit the number of SSIDs advertised in beacon frames. Every SSID increases the size of the beacon frame, which then takes more air time from everything else. Too many APs transmitting too many SSIDs, every 100ms, can begin to impact performance. One or two SSIDs are perfect. Many wireless administrators agree that staying below four is a reasonable target in a corporate production environment.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

802.11 Association

Protocol	Summary
802.11 Beacon	FC=.....,SN= 0,FN= 0,BI=100,SSID=CorpWLAN,DS=1
802.11 Beacon	FC=.....,SN= 0,FN= 0,BI=100,SSID=GuestWLAN,DS=1
802.11 Beacon	FC=.....,SN= 0,FN= 0,BI=100,SSID=CorpWLAN,DS=1
802.11 Beacon	FC=.....,SN= 0,FN= 0,BI=100,SSID=GuestWLAN,DS=1
802.11 Probe Req	FC=.....,SN= 28,FN= 0,SSID=CorpWLAN
802.11 Probe Rsp	FC=.....,SN= 0,FN= 0,BI=100,SSID=CorpWLAN,DS=1
802.11 Ack	FC=.....
802.11 Auth	FC=.....,SN= 29,FN= 0,Algorithm=0 (Open System)
802.11 Auth	FC=.....,SN= 0,FN= 0,Algorithm=0 (Open System)
802.11 Ack	FC=.....
802.11 Assoc Req	FC=.....,SN= 30,FN= 0,Listen=10,SSID=CorpWLAN
802.11 Assoc Rsp	FC=.....,SN= 0,FN= 0,Status=0,AID=1
802.11 Ack	FC=.....

Discover
the SSID

802.11 Open
Authentication

802.11
Association

The Wireshark frame capture in the figure shows a WLAN client connecting to an SSID (CorpWLAN). The client first discovers the network with beacons and probes. After discovery is complete, 802.11 authentication and association automatically begin.

MOD 9- 42


802.11 Management Frames

Leaving the AP

De-authentication & disassociation

De-Authentication

Disassociation Frame




MOD 9-43

When a client leaves an AP, it notifies the AP with a disassociation or de-authentication frame. For example, clients may decide to leave an AP because they are too far away and cannot meet minimum data rate requirements.

APs can also initiate the disassociation and de-authentication frames to clients in order to move a client to a better AP or to disconnect the client from the WLAN, perhaps because of a security violation.

802.11 CSMA-CA Control Frames and Timers



An RF channel is a contention-based, collision avoidance medium. This means that all clients are competing for a chance to send their traffic. Let's discuss how the collision avoidance medium works with Wi-Fi control frames.

The first mechanism you must understand is the count down timer on each client. When the countdown reaches 0, the client has a Transmit Opportunity or TXOP. When a client receives a TXOP, it will alert the AP and all other clients in the cell with a clear-to-send (CTS) message. The CTS-to-self announcement containing a network allocation vector or NAV. The Duration-ID portion informs the other clients of how long the transmission will take. It is important that the other clients can hear this notification. If two or more clients attempt to use the medium at the same time a collision will occur and both clients must re-attempt the transmission again.

Immediately after sending CTS-to-self with NAV distribution, the client will send its data frames and receive an ACK from the AP. If any part of the transmission is corrupted due to high noise or collisions, the entire process must start over. When a frame is lost and must be resent, it is called a L2-Retry or retried frame.

WLAN Design



Let's review some WLAN design ideas.

CONFIDENTIAL - FOR TRAINING

Coverage vs. Capacity Design


Coverage designs:

- Fewer APs
- APs set to higher power
- Typical: Few clients
- Low Throughput needs



Capacity design:

- More APs
- APs set to lower power
- Typical: Many clients
- High aggregate throughput needs



MOD 9-46

WLAN designs are often more complex than simply asking “is it a coverage or capacity design?” For the purpose of this course, we will look at coverage and capacity in a WLAN from a high-level perspective.

A WLAN designed for coverage can be implemented in environments where a few clients are spread over a large area such as a shipping yard or marina. These APs will have higher transmit power levels and typically only a few clients associated to them.

Note: During this discussion, when referencing the AP, it means the radios within the AP.

A capacity design requires serving many clients that are close together. Auditoriums and stadiums are examples of areas that require a high-capacity design. High-capacity designs require that APs use lower power settings. Lower power settings means the APs cover a smaller footprint on the floorplan, therefore servicing fewer clients. Now instead of, say 90 clients sharing the bandwidth of 1 AP, perhaps only 20 or 30 clients share that bandwidth. Yes, you spend more money on APs for this high-density deployment. But using fewer APs than needed is a bit like getting a great deal on a pair of shoes that do not fit, and the pain everyone feels will be just as aggravating.


Also, when you squeeze APs closer together, you need to carefully plan the layout and power settings to avoid co-channel interference. Two APs on the same or even adjacent channels can interfere with each other, drastically reducing performance. This is all about proper placement of APs within the physical space, along with proper tuning of power and channels.

settings.

Many seasoned Wi-Fi design and deployment engineers have been successful by placing APs where their power settings are optimal at about 30-50% of maximum power. This, of course, can vary greatly based on several factors beyond the scope of this course. This power setting range ensures that AP power can be adjusted up or down to accommodate inevitable changes in the RF space.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Channel Utilization



MOD 9- 47

Channel availability is a challenge to wireless network performance. High channel utilization will lead to a lack of transmit opportunities for clients. A channel is considered utilized when a device is making a transmission on the channel (see figure). When APs and clients send traffic, the channel is considered in use.

Interference will also contribute to channel utilization. While interference on a channel is present, clients may consider the channel in use and decide not to send their traffic at what would have been a transmit opportunity. The client will then back off and try again.

RF WLAN Interferers

Sources of Interference

Video security cameras

Cordless phone

APs

Microwave

Frequency hopping device (Bluetooth)




MOD 9- 48

Video cameras, cordless phones, third party APs, microwave ovens, and Bluetooth devices are some of the most common sources of WLAN interference. As we will see, a spectrum analyzer is an effective tool that identifies and locates potential sources of interference.

Things That Negatively Affect RF Signal


Free Space Path Loss



Distance

Absorption

5 and 6GHz bands affected more than 2.4GHz



Signal traveling through an object



MOD 9- 49

One of many things that affects an RF signal as it travels is free space path loss (FSPL). In the left side of the figure note how this fresh signal from the AP weakens and broadens as it moves further from the AP. As the signal's power decreases, the Signal to Noise Ratio (SNR) decreases, and the data rate drops.

Another factor that negatively affects an RF signal is absorption (see right side of figure): signal power decreases as it moves through an object, like a wall. The denser the object, the more the power decreases.

Note that the higher frequency bands like 5 and 6 GHz are more affected by absorption than the lower frequency 2.4 GHz band. Relate this to an audio system. Someone in a car near you might be playing their music obnoxiously loud. Their windows are up, as are yours, but what do you hear? You hear the low frequency bass sounds, which are less attenuated by obstacles than the high-frequency treble sounds.

AirMatch Function

Aruba Central hosted RF optimization service


Generates an RF solution for APs

Dynamically adapts to the changing RF conditions

Past 24 hours RF info used for calculation

Channel & Power deployment only once a day

Minimizes channel coupling, channel width



MOD 9- 50



Aruba has two features their APs support that deal with RF signal issues:

AirMatch

ClientMatch

AirMatch provides automated RF optimization to improve a user's experience by providing a stable network that dynamically adapts to the changing RF conditions. AirMatch, which runs in Aruba Central, is a centralized RF planning and optimization service. As shown in the figure, Aruba Central receives telemetry data from APs for radio measurements, channel range, EIRP range, operational conditions, and local RF events like radar detection or high noise. AirMatch also minimizes channel coupling, where adjacent radios are assigned to the same channel.

Aruba Central then models the RF network as a whole and devises a channel and power plan for the RF environment. AirMatch consumes the RF information and generates an RF solution that specifies new channel, bandwidth, EIRP, and mode of operation per radio. Central forwards this information to the APs and the APs apply the change in their radio configuration and operation. AirMatch calculates an RF channel and power level solution based on the past 24 hours RF information collected for each AP. The channel and power deployment process happens once a day at 5 AM, by default; however, deploying the channel plan can be changed based on your environment.

AirMatch also provides these additional features:

Providing the neighboring AP list to key management services to implement other features.

Provide AP partitions to implement a live upgrade service.

You will learn how to configure and use AirMatch during the lab activity.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY


What is Client Match?

Aruba Patented Technology

Load Balancing

Band Steering and Balancing

Sticky Clients



MOD 9- 51

ClientMatch is an Aruba patented technology that provides client load balancing across radio bands and AP radios. ClientMatch helps to influence sticky clients to roam to the best AP.

WLAN Setup

Let's get those users connected!

Here you will learn about how you configure a WLAN on your Aruba APs.

CONFIDENTIAL FOR TRAINING

WLAN Wizard

Create a New Network

1 General 2 VLANs 3 Security 4 Access 5 Summary

Name (SSID): p58t02-guest

> Advanced Settings

General	VLANs	Security	Access	Summary
• Name • Advanced Settings	• Forwarding Mode	• Security settings	• Role Based • Network Based • Unrestricted	• Review



MOD 9- 53

The WLAN wizard guides you through the SSID creation process. There are five tabs in this wizard: General, VLANs, Security, Access, and Summary.

Under the General tab you can give the new SSID a name. Notice that the SSID name is the same as the one your APs will broadcast. Under the advanced settings you may fine-tune the new SSID.

Under the VLANs tab, you need to choose the traffic forwarding mode and the client VLAN assignment.

Note: This course covers the bridge forwarding mode. The professional level certification covers two other forwarding modes.

Under the Security tab, select the SSID security level:

- Enterprise security means the network will use 802.1X for user authentication
- Personal means you'll use a pre-shared key.
- Visitors uses captive portal to authenticate users.
- Open means that no user or device authentication required.

Under the Access tab, you will configure the firewall settings which includes three options:

- Role-based, where a different firewall role could be applied to each user.
- Network-based, where all the users connected to the SSID will receive the same firewall

- role.
- Unrestricted, where no enforcement will be applied.

Note: For simplicity, the same steps are used to create WLANs for AOS10 bridge, tunnel, and mixed mode SSIDs, and AOS8 IAPs WLANs, with just a few changes on the wizard and configurations.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Traffic Forwarding Mode – Bridge & Native VLAN


Create a New Network

1 General 2 VLANs 3 Security 4 Access 5 Summary

Traffic forwarding mode:

☒ Bridge ☐ Tunnel ☐ Mixed

Client VLAN Assignment: ☐ Static ☐ Dynamic ☒ Native VLAN



MOD 9- 54

When you select bridge for the forwarding mode, you have three client VLAN assignment options (How the client device acquires IP addressing information):

- Static
- Dynamic
- Native VLAN

The next three sections cover these options.

If you select a Native VLAN for the client VLAN assignment, the corporate network will provide the client with an IP address on the same subnet as the AP. In the figure, the corporate DHCP server provides the AP the 10.1.4.49 address. The user associates to the AP and receives the 10.1.4.57 address, which is in the same subnet as the standalone AOS10 AP. This option is only available for bridged mode WLANs.

Traffic Forwarding Mode – Bridge & Static VLAN

Traffic forwarding mode:

☒ Bridge ☐ Tunnel ☐ Mixed

Client VLAN Assignment: ☒ Static ☐ Dynamic ☐ Native VLAN

VLAN ID:

10.1.11.150

SSID Employee
VLAN 11

10.1.12.140

SSID Guest
VLAN 12



10.1.4.49

802.1Q Native : 4

Allowed: 4,11,12



10.1.4.1



MOD 9- 55

If you select Static for the client VLAN assignment, you can statically set which VLANs the users are assigned, based on their SSID. In this configuration, the AP's IP address is in its own subnet. The administrator configures the SSID with a specific VLAN(s). If the WLAN is in bridge mode, then the VLANs are defined on the local switch.

In the figure, the AP is in VLAN 4. The users associated with the Employee SSID are assigned to VLAN 11, and the users associated with the Guest SSID are assigned to VLAN 12. For this to work, the switch port connected to the AP must be configured as a trunk. The AP management VLAN is the trunk's native VLAN (untagged), and the client VLANs (11 and 12) are tagged on the trunk.


Traffic Forwarding Mode – Bridge & Dynamic VLAN

☒ Bridge
 ☐ Tunnel
 ☐ Mixed

Client VLAN Assignment:
 ☐ Static
 ☒ Dynamic
 ☐ Native VLAN

VLAN Assignment Rules

Default VLAN: 11
■ If Aruba-User-Vlan equals 16 assign bridge VLAN: 16 ■ If Client MAC Address contains 00:11:22 assign bridge VLAN: 17



If you select Dynamic for the client VLAN assignment, the VLAN is derived based on the client connection parameters, such as the client's user account, device profiling, location, and other attributes.

The administrator needs to set up the VLAN derivation rules on the AP. You can set a default VLAN if the server does not return a value. Dynamic VLAN assignment is available for bridged and tunneled mode WLANs.

In the figure, if the user associating with the SSID receives a RADIUS attribute Aruba-User-VLAN 16, the user will be placed in VLAN 16. If the associating user has a MAC address that contains the following characters: 00:11:22 it will be assigned to VLAN 17. If the associating user does not match any of the assignment rules it will be assigned to VLAN 11.

Security

Enterprise	• 802.1x • WPA3, WPA2, Dynamic WEP • External RADIUS or Cloud Auth
Personal	• PSK • WPA3, WPA2, Static WEP, MPSPK AES, MPSPK Local
Visitors	• External Captive Portal (IE: ClearPass) • Cloud Guest (Central CP) • Encryption optional
Open	• Enhanced Open • Open




MOD 9- 57

You can configure an AOS10 WLAN with the following security profile types:

- Enterprise: for enterprise security, with user and/or device authentication.
- Personal: for personal networks, with a pre-shared key.
- Visitors: for guest networks with captive portal authentication
- Open: for open networks, with no authentication.

Note: Keep in mind that WEP is supported for backward compatibility. It should be avoided on new implementations since it is insecure and impacts and/or disables many of the modern functionalities such as MIMO/MU-MIMO, faster data rates, and much more. In short, WEP has been deprecated from the standards.

Access



MOD 9- 58

There are three main types of WLAN access rules; unrestricted, network-based and role-based.

An Unrestricted-WLAN does not enforce a firewall policy, allowing any traffic. All devices associated with this SSID can pass traffic unrestricted by the firewall.

Network-based WLAN has rules that are imposed on all devices associated to this specific WLAN. For example, if there is a rule that blocks pings, then all devices on the SSID would be denied ping traffic.

The role-based WLAN assigns roles to various devices in the SSID. Roles are derived on a user and/or device basis. If you have not assigned a role, then the AP will use the default role for the WLAN. For example, you may have one user receive the employee role, but the other users receive the IT role. In this situation both users are in the same subnet but have different access rights.

Monitoring Clients

Client Name	Status	IP Address	VLAN	Connected To	AP Role	Gateway Role	SSID/...
 employee1	 Connected	10.3.91.252	91	AP515-f0	employee		employee-T
 employee1	 Connected	10.3.91.254	91	AP515-7a	employee	employee	employee-T
 employee2	 Connected	10.3.71.254	71	AP515-D2	employee		employee-B
 employee1	 Offline	10.3.91.253	91	AP515-D2	employee	employee	employee-T



MOD 9- 59

The figure shows a typical List View output when monitoring clients in Aruba Central. You can also see this in a summary view. Notice that, by default, there is no way to differentiate bridged and tunneled clients. However, you may add a new column to the table by clicking the three-dot icon on the right and selecting the check box for tunneled. You will learn how to monitor clients during lab activities.



Knowledge Check

Self-check on key learning points



Its time for a knowledge check.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Question #1

Choose the options below that accurately compare the Wi-Fi radio frequency (RF) bands. (Select two.)

- a. 2.4GHz RF energy is less affected by absorption, as is caused when signals flow through office walls and other obstructions.
- b. 5 GHz RF energy supports more channels, which means there are less problems with co-channel interference.
- c. The 6 GHz band is supported in the 802.11ac and 802.ax standards.
- d. In a typical corporate deployment, channel bonding is used to improve performance for all RF bands.

a

Knowledge Check

MOD 9- 61

A,B

Question #2

Which statements below accurately compare Aruba AirMatch vs Aruba ClientMatch? (Select three.)

- a. AirMatch automatically tunes AP settings to optimize your RF environment.
- b. ClientMatch helps balance the client load across APs.
- c. ClientMatch helps balance the client load across RF bands.
- d. ClientMatch minimizes channel coupling, where adjacent radios are assigned to the same channel.

a

Knowledge Check

MOD 9- 62

A,B,C

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Question #3

When defining a WLAN, you can statically configure the VLAN, or allow the VLAN to be derived based on the client MAC address or RADIUS response attributes.

- True
- False

a

Knowledge Check

MOD 9- 63

T

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Question #4

Which of the following security mechanisms are most suitable for employees connecting to an enterprise-class deployment? (Select three.)

- a. 802.1X
- b. WPA3
- c. PSK
- d. Captive Portal
- e. External RADIUS server authentication

a

Knowledge Check

MOD 9- 64

A,B,E



Lab Activity

Consider the lab activity.

CONFIDENTIAL - FOR TRAINING

Lab Tasks

Lab 09.01: WLAN Fundamentals

Task 9.1

- Review Radio Default Channel and Power

Task 9.2

- Configure Radio Profiles

Task 9.3

- Configure AirMatch

Task 9.4

- Configure WLAN Zones

Module 10 – Implementing Secure WLANs



MOD 9- 66

The figure provides a brief review of lab tasks. Please see your lab guide for details. When complete, you will continue with Module 10 – Implementing Secure WLANs.

Rev.22.41

aruba
a Hewlett Packard
Enterprise company


Implementing Secure WLANs

Aruba Campus Access
Fundamentals



Welcome back! This is Module 10: Implementing Secure WLANs.

CONFIDENTIAL - FOR TRAINING



Objectives

- Explain AAA
- Describe 802.1X authentication
- Configure secure WLANs
- Discuss roles and access rules

WLAN RF data emanates out each AP's radio in all directions. Security of this resource is vital.

After completing this module, you will be able to explain authentication, authorization, and accounting (AAA) services and describe 802.1X authentication. You will also be able to configure secure WLANs and describe roles and access rules.

Introduction to AAA


We begin with an introduction to AAA.

CONFIDENTIAL - FOR TRAINING

Why Network Authentication

Zero trust networks require:


- Visibility into user/device identity
- Control over user/device access
- Access based on identity, device type, and other factors



MOD 10- 4

You have learned about Aruba's ESP platform. Zero-trust networking is one of the ESP elements for delivering security in zero trust environments. A key component of this security strategy involves network authentication. By identifying every user and device at the internal perimeter, a company gains visibility into their environment. Such visibility lays the foundation for control. First, the company can take action to control access to the internal network, ensuring that only authorized users and devices can connect. Next, the company can enhance this control by assigning the appropriate access controls to each client, based on knowledge about who is using the client, as well as other factors such as device type and location. This is called network access control or NAC.

AAA Services




A common NAC solution is based around running AAA services on the authentication Server. So, the three A's of AAA services.

First up is authentication: who are you? The RADIUS or TACACS+ server permits or denies access based on your credentials. Perhaps your username is Maria, with password 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).

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 be assigned to a certain VLAN, with certain firewall and quality of service (QoS) settings, which the switch port dynamically applies.

You have controlled who can access the network, and what they can do. Now we need accounting: what did they do? With accounting you 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.

Controlling Network Access




When a user joins the organization, they are given an account. This account holds details about their relationship to the organization and their authentication credentials. These credentials are often stored in a Microsoft AD or LDAP database.

During authentication, the user provides credentials, which the network validates against the data store. This is typically done via a AAA server like ClearPass. Then the network authorizes the user's account details and client device. The network may grant or deny access based on a complete picture of the client. If access is granted, the system continues to monitor client activity using start, stop, and interim accounting messages, with ongoing compliance checks.

By leveraging all available information and implementing sound access rules, ClearPass AAA-based deployment provides a complete access control system, with three common authentication methods.

Layer 2 vs Layer 3 Authentication



MOD 10- 7

You can divide network authentication methods into two broad categories: layer 2 and layer 3. Layer 2 authentication methods authenticate users and devices before they receive an active layer 2 connection. If the client fails authentication, the client is disconnected at layer 2, preventing the client from sending or receiving any traffic at all. MAC authentication (MAC-Auth) and 802.1X are examples of layer 2 authentication methods.

Layer 3 authentication takes place after the client receives an active layer 2 connection. Typically, the client will initially receive highly limited access. For example, the unauthenticated client might be allowed to use DHCP to obtain an IP address and DNS to resolve hostnames. To receive a higher level of access, the client must complete layer 3 authentication. Captive portal or IPsec VPNs are examples of a layer 3 authentication method.


MAC-Auth

Advantages

- No requirements or setup on client devices

Disadvantages

- Identifies device only
- Easily spoofed, not user
- Difficult to create and manage allow list



MOD 10- 8

MAC Authentication (MAC-auth) is one example of a layer 2 authentication method. A network infrastructure device that implements MAC-auth checks the source MAC address in a client's incoming traffic. Sometimes the authenticator has a local list of authorized MAC addresses. Sometimes, as shown in the figure, the authenticator submits the MAC address to an authentication server, which uses the MAC address to determine whether to grant the client access. For example, the server might have a allow list of MAC addresses that belong to the company; it only grants access to MAC addresses on that list.

Because wireless clients always include their MAC address in 802.11 or Ethernet frames, the client does not need to be configured to submit credentials. Therefore, this authentication method provides an option for devices that do not have a user interface, sometimes called headless devices, or lack support for 802.1X.

However, MAC-auth does not provide rigorous security. First it identifies the device only, not the user on the device. Worse, it is often trivial for a hacker to discover MAC addresses used in your system, whether by eavesdropping or accessing a legitimate device. The hacker can then easily spoof an authorized MAC address. For example, the Windows OS lets you change the MAC address in the NIC properties. Finally, it can be difficult to create and manage a allow list of the MAC addresses of all the company's authorized devices. Apple and Android devices are now starting to randomize their MAC addresses when they connect to different SSIDs, making it even harder to create a definitive allow list.


Rather than a security measure, MAC-auth is perhaps most useful for simplifying the

provisioning for devices such as printers, Internet of Things (IoT) devices, IP phones, and security cameras.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

802.1X Overview

- ✓ Requires clients to authenticate to a RADIUS server as soon as the endpoint connects
- ✓ Authenticates clients with individual credentials
- ✓ Supports a variety of credential types
- ✓ Most secure option



MOD 10-9

Another layer 2 authentication method is 802.1X, which is recommended for secure authentication in enterprises. 802.1X, an IEEE standard for Network Access Control (NAC), enables the network infrastructure to authenticate and authorize endpoints for network access. 802.1X requires users to authenticate as soon as connection is established. In a wired network, 802.1X takes effect as soon as the Ethernet connection comes up. In a wireless network, it takes effect after the initial 802.11 association.

To restrict the unauthenticated client's access, 802.1X divides the port into two logical ports. The controlled port permits all types of traffic. Its state depends on authentication status. Both sides of the connection should block the controlled port until the peer has successfully authenticated. After successful authentication, all types of traffic may flow through the controlled port from the point of view of 802.1X; however, the network infrastructure can further restrict access with firewall policies and ACLs. The uncontrolled port is always open, but it permits only Extensible Authentication Protocol (EAP) packets. It does not carry any other form of traffic, even DHCP messages and DNS messages. 802.1X allows an endpoint to authenticate before it has acquired an IP address.


802.1X Authentication



If you want the best authentication in a wireless and wired campus network, you need to implement 802.1X authentication.

CONFIDENTIAL - FOR TRAINING

802.1X Roles



MOD 10-11

The 802.1X process includes three parties. To understand these parties, you will consider an analogy in which a person is attempting to get into an exclusive, very important person (VIP) event in a secure building.

Supplicant

The supplicant is the 802.1X component that runs on the endpoint attempting to connect to the network. In our analogy, Bob, the person attempting to get into the event, is the supplicant. He is the one who wants to gain access to something.

In the real world, 802.1X supplicants are software that run on client devices. For example, Windows and MAC operating systems have native 802.1X supplicants, as do smartphone operating systems such as Android. Administrators can also install third-party supplicant software on endpoints.

Authenticator

The authenticator controls network access. It enforces decisions made by the authentication server, either permitting or denying access. In our analogy, the authenticator is the security guard at the building in which the event is being held. The security guard is responsible for guarding the building and only allowing authorized people in. Similarly, in 802.1X the authenticator stands at the internal network perimeter and either grants or denies access to the network. In Aruba networks, the authenticator could be an Aruba AP, an Aruba gateway, or

an Aruba switch.

Authentication Server

The authentication server directs the authenticator and ultimately takes responsibility for access decisions. You can think of the authentication server as the event organizer. She has the VIP list and tells the security guard who to let in.

In 802.1X, the authentication server is a RADIUS server. To make its access decisions, the RADIUS server checks users' credentials, but also considers other factors. Is the user connecting at the right time or location? An authentication server may, for example, deny users with legitimate credentials if they are trying to connect outside of regular work hours. The decisions all depend on the authentication server's policies.

For Aruba networks, it is recommended that you use ClearPass Policy Manager (ClearPass), which includes a powerful RADIUS server among many other features.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

802.1X High-level Analogy


Bob



Security guard



Event organizer with VIP list




MOD 10-12

You will now look at the 802.1X process. To help you understand the process, return to Bob, who is trying to get into the VIP event. Bob comes to the building and is met by the security guard. The guard informs him that the event is closed to the public and asks for Bob's name. Bob gives him his name and says that he is on the VIP list.

The security guard has an audio/visual communication channel open with the event organizer, who is inside the building. He's going to act as a courier for everything that Bob says, passing it on to the event organizer and similarly, passing on the event organizer's requests to Bob. The guard informs the event organizer that a person claiming to be Bob is outside. The event organizer tells the guard to collect Bob's ID. The guard passes on the message and has Bob hold his ID up to the camera for the event organizer to check. Checking the ID and her list, she sees that the person at the door is indeed Bob, who is on her list. She tells the security guard to let him in, which the guard does.

Please see Aruba's Security Courses for more detailed information about 802.1X.

802.1X Communication Protocols



MOD 10-13

In a moment, you will see how this process translates to 802.1X. To understand the process, you need to understand the two communication protocols used in 802.1X.

EAPOL

The supplicant and authenticator communicate with EAP over LAN (EAPOL) messages. EAPOL operates at the data link layer and uses EtherType 0x888e. Note that, despite the name, even wireless connections use EAPOL messages. In practice, the messages are often called simply EAP messages.

802.1X almost entirely shuts down the layer 2 link for unauthenticated clients. The client cannot even send DHCP or DNS messages into the network. The authenticator accepts nothing except EAP. This means that usually the client has no IP address. It can only communicate on the local link. The supplicant is either directly connected to the authenticator, or, if the authenticator is an Aruba gateway, an AP tunnels the supplicant's messages to the authenticator (gateway).

RADIUS

The authenticator reaches the AAA server over an IP network, and multiple layer 3 hops can exist between the devices. The authenticator acts as a relay between the supplicant and the authentication server much like the security guard in the previous example. The authenticator encapsulates EAP messages from the client with RADIUS access-requests, which it sends to

the AAA server. Similarly, the AAA sever encapsulates EAP messages for the client in RADIUS-challenge messages, which it sends to the authenticator. The authenticator decapsulates the EAP messages and passes them on to the client.

In this way, the supplicant and authentication server act as the two endpoints in an EAP exchange with the authenticator transporting messages between them.

Note that, in RADIUS terms, the authenticator is the network access server (NAS), network access device (NAD), or RADIUS client. However, the rest of this discussion will use “authenticator” to refer to the device for consistency.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Examples of EAP Methods

EAP method	Credentials	Comments
EAP-TLS	Server = Certificate Supplicant = Certificate	• Recommended option • Widely supported - Windows native • TLS handshake authenticates server and supplicant
PEAP	Server = Certificate Supplicant = Username and password (typical)	• TLS handshake creates secure tunnel (outer method) • A less secure authentication (MSCHAPv2) handles inner method
EAP-TTLS	Server = Certificate Supplicant = Username and password (typical)	• Similar to PEAP in concept
EAP-TEAP	Server = Certificate Supplicant = Username and password or certificate	• Emerging method (Windows 10, ClearPass 6.9 support) • TLS handshake creates a secure tunnel (inner method) • Enables "chaining" of machine and user auth with tokens



MOD 10- 14

Many specific EAP authentication methods exist within the flexible framework. Some common ones are listed in the table. One important point is that each of these methods provides mutual authentication—both the server and supplicant authenticate. It is important for the supplicant to authenticate the server to avoid man in the middle (MITM) attacks in which the supplicant connects to a hacker device rather than the legitimate network. EAP-TLS is the preferred method and the one that you should implement for enterprise networks.

EAP-Tunneled Layer Security (TLS)

Considered one of the most secure EAP methods, EAP-TLS uses a TLS 1.2 handshake to authenticate the server and supplicant with digital certificates. The Windows native client supports EAP-TLS, as do many other supplicants. ClearPass also supports EAP-TLS. While deploying certificates to clients can be challenging, methods exist for easing and automating this deployment. You will look at some of these methods later.

Protected EAP (PEAP) and EAP-Tunneled TLS (TTLS)

PEAP was developed to provide some of the security of EAP-TLS without forcing stations to have digital certificates. It is the option that the Windows native client uses when the user logs in with a username and password.

Tunneled Transport Layer Security (TTLS) is a similar, but less commonly used, method. (TTLS was developed by Funk Software and Certicom; PEAP was developed by Microsoft,


Cisco Systems, and RSA Security.) ClearPass supports both PEAP and EAP-TTLS.

TTLS and PEAP function in similar ways. Both methods involve a two-phase authentication process. In the first phase, the outer method creates a secure tunnel using TLS. The authentication server then initiates a second authentication method, called the inner method, inside the secure TLS tunnel. PEAP and EAP-TTLS can use several different options for the inner method. (In fact, they can even use EAP-TLS as the inner method; however, doing so does not typically make sense because the goal of PEAP and TTLS is to protect a less secure inner method.) EAP-Microsoft CHAPv2 (MSCHAPv2) is the most common option. EAP-MSCHAPv2 is not secure enough to run over an untrusted network but is acceptable when protected by the TLS tunnel.

EAP-Tunneled EAP (EAP-TEAP)

Like PEAP and EAP-TTLS, EAP-TEAP uses TLS as an outer method to create a secure tunnel. It then uses an inner method to authenticate the client. Depending on the inner method implemented, the client can authenticate with a username/password or a certificate. EAP-TEAP also supports “chaining” of multiple authentications, which can be useful when a company wants to ensure that users connect on authorized machines. The client can authenticate with machine credentials first, proving the device is authorized. The client obtains a token during this authentication. The client can then authenticate with user credentials, including the token from the machine authentication, proving that both the user and device are authorized. EAP-TEAP is supported in Windows 10 and ClearPass 6.9.

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 connection. An AP periodically sends out beacons to advertise the WLANs that it supports. The beacons include information about the network name, or SSID, and the security options. A client typically also sends a probe request for the specific SSID to which it wants to connect. If the AP supports that SSID with compatible options, it will send a probe response.

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. 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. Access Point is the 802.1X authenticator. As you learned earlier, the 802.1X process can involve many packets, but in the end the client is authenticated.

The final step for all WLANs that use encryption is an exchange called the four-way handshake. The client and AP use this handshake to derive the keys used to secure communications; 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.

802.1X Authentication Process



MOD 10- 16

Let's look at the 802.1X authentication process in detail, to improve your theoretical knowledge and real-world diagnostic abilities. 802.1X occurs after a successful 802.11 association.

802.1X carries EAP messages which are used for authentication. EAP defines a flexible framework that supports several authentication methods, to meet your particular environment and policy needs. The figure shows this general framework. As you follow the process, remember that the NAS authenticator relays all messages from the client to the authentication server, encapsulating them with RADIUS as necessary. The vertical line underneath the switch shows the point at which the authenticator encapsulates the messages.

The client establishes a connection. Network access is controlled, blocking all traffic except EAP, and sends an EAP Request/Identity challenge. This initiates authentication and asks the client only to identify itself – do not send any other information.

The client responds with an EAP Response/Identity packet, which typically includes a username. This is encapsulated in a RADIUS message and forwarded to the RADIUS authentication server.


The authenticator includes the user's identity in all future frames so that the server can keep track of which EAP messages belong to which user. The user's identity also marks any accounting frames for the connection - important for companies with rigorous auditing requirements. Depending on the EAP method, though, the username submitted at this point

might not be the actual username. Some methods have the supplicant submit a false name at this point and the true name later in the process after a secure, encrypted session is established. This protects the username from hackers.

The server initiates an EAP method for exchanging credentials. Message exchange details, flow, and authentication credentials depend on the EAP method. Based on information received in the exchange, the authentication server grants or denies access and informs the switch. The authenticator then either activates the connection and transmits an EAP-Success or leaves the connection deactivated and transmits an EAP-Failure.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

EAP and RADIUS Together



Now let's look at how the authenticator relays messages between the client supplicant and authentication server. You must configure the authenticator (sometimes referred to as a NAS) and the authentication server to work together. You must configure the NAS with the authentication server's IP address, so the NAS knows where to send packets. And you must configure the authentication server with the IP addresses of acceptable NAS devices, so it knows who to accept requests from. Both the authenticator and the authentication server must be configured with a matching shared secret: a simple alphanumeric string of your choice.

The authenticator encapsulates a supplicant's EAP message in a RADIUS access-request packet. It also hashes the packet with MD5, using the RADIUS shared secret as the key. It adds the hash to a message authenticator field. The authenticator then transmits the packet to the authentication server. This communication is between the NAS IP and the authentication server IP address: the client does not need an IP address or any network access.

The authentication server checks that the message has arrived from a valid network device (sometimes called a Network Access Server [NAS] or RADIUS client); that is, the switch IP address must match a network device address configured on the server. The authentication server also checks the message authenticator field by hashing the packet with its shared secret. If the request is not from a known network device with a correct shared secret, the authentication server does not respond to the message: it silently discards the request.

If the switch is known and the secret is correct, the server decapsulates the EAP messages and formulates the appropriate responses, which it encapsulates in RADIUS challenge messages. These messages are also authenticated with the message authenticator field.

At the end of the process, the authentication server sends either:

An access-accept message, if requirements are met:

- The client has successfully authenticated.
- The policy that applies to this access request allows access.


An Access-Reject message, if one of these situations occurred:

- The client could not complete authentication, whether because the authentication server does not trust the authenticator, because the client and server do not support compatible EAP methods, or some other problem.
- The client does not trust the authentication server's certificate, so the server failed authentication.
- The client failed authentication.
- The policy that applies to this access request does not allow access.
- No authentication server policy applies to this access request.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

EAP-TLS Exchange

④ Exchange: EAP-TLS



As mentioned earlier, step 4 in the EAP/RADIUS process is where users prove their identity. The server's EAP request/method packet starts the process and indicates the method that the server requires. You can configure most authentication servers to select methods according to conditions such as user identity and location.

If the station supports the requested method, it continues the exchange. Otherwise, the station sends an EAP NAK packet, which can suggest a different method. The server, if it supports the alternative method, may then initiate the exchange with the new method. The station and the authentication server must agree upon the method for the process to continue.

The server sends their certificate to the user, which validates the supplicant against a local store of acceptable certificates. The EAP Request/TLS and EAP Response/TLS packets include information such as:

- Digital certificates and certificate verifications
- Supported encryption suites
- Values for generating encryption keys (not the keys themselves)

The digital certificate includes the device's identity, and it is signed by a Certification Authority (CA). The authentication server validates the certificate submitted by the user, and the

supplicant validates the authentication server's certificate. Validation occurs through asymmetric encryption with public and private keys, which is beyond the scope of this course. The validation process also includes checking the CA signature, so the certificate must be signed by a CA that the device is configured to trust.

After the server validates the user certificate, it sends a message to indicate a successful finish, and the client responds. Recall that after the EAP exchange successfully completes, step 5 occurs: the authentication server sends a RADIUS access-accept message.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

EAP-TLS Considerations

Advantages

- Very strong security
- Mitigates risk of password-based authentication

Concerns

- Overhead of managing CA services
- Staff training to manage CA services
- Certificate distribution and management



MOD 10-19

The figure summarizes the advantages and concerns related to EAP-TLS. Please click on each box in the figure to explore.

Advantages

EAP-TLS is one of the most secure EAP methods because it provides mutual authentication with Public Key Infrastructure (PKI) digital certificates. Digital certificates rely on extremely strong asymmetric keys and are validated by trusted Certificate Authorities (CA). The digital certificate is typically stored on a laptop or a smart card, providing stronger security. Although you can set up requirements to force users to create stronger, and therefore more secure, passwords, you cannot prevent users from telling people their password or writing it on a paper displayed in plain sight. Generally, it is harder for users to leak certificates inappropriately; systems administrators might, for example, implement security measures to prevent users from copying certificates.

Hackers often run dictionary attacks on passwords, phish for passwords, or use other social engineering techniques to get users to leak their passwords. Although no credentials can be totally secure—malware might be able to steal certificates, for example—digital certificates are more resistant to many more common attacks.

Someone who steals a laptop or smart card might be able to gain access to the certificates within, but they are usually protected by an additional layer of security such as a password, which the user must enter before the supplicant can use the certificate. In addition, the user

can immediately report the theft, allowing you to revoke the compromised certificate.

Concerns

EAP-TLS is not vulnerable to eavesdropping and replay attacks that affect less secure EAP methods such as EAP-MD5, but security comes at the cost of purchasing and managing the digital certificates—substantially more expensive than managing passwords. Maintaining a large number of certificates requires specialized software and trained IT staff.

In addition, in environments that have multiple operating systems, installing certificates on all clients can be difficult. Bring Your Own Device (BYOD) environments can make managing certificates even more difficult. However, Aruba ClearPass BYOD onboarding capabilities can help companies surmount these challenges. ClearPass Onboard can provide new user devices with limited access and helps users.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

④ Exchange: PEAP



Step 1—outer method: Like EAP-TLS, TTLS and PEAP use a TLS handshake to generate encryption keys and negotiate a tunnel secured by those keys. Only the server authenticates itself with a digital certificate during this exchange, but you must ensure that the clients trust the CA that signed the server certificate. The encryption keys for the tunnel are derived from the public key in this certificate.

The authentication server sends a new EAP request/identity packet, and the supplicant responds with their identity. The server then begins the inner authentication method.

The authentication server sends a challenge. The supplicant responds with a hash of the

password. The server validates that the hash is correct, and, if so, it sends a message indicating the success. The supplicant responds. The figure shows only the PEAP part of the complete 802.1X authentication process. After these messages, the tunnel is taken down, and the authentication server sends a RADIUS access-accept to the switch.

Note: For 802.1X on wireless connections, whether implemented through EAP-TLS, EAP-TTLS, or PEAP, the authenticator sends the encryption material. The wireless client (supplicant) and authenticator continue to use encryption to secure the connection. On Ethernet connections, however, the client and switch port do not encrypt traffic in this way.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

EAP-PEAP Considerations

Advantages

- Strong security
- Easier to implement than EAP-TLS

Concerns

- Slightly weaker security than EAP-TLS



MOD 10- 21

EAP-PEAP Advantages


Like EAP-TLS, EAP-TTLS and PEAP provide strong, mutual authentication. Because TTLS and PEAP use encrypted tunnels to secure usernames and passwords (rather than require digital certificates on stations), you can implement these methods more easily than you can TLS.

TTLS has one unique benefit: it always protects the username. Depending on how PEAP is implemented, the username might be transmitted in plain text in the initial EAP/response identity packet, allowing a hacker to detect the user's identity and possibly lock the user out of his or her account.

Concerns

As compared to EAP-TLS, EAP-TTLS and PEAP have the disadvantage of a weaker authentication method.

Options For Deploying Certificates On User Devices



As you see, ensuring that devices have the correct certificate can be challenging. Maintaining a large number of certificates requires specialized software and trained IT staff. In addition, in environments that have multiple operating systems, installing certificates on all clients can be difficult. There are many options for deploying and managing certificates. In the figure you see just a few of those options.

Typically, customers set up a private CA to issue the certificates required for EAP-TLS. Microsoft Windows CA is a common option for Windows domains that use Active Directory (AD). AD administrators can set up group policy objects (GPOs) to automatically push down certificates to devices joined to the domain. The Windows CA also supports Active Directory Certificate Services (ADCS) Web Enrollment, which enables users to obtain certificates through a self-service web portal.

However, the Windows CA options are geared for a world in which the only devices are computers joined to the domain. A modern network has far more types of devices, including smartphones and Internet of Things (IoT) devices. Bring Your Own Device (BYOD) environments can make managing certificates even more difficult. There are mobility device managers (MDMs) that help to automatically enroll mobile devices for certificates and push certificates to devices.

If a company needs to deploy certificates to many unmanaged devices, Aruba ClearPass Onboard can help. ClearPass Onboard helps users to easily download the correct certificates to their devices. Onboard can integrate with multiple CAs. It can act as a Intermediate CA or

issuing certificates for devices. When Onboard integrates with a Windows CA PKI, Onboard can alternatively act as an intermediary with ADCS, in which case the Windows domain signing CA issues client certificates, but the Onboard client makes the process of obtaining the certificate transparent to the user.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY



Secure WLAN Setup

Now that you are familiar with the 802.1X authentication process, this section introduces you to configuring WLANs that use 802.1X authentication.

CONFIDENTIAL - FOR TRAINING

Configuring Secure WLAN

Create a New Network

1 General

2 VLANs

3 Security

4 Access

5 Summary

Security Level	Enterprise	Personal	Visitors	Open
Key Management	Dynamic WEP	WPA	WPA2	WPA3
RADIUS Server	AP must direct RADIUS requests to proper server Create RADIUS Profile name, IP address, and shared secret			



MOD 10- 24

During the WLAN configuration in Aruba Central, you will have an opportunity to choose various security options, as shown in the figure.

The Enterprise security level is most appropriate for employee WLANs. Recall that WPA3 is the most secure key management option. You must also specify an authentication server: the AP must know where to direct RADIUS requests during client authentication. Remember, 802.1X/EAP authentication is between the end client device and the authentication server (RADIUS server), with the authenticator (AP) acting as an intermediary for this process.

During lab activities, you will perform this process, including specifying a RADIUS profile with a name, IP address, and shared secret. This shared secret must match for what the authentication server has been configured for. If not, the authentication server will reject requests from the AP.




Roles and Rules

Now you explore Aruba roles and firewall rules: these are used to control the supplicant's access once they are connected (authenticated) to the network.

CONFIDENTIAL - FOR TRAINING

Role-Based ArubaOS Firewall



The role-based ArubaOS firewall can enforce different rules based on user identity and assigned role, rather than simply on features in packet headers. For example, the firewall can restrict contractors web traffic to a limited range of IP addresses while allowing engineers' traffic to flow more or less unrestricted.

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 rules based on source IP address. 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 users in the same subnets, but easily create different rules for different 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 ClearPass, the Aruba authentication server. 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 10- 27

An Aruba role is a group of settings that controls the client access. They bundle 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 in firewall policies. These rules filter traffic based on layers 2 through 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. 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 access control and application control 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, 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.

Example Role Definitions

Role	Access Rules For Selected Roles
test-role 	→ Enforce Captive Portal
contractor	 ● Deny Suspicious Web Reputation Index
default_wired_port_profile	 ● Deny Social Networking to all destinations
employee	 ● Deny Peer to Peer to all destinations
p58t01-employee	 ● Deny ssh on server 10.10.10.10/255.255.255.255
p58t01-guest	 ● Allow any to all destinations



MOD 10- 28

The figure shows some typical roles that you might define.

Access Network Rules

Access rules

Rule Type:

- Access Control
- VLAN Assignment
- Captive portal
- CALEA
- Bandwidth Contract
- Blocked Page URL
- Redirect Blocked HTTPS

Options:

☐ 802.1p priority

☐ Denylist

Service:

☒ Network

☐ Application

Category:

☐ Application

☐ Web Category

☐ Web Reputation

Action:

- Any
- Any
- CUSTOM
- adp
- bootp
- cftp-tcp
- cups
- dhcp
- dhcp-v6
- dns
- esp
- ftp
- gre
- h323-tcp
- h323-udp
- http-proxy2
- http-proxy3
- http
- https
- icmp
- icmp-v6

Destination:

- Allow
- Deny
- Destination-NAT
- Source-NAT

Destination:

- To all destinations
- Select --
- To all destinations
- To a particular server
- Except to a particular server
- To a network
- Except to a network
- To a Domain Name
- To AP IP
- To AP network
- To conductor IP
- To AP IP all
- To a network alias

Log

☐ Log

Cancel **OK**

When you select either a network-based or roles-based SSID you can also select which type of access rule you want to assign to users. The rule type can be one of the following:

- Access Control is where you can choose from several services and actions.
- VLAN assignment is the process of assigning a configured VLAN to the users.
- Captive Portal allows you to assign a captive portal page to the user.
- CALEA is for law enforcement.
- Bandwidth contracts are only available if you're placing a rule into a role-based SSID and not a network-based SSID.
- You can redirect to a URL page or redirect a blocked HTTPs service.

Once you've selected your rule type then you must select your Service, Action and Destination.

The network service option allows you to either choose any service, or you can choose the custom option and specify TCP, UDP or other port numbers. Many of the standard protocols have been predefined for your selection. The action can be Allow, Deny, Source-NAT or Destination-NAT. Source-NAT and Destination-NAT are forms of allowed traffic. The destination choices are simple and will depend on your preferences.

Application Based Rules

Access rules

NOTE: Application and Application category rules will only be applied on supported platforms.

Rule Type: Access Control

Service: Network, Application

Category: Antivirus, Gaming, Webmail, etc., Yahoo-games, Gmail-chat etc., Shopping, Travel, Games, etc., Trustworthy, high risk etc., Encrypted

Action: Allow

Destination: To all destination

☐ 802.1p priority ☐ Denylist ☐ DSCP TAG

☐ Application Throttling ☐ Disable Scanning ☐ Log



MOD 10-30

APs support both Aruba's in-built Deep Packet Inspection engine (AppRF) and a cloud-based web filtering service (Web_CC) as a single AppRF visibility feature. WebCC is part of the Aruba Central Foundation license. The Deep Packet Inspection engine has over 1,500 applications.

You can use enforcement at the rule level to allow or block applications or app categories after classifying the packets. Application classification requires matching packet contents. It needs one or more packets in both directions for application classification.

You can have a maximum of 8 or 9 packets to classify the flow. Therefore, rules matching for any app rule must leak packets until the classification is complete. If not, then it will classify as not-classified.

Extended Actions

Options:

☐ 802.1p priority	☐ Denylist	☐ DSCP TAG
☐ Application Throttling	☐ Disable Scanning	☐ Log



MOD 10- 31

The extended actions allow you to perform actions beyond the main action. The extended actions will vary according to the service you select. For example, if you chose the service application then there is also an option for application throttling.

Here are common actions you can choose:

The Log action means the AP will send all logs as security logs to a server. You can configure the logging server under System – Monitoring.


The Blacklist action removes this user from the network. The user (and any devices) would not be allowed back on the network unless un-blacklisted or the blacklist time has expired. The blacklist time is set in the security window as Auth failure blacklist time.

Disable scanning disables ARM scanning when the rule is invoked.

DSCP tag sets the priority DSCP priority that can be configured 0 to 63.

The 802.1p priority option tags the level two frame with assigned priorities 0-7.

Radius Server Derived Role



ClearPass server gives you more options


MOD 10- 32

For server-derived roles there must be an exchange of information between the AP and the server. What is sent depends on what type of server you're using and if you've installed RADIUS vendor-specific attributes (VSAs).

The AP sends attributes along with the RADIUS request. The server can use this information to derive a role and reply with Aruba VSAs or RADIUS attributes. When a server returns an Aruba VSA with the RADIUS reply, the Virtual Controller (VC) places the user in the Role or VLAN specified. The AP does not require configuration to look for the VSA. When the server returns a RADIUS attribute, you must configure the VC to look for this attribute. If you use an Aruba ClearPass server, AP-ClearPass communication is automatic. The AP sends a RADIUS request with the Aruba VSAs.

ClearPass can aggregate the information and perform many actions, like profiling, posture and onboarding. ClearPass can also interact with the Active Directory and retrieve attributes. ClearPass will return a RADIUS message using Aruba VSAs with enforcement actions, such as a role name and VLAN identifier, enforcement of bandwidth limitations, or redirect the user to a web page.

RADIUS Role Assignment Roles




Create a New Network

1 General 2 VLANs 3 Security 4 Access 5 Summary

Access rules

Role Based Network Based Unrestricted



Role Assignment Rules

- If Filter-Id equals contract assign role: Contractor
- If Filter-Id contains empl assign role: Employee

Default role: Corporate



MOD 10-33

When using the wizard to create a WLAN in Aruba Central, you can use role derivation to assign a role to the users/devices associated to a WLAN network from the Access tab (see the figure). When the network receives an association request, it can assign the user/device a role based on the AP or the AP-group. If the server requires authentication, the network can assign the role from attributes it receives from the authentication server.

For example, if a user associates to an AP and the authentication server returns an Aruba-User-Role RADIUS filter-ID or VSA with value employee, the AP assigns the employee role. If RADIUS returns the value contractor, the AP assigns the contractor role. This requires coordinating RADIUS values with role assignment rules. If RADIUS does not send a filter-ID, nor an Aruba-User-Role VSA, or the value does not match with a role defined on the AP, then the AP assigns the default role defined for the WLAN. Typically, you want to use the server-derived role and not the default one. The default one should only be used as a fallback in case there is a misconfiguration on either the AP or the authentication server, and therefore, should be defined fairly strict in nature.

Once you select an attribute, the received information is compared to the strings and a role is assigned. This choice depends on operator settings: either match, equals, contain, start-with, ends-with, or not equal. When no match is found, the user device is assigned the default role.



Knowledge Check

Self-check on key learning points




Its time for a knowledge check.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Question #1

Match the AAA service to an appropriate description.



Knowledge Check


Q

MOD 10-35

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Question #2

In the diagram, which device would host AAA services?



a

Knowledge Check

MOD 10-36

RADIUS

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Question #3

Which of the options below accurately compares 802.1X authentication vs. captive portal authentication? (Select three.)

- a. They both occur at Layer 3 of the OSI model.
- b. 802.1X authentication is typically used for corporate employee access.
- c. Captive portal authentication is often used for guest access.
- d. 802.1X authentication requires clients to authenticate to an authentication server as soon as the endpoint connects.

a

Knowledge Check

MOD 10-37

B,C,D

Question #4

During 802.1X client authentication, which devices can perform the Authenticator Role? (Select two.)

- a. Supplicant
- b. AP
- c. Switch
- d. Authentication Server
- e. RADIUS

a

Knowledge Check

MOD 10-38

B,C

Question #5

There are a few popular Extensible Authentication Protocols (EAP). Match the advantage or disadvantage with the appropriate EAP method.

EAP-Transport layer
Security (EAP-TLS)

Protected EAP
(PEAP)

EAP-Tunneled EAP
(EAP-TEAP)

Very strong security, mitigates risk of password authentication.

Easier to implement than some other methods, but slightly weaker security.

May have relatively limited support as compared to other methods.

Enables chaining of machine and user authentication with tokens.

Requires the overhead of managing CA services, certificate distribution and management



Knowledge Check

MOD 10-39

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Question #6

Your colleague is asking about the role-based ArubaOS Firewall. Which statements below can you leverage to improve their understanding? (Select three.)

- a. It can enforce different rules based on user identity and assigned role but increases the reliance on VLANs and subnets for security.
- b. A role is a group of settings that controls client traffic.
- c. A role can include rules that leverage Deep Packet Inspection (DPI) and Web Content Classification (WebCC).
- d. It can leverage roles derived from an authentication server.

a

Knowledge Check

MOD 10-40

B,C,D

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY



Lab Activity

Consider the lab activity.

CONFIDENTIAL - FOR TRAINING

Lab Tasks

Lab 10.01
Implementing a Corporate WLAN

Lab 10.02
Implementing Access Control

Module 11 – Guest
Access



MOD 10- 42

The figure provides a brief review of labs for this module. Please see your lab guide for details. When complete, you will continue with Module 11 – Guest Access.


Rev.22.41

aruba
a Hewlett Packard
Enterprise company

Implementing Guest Access

Aruba Campus Access
Fundamentals

Welcome back team! This is Module 11: Implementing Guest Access.



Objectives

- Describe guest access
- Setup captive portal authentication
- Configure guest WLANs

This module will focus on guest access, where a web browser is used for the authentication process. This is commonly referred to as captive portal. Proper guest access deployment helps to accommodate less trusted users that require limited access: typically, internet only. It also improves security by helping to maintain strict separation between employee and guest users. Captive portal is used to secure guest access as well as onboard employee devices.

After completing this module, you will be able to describe guest access advantages, features, and the overall solution. You will also be able to setup captive portal authentication and configure guest WLANs.

Guest Access Introduction

You start with an introduction to guest access solutions and captive portal options.

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



Aruba allows for a simple collapsed architecture: 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 APs used for the guest 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. Access point traffic is directed via secure tunnels straight to the mobility gateway, where an integral stateful firewall maintains strict segregation between different traffic classes (employee and guest, for example).

Internal traffic gets access to core LAN and corporate resources, while guest traffic travels through a secure tunnel to a mobility gateway in the DMZ, 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 use the Aruba ClearPass server for more sophisticated guest access solutions - credit card processing, access code authorization, property management systems, and more.

Captive Portal Options

Cloud Guest

- Splash page hosted in Cloud
- Managed by Aruba Central
- Can be customized

External Captive Portal

- Splash page hosted on external server
- Customization depends on server
- Example : ClearPass guest



MOD 11- 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 Policy Manager, which can provide guest captive portal authentication with more customization options and scalability.


Captive Portal Authentication



Now you explore captive portal authentication.

CONFIDENTIAL - FOR TRAINING

Captive Portal Authentication




Captive portal authentication is the primary example of layer 3 authentication and is often implemented for guest wireless networks: airports, hotels, and coffee shops. It's an easy way for users to log in without special client device requirements or user expertise: they simply connect to an open network and open a web browser, where the web browser is redirected to a login page (commonly called a splash page). The user then opens a Web browser, which is directed to a login page. The user enters credentials or maybe just accepts an agreement to log in. The user's client then receives more complete network access.

Captive portal deployments can be quite nuanced, based on your specific needs, but let's examine a high-level process here, illustrated in the figure. In this example, the client connects to an Aruba AP, acting as the AAA authenticator.

Upon initial connection to the captive portal WLAN, the AP grants limited client access: enough for the client to get an IP address via DHCP, resolve names via DNS, and just enough web access to access the guest captive portal page. Once the client receives the IP addressing information the user opens a web browser and attempts to browse to a web page. This communication flows through the ArubaOS firewall (on the AP), which redirects the client to a login page on ClearPass Guest or Aruba Cloud in Central. The user logs in. Then a RADIUS message is communicated from the external server or Aruba Cloud back to the AP (authenticator). The RADIUS attributes will change the firewall policy, minimally so that the guest can browse the Internet or use the network resources, as dictated by their updated user role.

Captive Portal Splash Pages



MOD 11- 8

You can use the Aruba Central Guest app to configure splash page profiles. The splash page originates from the Central account. You can configure several splash pages to allow different groups or guest SSIDs to send different splash pages as needed.

The guest management feature allows guest users to connect to the network and allows you to control guest user access. You create a splash page profile for guest users in Aruba Central. Then guest users access the internet by providing either credentials, which can be configured by a guest operator, or perhaps their respective social networking login credentials.

You can easily customize your splash page with a corporate logo, color scheme, and terms of service. Enable user login with a social networking service like Facebook®, Google +®, Twitter®, or LinkedIn®. Facebook's Wi-Fi feature allows you to pair your network with a Facebook business page.

Guest Splash Page: Authentication Type

NEW SPLASH PAGE

1

Configuration

2

Customization

3

Localization

Guests log in to splash page, no credentials, optional PSK
To enable PSK, set the Guest Key to ON, specify a password

Type:

Anonymous

Authenticated

Facebook Wi-Fi

Must set authentication/authorization attributes and login credentials
Authenticate using sponsored access and social networking login profiles

Pair your network with a Facebook business page, guests can use Facebook credentials
Users attempt access, browser redirects to Facebook page
User logs in, then checks in to access free Internet, or skip checking in and continue.

Social Login:

MOD 11-9



To create a new splash page, first configure its name. Then select the type of splash page: Anonymous, Authenticated, or Facebook Wi-Fi.

- The Anonymous login method allows guest users to log in to the splash page without providing any credentials. You can also enable a pre-shared key for anonymous user authentication to allow access. To enable a pre-shared key-based authentication, set the Guest Key to ON and specify a password.
- Authenticated requires you to configure authentication and authorization attributes and login credentials that enable users to access the Internet as guests. You can configure an authentication method based on sponsored access and social networking login profiles.
- Facebook Wi-Fi gives you the option to enable network access through the free Wi-Fi service Facebook offers. This feature allows you to pair your network with a Facebook business page, thereby allowing the guest users to log in from Wi-Fi hotspots using their Facebook credentials. If you have a Facebook Wi-Fi business page set up, when the users try to access the Internet, the browser redirects the user to the Facebook page. The user can log in with their Facebook account credentials and can either check in to access free Internet or skip checking in and continue.

Guest Splash Page: Authentication Method

Username/Password:	☒	Authenticate with user credentials created on Aruba Central
Self-Registration:	☐	Allow guests to self-register
Social Login:	☐	Social Login: Facebook, Twitter, LinkedIn, Google
Allow Internet In Failure	☐	
Override Common Name	☐	
Authentication Success Behavior:	☒	Return to Original URL
	☐	Redirect URL

MOD 11- 10

When you set the splash page type to Authenticated, you should configure an authentication method:

- By Default, when you set a username/password, an administrator or guest administrator user will need to manually configure the guest user accounts.
- An option to avoid an overhead of creating new guest accounts is to enable self-registration, which allows guest users to register their own accounts. In addition, as a security measure, email and phone verification are available to validate the guests' identities.
- Another option is to enable the social login, which uses single sign-on with a social media pages to authenticate the guests.

You can enable the social login option to allow guest users to use their existing login credentials from their social networking profiles and sign in through a third-party website. When Central configures a social login-based profile, the guest does not have to have a new guest account to access the guest network or third-party websites.

Facebook, Twitter, and LinkedIn allow guests to use their social media credentials to log in to the splash page. To enable site integration, you must create an app for each and obtain the app ID and individual secret key as well. To complete the integration, you will need to enter the app ID and secret key for client ID and client secret respectively (for more information on app creation, search [Creating a Facebook App](#), [Creating a Twitter App](#), or [Creating a LinkedIn App](#).)

Similarly, Google+ allows guests to use their Google+ credentials to log in to the splash page. To enable Google+ integration, you must create a Google app and obtain the app ID and secret key (for more information, search [Creating a Google App](#).)

To restrict authentication attempts to only Google host domain members, enter the domain name in the Gmail for Work Domain text box. Ensure that you have a valid Google Domains or Google Apps licensed domain account. For more information see:

https://apps.google.com/intx/en_in/

<https://domains.google.com/about/>

Then specify the text for the login button.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Guest Splash Page Customization

NEW SPLASH PAGE

1 Configuration 2 Customization 3 Localization

Layout: Horizontal, better for computers ▼

Background color: #ffffff

Button color: #0096d6

Header fill color:

Page font color: #bbbbbb

Logo: Browse...

Background Image: Browse...

+ Terms and Conditions Settings

+ Ad Settings

Sign in
airheads
COMMUNITY

Username

Password

☐ I accept the [terms and conditions](#)

Sign in

Register >

Self-Registration

MOD 11- 11

After creating a guest splash page, you must customize it to your specifications. For example, you can set your own colors and enter your logo and background image. You can also set up the terms and conditions and add an image with a URL to an external site.

Here are the items you can customize:

Layout: To customize the page layout based on the device type, specify a layout by selecting one of the following options:

- Horizontal, better for computers
- Vertical, better for phones

The horizontal layout is selected by default. To change the layout, click the drop-down list and select the required layout type.

Background color: To change the splash page color, select a color from the Background Color palette.

Button color: To change the sign-in button color, select a color from the Button Color palette.

Header fill color: Select the splash page header's fill color from the Header fill color palette.

Page font color: To change the text font color on the splash page, select a color from the Page font color palette.

Page font Color: Select the splash page's font color from the palette.

Logo: To upload a logo, click Browse, and browse the image file. Ensure that the image file size does not exceed 256 KB.

Background Image: Click Browse to upload a background image. Ensure that the background image file size does not exceed 512 KB.

Page Title: Add a suitable title for the splash page.

Welcome Text: Enter the welcome text to be displayed on the splash page. Ensure that the welcome text does not exceed 20,000 characters.

Terms & Conditions: Enter the terms and conditions to be displayed on the splash page. Ensure that the terms and conditions text does not exceed 20,000 characters. The text box also allows you to use HTML tags for formatting text. For example, to highlight text with italics, you can wrap the text with the `<i>` `</i>` HTML tag. Specify an acceptance criteria for terms and condition by selecting any of the following options from the Display "I Accept" check box:

- No, Accept by default
- Yes, Display check box

If you require the I ACCEPT check box to display on the splash page, select the display format for terms and conditions. Ensure that Display Option For Terms & Conditions has the Inline Text option auto-selected and displayed as an un-editable text.

Ad Settings: If you want to display advertisements on the splash page, enter the URL in the Advertisement URL.

Advertisement Image: click Browse and upload the image.

Localization

Splash Pages Visitors

NEW SPLASH PAGE

1 Configuration 2 Customization 3 Localization

Optional fields to localize (translate) guest portal's labels and text

Login Section

Login button title:

Network login title:

Login page title:

Access denied page title:

Logged in title:

Username label:

Username placeholder:

Password placeholder:

Email address placeholder:

Register button title:

MOD 11- 12

The localization page allows you to localize (or translate) the cloud guest portal content to the guest's local language: configure the parameters and enter titles for the various buttons.



Guest WLAN Setup

Configuring a guest WLAN is very similar to creating an employee WLAN: the main difference is that the captive portal process is used (a splash page) instead of 802.1X.

CONFIDENTIAL - PROPRIETARY

Guest WLAN – Cloud Guest

Create a New Network

1 General

2 VLANs

3 Security

4 Access

5 Summary

Security Level

Enterprise

Personal

Visitors

Open

Access Type

External Captive Portal

External server hosts splash page

Cloud Guest

Aruba Central Splash Page

Choose Splash Page

AP must direct RADIUS requests to proper server

Create RADIUS Profile name, IP address, and shared secret

Access Control

Allow guests to get an IP address, use DNS, access internet

Deny guest access to any and all corporate resources

MOD 11- 14

The main difference, from a configuration perspective, for creating a guest WLAN versus an employee WLAN is the need for a splash page and the information you define under the Security tab in the WLAN wizard. When creating a new WLAN for guests, the visitors' security level is chosen under the Security tab: it enables the SSID's captive portal options (see the figure).

There are then two access type options:

Choose External Captive Portal if you use some external server to host the splash page for authentication. You might prefer this option if your organization has very complex requirements or requires heavy customization. You must configure a Captive portal profile with authentication type, server address, and URL of the splash page. You must also configure a RADIUS server, allowing the APs and gateway communication with the RADIUS server to authenticate guest users.

Choose Cloud Guest to have Aruba Central provide the splash page. This is usually powerful enough for most organizations and is highly customizable.

Then choose a splash page. Some customers may have different splash pages for different locations. You will learn to create splash pages in Aruba Central during lab activities, along with other required steps to create fully functional guest WLANs.

As a best practice, you may also navigate to the WLAN wizard's Access tab to configure

firewall rules: to restrict guests from accessing the corporate network. For example, you might want to allow guests to obtain an IP address using DHCP and have access to DNS servers. However, you want to deny any and all access to corporate resources.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Guest Visitors

Session List (2)							
Visitors	Login Type	Browser	MAC Address	Device Type	OS Name	Login Time	Session Time (hh:mm:ss)
guest2	web		FE-21-C9-9...	iPad	iOS	12/08/21 0...	0
guest1	web		7A-A5-8C-A...	iPhone	iOS	12/08/21 0...	0

5102550

Per Page

⏪

<

>

⏩

Page: 1/1



MOD 11- 15

The figure shows typical information on guest session and account details of the visitors who access the splash page. During the lab activities, you learn to create and manage visitors. You will also learn to monitor their sessions and get precisely the information you need, using filters and the search tool.



Knowledge Check

Self-check on key learning points



Its time for a knowledge check.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Question #1

Which options below can help you explain Aruba guest access captive portal options to your colleagues? (Select four.)

- a. Cloud Guest is relatively simple to implement, even for multisite environments.
- b. Cloud Guest includes self-registration and social login features.
- c. The Cloud Guest splash page is not customizable.
- d. External captive portal can be implemented with ClearPass or a third-party solution.
- e. The external captive portal performs client authentication, while the splash page remains hosted on Aruba Central.
- f. You might use the external captive portal if you have very high customization and/or scalability requirements.

a

Knowledge Check

MOD 11-17

A,B,D,F

Question #2

It is time to play **True** or **False**!

- You can use the Cloud Guest to configure profiles for splash pages that originate from your Aruba Central account. **TRUE**
- For guest users to access the internet, they must enter a username and password. **FALSE**
- You can set a guest splash page to Type = Authenticated, so guests can login using sponsored access or social networking credentials. **TRUE**
- Instead of creating guest user credentials yourself, you can allow guests to self-register. **TRUE**

Q

Knowledge Check

MOD 11- 18

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY



Lab Activity

Consider the lab activity.

CONFIDENTIAL - FOR TRAINING

Labs

Lab 11.01
Implementing Guest Access with Cloud Guest

Lab 11.02
Implementing Guest Access with ClearPass Guest

Module 12 –
WLAN Security



MOD 11- 20

The figure provides a brief review of labs for this module. Please see your lab guide for details. When complete, you will continue with Module 12 – WLAN Security.

Rev.22.41


aruba
a Hewlett Packard
Enterprise company

WLAN Security

Aruba Campus Access
Fundamentals

Welcome back everyone. This is Module 12: WLAN Security

CONFIDENTIAL - FOR TRAINING



Objectives

- Describe WLAN security
- Explain certificates
- Describe cloud authentication

Now you learn to deploy secure WLAN systems.


After completing this module, you will be able to describe WLAN security concepts, explain certificates as relates to WLAN security, and describe cloud authentication.

WLAN Security

This module begins by discussing the need for wireless security and exploring the different, common options that are used in today's WLANs.

CONFIDENTIAL - PROPRIETARY

Need for Secure Wireless Communications




Encrypt and authenticate data
Protect against eavesdropping and tampering



MOD 12- 4

Because it is trivial for anyone in the area to collect wireless communications, you must take extra measures to secure all communications with encryption and authentication. Otherwise, you are susceptible to hacks and attacks.

WLAN Security Options



MOD 12- 5

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: users need not submit authentication credentials to connect. The client requests a connection, and the AP permits that connection, which is very simple to implement. 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 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): it is used as the foundation to encrypt and authenticate traffic.

WPA-Enterprise is recommended, due to its superior, 802.1X-based authentication. With 802.1X, a client authenticates to a backend authentication server using individual credentials: a username and password and/or a digital certificate.

Both the Personal and Enterprise flavors do 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 RC4 encryption. WPA and TKIP are now deprecated.


WPA2 and WPA3 both support AES, although, 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 will 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 main focus of 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 used 802.1X for authentication and a key established during the 802.1X exchange to encrypt data. However, both forms of WEP are proven to be insecure and should not be used in today's networks.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

WPA Key Hierarchy



MOD 12- 6

With this brief background in wireless encryption, you are ready to look at the three categories of security in more detail. Let's begin an introduction to WPA. To understand the four-way handshake and data encryption and authentication keys used for all encrypted WLANs, you must have an understanding of WPA key hierarchy.

At the top (see figure) is the Pairwise Master Key (PMK). The PMK itself does not encrypt or decrypt data. Wireless devices derive a Pairwise Temporal Key (PTK) from the PMK. Besides the PMK, other inputs for PTK creation are a random number used once (nonce) from the AP, a random nonce from the client, the AP's MAC, and the station's MAC. The ability to derive different PTKs from the PMK at different times keeps encryption keys from being reused too much.

Wireless devices then derive five more keys from the PTK. You do not need to memorize the names of these keys, but you should understand their basic purpose:

- Temporal Key (TK), which the client and AP will use to encrypt unicasts in this session
- MIC transmit key, which the AP will use to provide integrity for data that it transmits in this session (client will use this key to check the integrity)
- MIC receive key, which the client will use to authenticate data that it transmits in this session (AP will use this key to check the integrity)
- Key Confirmation Key (KCK), used during the 4-way handshake
- Key Encryption Key (KEK), used during the 4-way handshake

WPA Four-Way Handshake

- 1 Exchanges messages for generating PTK
- 2 Proves knowledge of the PMK
- 3 Distributes an encrypted GTK to the client
Used to encrypt/decrypt broadcasts and multicasts
- 4 Sets first initialization vector (IV)
IVs help prevent key reuse



MOD 12- 7

The four-way handshake provides several functions.

First, it lets the client and AP exchange the nonces, which are needed to derive the PTK from the PMK.

Second, it enables the station and AP to prove that they know the PMK without transmitting the PMK. They do this by using the KCK to create a message integrity code (MIC), which authenticates the third and fourth packets in the handshake. The authentication will only check out if the device knows the correct KCK, which demonstrates that it knows the PMK.

Third, during the handshake, the AP also distributes a groupwise transient key (GTK) to the client, encrypting this key with the KEK. The GTK is shared by the AP and all clients connected to the WLAN. The AP uses it to encrypt broadcasts and multicasts.


Finally, the handshake enables the AP and client to agree on the first initialization vector (IV) for their encrypted messages.

Open WLAN

Typically used for guest networks
Often with captive portal authentication

Sometimes used with MAC-Auth
Minimal authentication for devices with:

- No user interface
- No 802.1X support



802.11 connection **UP**



MOD 12- 8

An open WLAN simply allows clients to connect. The client passes through 802.11 open-system authentication, which permits all clients, and then 802.11 association. No further steps occur; the 802.11 connection is up, and the client can receive an IP address and start communicating. As you can see, this is a fairly simple process.

After that, though, layer 3 authentication often takes place. For example, companies often implement captive portal authentication on open networks. Companies also sometimes combine open WLANs with MAC-auth to provide minimal authentication for devices with no user interface or 802.1X support.

Opportunistic Wireless Encryption

No change in user experience


- Select and connect

Protection against eavesdropping

- Auto-negotiation of PMK with non-authenticated Diffie-Hellman (DH)
- Unique, secret PMK per-session

Users still susceptible to MITM

- If hacker poses as AP




Now you examine how Opportunistic Wireless Encryption (OWE), defined in RFC 8110, enhances open WLANs. OWE is not technically part of WPA, but instead rolled into the 802.11 standard. However, OWE uses WPA concepts like the key hierarchy and 4-way handshake.

An AP's beacons for an OWE WLAN indicate that it supports OWE. When a station connects to the WLAN, it passes through open-system authentication as usual. But the 802.11 association frames include extra TLVs, which carry an unauthenticated Diffie-Hellman (DH) exchange to derive a PMK. The station and AP then conduct the four-way handshake to derive all the keys for encrypting and authenticating future data. From the user's viewpoint, an OWE SSID works like earlier forms of open SSIDs. The user clicks connect, and the connection happens without the user having to submit any credentials. However, the traffic is encrypted and authenticated (digitally signed).

OWE provides best effort security—not true security. Remember that unauthenticated Diffie-Hellman exchanges are subject to man-in-the-middle (MITM) attacks, so a savvy hacker can still use a honeypot AP to lure clients to connect to it and generate a PMK with it. However, OWE does protect traffic against more casual snooping. Even though anyone can connect, the PMK for every session is unique and protected.

WPA/WPA2-Personal



You will next look at WPA/WPA2-Personal WLANs. A client that wants to connect to a WPA/WPA2-Personal WLAN passes through open-system 802.11 authentication and then 802.11 association. It then conducts the four-way handshake. Remember that part of the four-way handshake is proving knowledge of the PMK. For WPA/WPA2-Personal, the PMK is expanded out in a predictable way from the passphrase configured for the WLAN. Thus the four-way handshake proves that the client and AP are using the same passphrase, providing implicit authentication. If the four-way handshake is successful, the connection comes up.

WPA/WPA2-Personal Issues

PMK derived directly from admin-selected passphrase

- Subject to offline dictionary attacks from captured handshakes

PMK shared by all users

- Passphrase often shared publicly
- Easy to derive another client's PTK



MOD 12- 11

WPA2-PSK has two main, related flaws. First, the PMK and all other keys ultimately derive from a human-selected passphrase. Humans rarely select good passphrases because they cannot remember good passphrases. So, it is important that any key that derives from a passphrase cannot be subject to a brute force dictionary attack. However, WPA2-PSK is subject to brute force dictionary attacks!

An attacker can capture a valid handshake between a client that knows the password and the AP. Two of the handshake messages includes a MIC, which is a hash of the plaintext message with the KCK (one of the keys derived from the PMK). An attacker takes this data offline and uses the massive parallel processing power now available with GPUs and hacking farms to simply try every possible passphrase until one produces keys that creates the right MIC.


Even if the passphrase is not a common word and includes special characters, if it is 8 characters or under, processing power is high enough now that brute force attacks can relatively quickly crack the PSK. While companies could attempt to use very long and random strings, doing so is not feasible in practice.

Another issue with WPA2-Personal is that the password is often freely shared. While every client technically has its own PTK, in practice, anyone can easily derive another client's PTK. You just need the client's and AP's MAC addresses and the nonce values which are transmitted in plaintext during the handshake. This means that anyone who knows the passphrase can decrypt other devices' traffic. While you could hope that no one trusted with

the passphrase would do this, that is often not the case.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

WPA3-Personal



Instead of using a passphrase as the source of a PMK, WPA3-Personal uses Simultaneous Authentication of Equals (SAE) to generate a different PMK each time that a station connects.

The AP advertises SAE as the authentication method for a WPA3-Personal SSID in its beacon and probe response frames. This is a new 802.11 authentication method (value 3); WPA3-Personal is the only security option that uses an 802.11 authentication method other than open-system. (Shared-key is a deprecated option that is insecure and should never be used.)

The 802.11 SAE process involves four frames. You do not need to know the precise exchange but should understand the high-level concepts. SAE is based on the Dragonfly Key Exchange defined in IETF RFC 7664. This key exchange process uses some of the same cryptographic concepts as Diffie-Hellman and can use either Finite Field Cryptography (FFC) or Elliptic Curve Cryptography (ECC). ArubaOS uses group 19, a 256-bit Elliptic Curve group. The added twist, distinguishing SAE from Diffie-Hellman, is that the peers can only derive the same shared key—the PMK—if they also know the same passphrase.

After going through SAE, the client and AP have authenticated each other and derived a PMK that is unique to the session. The client associates. The familiar four-way handshake then occurs.

WPA3-Personal Benefits

Passphrases not subject to offline dictionary attacks

- New SAE process for every guess at passphrase

Enhanced security for short and weak passphrases

- Strong passphrases still recommended

No eavesdropping by users who know passphrase

- Unique per-session PMK known only by participants in that SAE exchange

Forward secrecy, if attacker learns passphrase:

- Can conduct MITM attacks
- CANNOT decrypt traffic she collected earlier

MOD 12-13

WPA3-Personal offers many improvements over WPA/WPA2-Personal.

The SAE exchange is not subject to offline dictionary attacks because it is computationally infeasible to derive the passphrase from the values transmitted in the SAE exchange. While an attacker can still take the four-way handshake offline, the PMK is no longer directly derived from the passphrase, so it is more resistant to a dictionary attack. And cracking the PMK will not be useful because it is different for each session. If an attacker wants to launch a dictionary attack, she must actually go through the SAE process and handshake for each guess of the passphrase, which makes the attack infeasible. Therefore, administrators can choose weaker passphrases while still maintaining good security. In fact, SAE is designed to work with a large pool of passphrases known to hackers. However, it is still best practice to choose good passphrases.


WPA3-Personal also protects users' communications from eavesdropping by other authorized users. Every session has a unique PMK, known only to the client and AP that established the session.

WPA3-Personal further provides forward secrecy, which means that a future compromise of the passphrase does not compromise data transmitted earlier. If an attacker learns the passphrase somehow—perhaps through a social engineering attack—they can conduct honeypot AP and MITM attacks. However, they cannot go back to encrypted data that they collected earlier and decrypt it. This is because the attacker was not part of the SAE

exchange that derived the PMK for those exchanges.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Aruba mPSK



Some companies have headless devices that do not support 802.1X authentication, so these devices must connect to WLANs in Personal mode. However, the devices might also lack support for WPA3, so WPA2-Personal, with all its flaws, is the only option for them. Aruba provides multi Preshared keys (mPSK) to enhance security for these devices.

A detailed explanation of mPSK configuration is beyond the scope of this course. The figure shows at a high level how the solution typically works.

- Step 1: Before attempting to connect the headless device, a user connects to the ClearPass Self Service Registration page on a different connection. In 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.
- Step 2: The user then configures the client device to use the passphrase.
- Step 3: 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.
- Step 4: However, at the point at which the client sends its 802.11 association request, the ArubaOS Aruba Instant AP or Mobility Controller performs MAC-auth to ClearPass. ClearPass responds with the passphrase assigned to the client in an Aruba-MPSK-Passphrase VSA. The ArubaOS device then sends the 802.11 association response and performs the 4-way handshake using the PMK derived from the passphrase individual to


the client device.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

WPA/WPA2/WPA3-Enterprise

Security based on 802.1X

- RADIUS-EAP authentication
- Unique PMK for each session



You will now move on to WPA/WPA2-Enterprise.

A connection to a WLAN that enforces this form of security begins with open-system 802.11 authentication and an 802.11 association. 802.1X then shuts down the connection to all but EAP traffic. EAP authentication proceeds as you examined previously, enabling the client and RADIUS/EAP server to authenticate each other. In addition to authentication, the EAP process also generates a TLS master secret (whether EAP-TLS, PEAP, or another EAP method is used). The client and RADIUS/EAP server each use this secret to derive the PMK.

Because the AP or gateway is the authenticator that needs the PMK for encrypting and decrypting wireless traffic, the RADIUS server sends the PMK to the authenticator in the RADIUS Access-Accept message. It encrypts the PMK with the RADIUS shared secret. Now both the client and the authenticator know the PMK for that session. The client and AP then complete the four-way handshake, and the connection is up.

WPA3-Enterprise: Overview

Same basic framework as WPA2-Enterprise

Supports WPA-2 AES CCMP: 128-bit key encryption, 64-bit key integrity check

Supports new operation modes, longer key lengths

AES-GCM-256 for encrypt/integrity – 256-bit keys

Optional CNSA mode

Mandates more secure algorithms



MOD 12- 16

WPA3-Enterprise continues to work much like WPA2-Enterprise. Stations use 802.1X with EAP to authenticate, and the PMK is derived during the exchange. However, WPA3-Enterprise has added support for new data encryption and authentication algorithms. WPA2-Enterprise used AES CCMP with a 128-bit key for encryption and a 64-bit key for ensuring integrity. WPA3-Enterprise still supports this mode, but also supports AES-GCM-256, which simultaneously encrypts and ensures integrity for data with 256-bit keys. The key new feature is Commercial National Security Algorithm (CNSA) mode.

Certificates



You can improve security with certificates.

CONFIDENTIAL - FOR TRAINING


What is a Digital or X.509 Certificate?

Bind an identity to a public key

Entity that vouches for the subject's identity; Certification Authority (CA)

Subject = Entity identified in the certificate; entity that owns the public/private keypair

Proof that issuer really issued this certificate



MOD 12- 18

Certificates provide key management: they tie a specific identity to a public key. In other words, they verify to everyone that the public key posing as “User1 key” really belongs to User1. And they do this without User2 having to install User1 public key on his machine in advance. This enables User1 and User2 to set up secure communications on the fly.

Digital certificates follow a standard format based on industry standard X.509, which is why you will sometimes hear certificates called X.509 certificates. The figure shows some of the key fields within an X.509 certificate.

The serial number uniquely identifies the certificate among certificates issued by the same CA. The validity period gives the certificate a specified period in which the certificate is valid. The longer that a device can retain a certificate, the less secure the certificate becomes. Perhaps the private key is compromised or the certificate owner is no longer approved for the rights that the certificate provides. Just as you must renew your driver's license or passport periodically, devices must renew their certificates periodically.

The issuer identifies who issued this certificate to the subject. This is a crucial component of certificate trustworthiness. Hackers could create a certificate with his own public key and any identity he wants to claim. But unless the certificate has been issued by a source that you trust, you would know to reject the certificate.

The subject identifies the entity that owns the certificate and the keypair associated with the

certificate.

The public key is the subject's public key.

A certificate also includes the issuer's signature. Much like watermarks help people to validate that the proper governmental body truly created an ID like a driver's license, the certificate signature validates that the issuer claimed by the certificate truly issued the certificate.


Just like specific governmental bodies are responsible for issuing documents such as driver's licenses and passports, Certification Authorities (CAs) are responsible for issuing digital certificates. We trust that the government has taken the proper steps to validate the passport holder's identity, so we trust the passport. Similarly, devices and applications trust that the CA has validated the subject's identity, so they trust the certificate issued by the CA.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Certificate Usage


Photo ID usage

- Prove ability to drive a vehicle
- Prove identity to board an airplane



Certificate Usage

- Data encryption
- Data signatures
- Certificate signing
- Secure Email
- **Server Authentication**
- **Client Authentication**



MOD 12- 19

You know that a certificate binds an identity with a public key (and implicitly with the private key paired with the public key). And you know that you can use private and public keys to sign and encrypt data. These are called key usages, and they are specified within the certificate. Certificate signing is another example of a key usage. Certificate signing is actually data signing, but for the specialized purpose of issuing certificates. Only CA certificates should have the certificate signing usage.

Also, consider a broader viewpoint: why the device must encrypt or sign data and how the device will use the certificate. People might use a driver's license to prove they can drive a car, or to prove their identity when boarding an airplane. Likewise, users and devices can use their certificates for multiple purposes, which are defined in extended key usages (EKUs), which can also be included in the certificate. For example, the public key in a certificate with the Secure Email extended key usage can be used to encrypt an email destined to the owner of certificate.


The most relevant EKUs for your purposes are Server Authentication and Client Authentication. To authenticate itself with a certificate, either a server or a client signs data with the private key. By signing the data, the device proves that it owns the private key associated with the certificate and is the source of the signed communications.

Note: It is optional to define key usage – certificates function without a usage defined. But this means the certificate can be used for anything! This widens the possibility of abuse. When defining how certificates shall be created, it is a best practice to lock down how they can be

used.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Example Certificate




In the figure you see an actual certificate, as shown on a Windows machine.

Some of the important fields include:

- Subject name, which identifies the owner of the certificate
- Issuing entity (the name of the CA that signed the certificate)
- Certificate serial number, which is unique to this certificate and issuing entity
- Validity dates (just like passwords should expire, secure certificates should expire too, forcing a renewal in which the entity generates a new keypair and once again proves its identity)
- Public key—Remember that the certificate includes the public key itself.

You can also look through these fields to see the key usages and EKUs.

CAs and the Public Key Infrastructure



MOD 12- 21


Now that you have looked at certificates, you examine the mechanisms for issuing those certificates.

Certificates exist within a public key infrastructure (PKI) hierarchy that enables devices to determine which certificates they should trust. Every certificate within the PKI is signed by an entity that vouches for the authenticity of the certificate holder's identity.

Root CAs are at the top of the hierarchy. They sign their own certificates, called root CA certificates. Root CAs are well-known entities that are widely believed to be trustworthy. For example, DigiCert and Comodo are root CAs. A company can also set up its own internal or private root CA, which is believed to be trustworthy within the company.

As the ultimate source of trust, root CAs must be rigorously secured. The root CA issues just a few certificates itself, issuing and signing certificates for intermediate CAs. This ensures minimal contact between the root CA and other entities. In fact, the root CA is typically kept offline for security purposes. The intermediate CAs, sometimes also called subordinate or signing CAs, are responsible for issuing and signing certificates for end entities such as Web servers, network infrastructure devices, and users. The end entity certificates are sometimes called server certificates, user certificates, or client certificates, depending on the type of entity to which they are assigned. They are also sometimes called leaf certificates because they are at the other side of the PKI tree from the root. A PKI could also have multiple tiers of intermediate CAs, as shown in the figure.

Chain of Trust



MOD 12- 22

You can also look at the PKI as a chain of trust, which flows back up through the intermediate CAs to the root CA. Devices and applications can trust the server's certificate because they trust the intermediate CA's certificate. In other words, to validate the server's certificate, applications check the signature in the certificate using the issuing intermediate CA's certificate.

In turn, devices and application trust the intermediate CA's certificate because they trust the root CA's certificate. Because the root CA certificate is self-signed, applications and devices must accept this certificate as trusted through another mechanism.

Typically, they are pre-configured to do so. This structure solves the practical problem of using asymmetric encryption. Rather than having to learn the public key for every entity with which you might want to communicate over a secure channel in advance, you need only learn the certificate, and included public key, for a few root CAs.

Cloud Authentication



Cloud authentication is convenient and scalable.

CONFIDENTIAL - FOR TRAINING

Cloud Authentication Key Benefits and Features

Quicker configuration

- Authentication
- Authorization
- Onboarding

Integration into cloud identity stores

- Google Workspace
- Azure Active Directory

Simpler user experience

- Client app supports a broad range of devices
- Mobile, desktop, laptop

Enhanced security

- AI-based client insights
- PassPoint technology




MOD 12- 24

It is critical to apply consistent security controls and ensure users have seamless access to apps and data in the office, at home, and on the go. Aruba Central simplifies this IT process with cloud-based authentication, extending its ability to deliver a single point of visibility and control over all network infrastructure and related security services. Featuring an easy-to-use interface and dashboard, Aruba Central makes it easy to onboard new clients, and to monitor and troubleshoot user connection issues.

You configure fine-grained policies in Aruba Central. Thus, when end users authenticate, they are authorized for appropriate network access. To help enhance security, you get MAC-based authentications and AI-based client insights, which capture and profile all connected devices. This gives you continuous, in-depth visibility into the connected client behavior, enabling a stronger security posture than by using MAC addresses alone.

Cloud Authentication & Policy for Central



Cloud authentication and policy creation in Aruba Central involves four main components:

Cloud Identity

Aruba Central enables end users to connect to Wi-Fi networks securely and automatically. The application integrates with a company's existing cloud identity store such as Google Workspace or Azure Active Directory to authenticate user information and assign appropriate access. You get frictionless Wi-Fi connectivity for known and new users, visibility into traffic patterns, and fine-grained policies to control access.

End-User device onboarding

Easily configure employee devices for seamless wired and wireless connections. Managed through Aruba Central, end users are authenticated through the company's cloud identity store with an enrollment link provided on Aruba Central.

Using company credentials to log in, the user is redirected to the identity store for authentication. Configure client devices using Aruba Onboard, a client app that is supported on macOS, Windows, iOS, and Android operating systems. With the Enterprise PassPoint Profile installed on the device, when user walks into network range, the device automatically connects with appropriate access, based on your Aruba Central configuration. The client app provides automatic renewals: no additional onboarding steps and upkeep from the end user. And you can change and update policies at any time.

Automated onboarding for guests and visitors

Aruba Air Pass is the industry's first seamless cellular roaming solution designed to unify enterprise and mobile network experiences. You manage Air Pass in Aruba Central, extending the 5G experience to the enterprise using Aruba Wi-Fi technology. Air Pass uses pre-negotiated agreements with mobile network operators (MNOs) powered by PassPoint. Air Pass gives guests and visitors an uninterrupted Wi-Fi experience in indoor venues with the added benefit of enterprise-grade security.

Network management

Covered throughout the course.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY



Knowledge Check

Self-check on key learning points




Its time for a knowledge check.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Question #1

Match the authentication and encryption options to the appropriate WLAN security option.



Knowledge Check

MOD 12- 27

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Question #2

It is time to play **True** or **False**!

- During the WPA four-way handshake, the Pairwise Master Key (PMK) is used to encrypt data, while the Message Integrity Check (MIC) keys are used for hashing – a frame integrity check. **FALSE**
- Opportunistic Wireless Encryption (OWE) adds encryption to an open SSID that eliminates Man-in-the-Middle (MITM) attacks. **FALSE**
- When using WPA2-Personal as your WLAN security option, your security could be subject to offline dictionary attacks. WPA3-Personal mitigates such attacks. **TRUE**
- Aruba provides a multi Pre-Shared Key (mPSK) feature to improve security for devices that lack 802.1X authentication support. **TRUE**

Q

Knowledge Check

MOD 12- 28

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Question #3

Which of the statement below help you accurately describe digital certificates to your colleagues? (Select three.)

- a. Certificates can serve to prove user and device identity, and to provide key management.
- b. Only certificates that are properly signed by a Certificate Authority (CA) should be trusted.
- c. When creating certificates, you are not required to specify how it should used, but it is a good idea to do so.
- d. Entities must validate all certificates by communicating directly with the Root Certificate Authority (CA).

a

Knowledge Check

MOD 12- 29

A,B,C

Question #4

Which options below describe significant features of Aruba Central's Cloud Authentication? (Select all that apply)

- a. It offers a single point of visibility and control over all network infrastructure and related security services.
- b. It makes it easy to onboard, monitor, and troubleshoot client connections.
- c. It integrates in cloud identity stores like Google Workspace and Azure Active Directory.
- d. Artificial intelligence (AI) Insights give you continuous, in-depth visibility into connected client behavior.

a

Knowledge Check

MOD 12-30

A,B,C,D



Lab Activity

Consider the lab activity.

CONFIDENTIAL - FOR TRAINING

Lab Tasks

12.01: WLAN Security Features

Task 12.1

- Create PSK WLAN

Task 12.2

- Test PSK Access with IoT air sensor

Task 12.3

- MPSK with Differentiated Access Control

Task 12.4

- Test MPSK Access

Task 12.5

- Troubleshoot clients connecting to PSK

Module 13 – Monitoring and Maintenance



MOD 12- 32

The figure provides a brief review of lab tasks. Please see your lab guide for details. When complete, you will continue with Module 13 – Monitoring and Maintenance.


Rev.22.41

aruba
a Hewlett Packard
Enterprise company

Monitoring and Maintenance

Aruba Campus Access
Fundamentals

Hello good people, and welcome back. This is Module 13, Monitoring and Maintenance.



Objectives

- Use Aruba Central monitoring capabilities
- Identify LED status
- Perform firmware upgrades
- Enable SNMP on devices
- Describe AI Insights
- Describe Alerts & Reports
- Explain UXI

Maximizing situational awareness means understanding your network health, who's connected, the health of those connections, and more. This awareness prepares you to more quickly diagnose issues, often proactively, before they impact users.

After completing this module, you will be able to use Aruba Central monitoring capabilities, identify device LED status, perform firmware upgrades, and enable SNMP on your devices. You will also be able to describe AI insights and explain Aruba's User eXperience Insights (UXI) deployments, which can really improve your awareness of end user performance issues, and quickly nail down the cause.

Monitoring with Aruba Central

This section is focused on monitoring with Aruba Central.

CONFIDENTIAL - FOR TRAINING

Overview - Summary



MOD 13- 4


The figure shows typical Overview page info: Network Health, WAN Health, Summary, Wi-Fi Connectivity and AI Insights.

The Network Health information includes a site map, with locations in various colors.

The Summary info shows bandwidth usage, client count, top APs by usage, top five charts, top IAP clusters by usage and clients, application usage as well as WLAN network detailed from a selected group. By default, Central plots the graphs for a time range of three hours.

You will learn more about accessing this information during lab activities.

Overview – Network Health




MOD 13- 5

The figure shows typical network health dashboard information. This Summary view shows vital network information of individual sites on cards mapped by geographical location. Sites are marked with location pins—green pin for a site with no issues and red pin for a site with potential issues. Hover over a pin to view more details about a site in a card view. When there are multiple sites in a region that cannot be displayed as individual pins on the map view, they appear as numbers on a round pin in the condensed view. The pin color is red if it has one site in the down status and green if all the sites are up. Click the pin to zoom in and view the sites on the map.

There is also a List view, which is not shown here. With a list view, you get a global network report in a list sorted according to individual sites.

Overview – Wi-Fi Connectivity




MOD 13- 6

Aruba Central has a Wi-Fi Connectivity page, formerly known as Clarity, to see an overall view of client connection details: those that are connected to or tried to connect to each connection phase. The connection phases include information concerning wireless association, authentication, DHCP, and DNS.

You will explore this and other information during the lab activities.

Site - Topology



MOD 13- 7

The Topology pane gives you a topographical view of your network. The topology is created based on the LLDP neighbor list of managed devices. You need to select a filter based on labels and sites to see the topology.

In this example, you can see three IAPs and one switch. You can see the connections because the IAPs and switch are both at the OKC site. If the switch or IAP were in different sites or labels you would only see the devices, but no connections.

If you select a device, then the device information is on the right-hand side of the screen. You can use the filter option to display only the IAP, switches, branch controllers, or a combination of them.

DEVICE

AP MODEL

AP-515

MAC

e8:26:9c:c8:ca:aa

UPTIME

2 Hours 13 Seconds

FIRMWARE VERSION

10.3.0.0-10.3.0.0-beta_82025

Last Updated on Nov 04, 2021, 13:39

DUAL SELECTION

Dual Band

POWER NEGOTIATION

802.3 at

LABELS

First Floor

SITE

Sao Paulo

COUNTRY CODE

US

SERIAL NUMBER

USK2KD53N7

LAST REBOOT REASON

AP Reboot

CONFIGURATION

Synchronous

Last Config

POWER DISTRIBUTION

13.11 W

GROUP

Branch C

LEDs on AP

☒ Blue
 ☒ Blue
 ☒ Blue

NETWORK

ETH0	SPEED (Mbps) / DUPLEX	VLAN
Up	1000 / Full	Trunk (all) LLDP Detail

ETH1	SPEED (Mbps) / DUPLEX	VLAN
Down	-	-

CURRENT UPLINK

UPLINK CONNECTED TO

CHANNEL UTILIZATION

Legend: ☒ Transmitting ☒ Receiving ☒ Non-Wi-Fi Interference

ows you to view managed devices. They are divided into gateways. By selecting the device type, a list of devices and offline devices. Pick a device to drill down and get more information. Of course, you can easily get information for any device in more depth during lab activities.


LED Status



Now its time to look at all the pretty lights!

CONFIDENTIAL - FOR TR

LEDs on AP



MOD 13-10

The LED displays located on the front panel of the access point indicate the system status of the access point.

System Status LED



	Off	Device Powered off
	Green: solid	Device ready, fully functional, no network restrictions
	Green: blinking¹	Device booting, not ready
	Green: flashing off²	Device ready, fully functional, uplink negotiated in sub-optimal speed (<1Gbps)
	Green: flashing on³	Device in deep-sleep mode
	Amber: solid	Device ready, restricted power mode, no network restrictions limited PoE or IPM restrictions applied
	Amber: flashing off	Device ready, restricted power mode, uplink sub-optimal speed (<1Gbps)
	Red	System error condition - Immediate attention required

MOD 13- 11

The figure show system status LED operation. Here's an explanation of what is referred to as blinking or flashing:

1. Blinking: one second on, one second off, two-second cycle.
2. Flashing off: mostly on, fraction of a second off, two-second cycle.
3. Flashing on: mostly off, fraction of a second on, two-second cycle.

Radio Status LED



Off

Device powered off or radio disabled



Green: solid

Radio enabled in access (AP) mode



Green: flashing off²

Radio enabled in uplink or mesh mode




Amber: solid

Radio enabled in monitor or spectrum analysis mode

MOD 13- 12

The Radio Status LED table is applicable to 2 GHz, 5 GHz, and 6 GHz labels.

Switch LEDs



1 - Switch port LEDs

4 - Stk LED

7 - Global Status LED

2 - Back Module status LED

5 - Reset button

8 - Mode select button

3 - Speed mode selected LED

6 - Unit Identification LED

9 - OOBM port LED



MOD 13- 13

The figure shows the locations of typical switch LEDs.

Switch LEDs Meaning

Switch LEDs	Function	State	Meaning
Back LED	Status of modular components installed in the back of the chassis (not applicable for 6300F switches)	On - Green	Normal
		Slow Flash - Amber	Fault in one of the modules in the back of the chassis
PoE LED	Indicates Port LEDs are showing PoE information (not applicable for non PoE switches)	Off	PoE mode not selected
		On - Green	PoE mode selected
		Slow Flash - Amber	One or more ports experiencing PoE failure. PoE mode not selected.
		On - Amber	One or more ports experiencing PoE failure. PoE mode selected.



MOD 13- 14

Here you see operation of the back and PoE LEDs.

Switch LEDs Meaning

Switch LEDs	Function	State	Meaning
Spd LED	Indicates Port LEDs are showing speed information	Off	Speed mode not selected
		On - Green	Speed mode selected
		Not Implemented	No fault defined
Stk LED	Indicates Port LEDs are showing stacking mode information	Off	Stacking mode not selected
		On - Green	Stacking mode selected
		On - Amber	One of the switch is experiencing stacking failure. Port LEDs will be used to indicate the number of the stacking member.
		Slow flash Amber	A port has a stacking failure. Stacking mode not selected



MOD 13-15

The state of the switch port LEDs is controlled by the LED Mode select button. The current view mode is indicated by the mode LEDs next to the button. To step from one view mode to the next, press the button to cycle through the different modes (speed or Spd and stacking or Stk).

Switch LEDs Meaning

UID LED	Locator LED	Off	User configured the located LED : OFF
		On/Flash Blue (for 30 min)	User configured the locator LED: On/Flash
Global Status Indicator LED	Overall status of the product	Flash - Green	Self-test in progress during UBOOT and SVOS
		On - Green	Successfully booted AOS-CX
		Flash - Amber	Recoverable faults (e.g. fans, PSU fault)
		On - Amber	Critical faults (e.g. exceed temperature limit)
		Snoring (Dim - Bright periodically)	System is in hibernation



MOD 13- 16

And here are the UID and Global Status Indicator LEDs.

Switch Reset Buttons

To Accomplish this:	Do this:	This will happen:
Soft reset	Press and release the Reset button.	The switch operating system is cleared gracefully. The switch then reboots and runs self-tests.
Hard reset	Press and hold the Reset button for more than 5 seconds, then release.	The switch reboots, similar to a power cycle. A hard reset is used, for example, when the switch CPU is in an unknown state or not responding.



MOD 13- 17

The Reset button is recessed from the front panel; this design protects it from being pushed accidentally. The button is accessible through small holes on the top of the front panel. Use pointed objects, such as unbent paper clips, to press it. There are two types of resets: soft and hard, see figure for explanation.



Firmware Upgrades

You really should keep up with those firmware upgrades. Aruba Central makes it easy.

CONFIDENTIAL - FOR TRIAL

Firmware Management and Upgrade Overview

Access Points (2)						SET COMPLIANCE	UPGRADE ALL
Name	Site	Firmware Version	Recommended Version	Upgrade Status	Compliance Status		
VC1-B1	Building 1	8.9.0.0_81161	8.6.0.9_79813	Firmware up to date	Not Set		
VC2-B2 (2)	Building 2	8.9.0.0_81161	8.6.0.9_79813	Firmware up to date	Not Set		

Switches (3)										SET COMPLIANCE	UPGRADE ALL
Name	Family	Site	MAC A...	Model	Firmware...	Recommended Ve...	Upgrade ...	Compliance Status			
2930F-8P	AOS-S	Sao Paulo	f4:03:43:07:c3...	Aruba2930F-8G-PoE+-25FP+ S...	16.10.0005	16.10.0015	Newer firmware av...	Not Set			
6200	CX	Sao Paulo	b8:d4:e7:b4:1...	6200F 24G CL4 45FP+370W S...	10.08.1021	10.07.0020	Firmware up to date	Not Set	Not Set		
6300	CX	Sao Paulo	b8:d4:e7:0e:a...	6300F 24G CL4 PoE 45FP56 Sw...	10.08.1021	10.07.0020	Firmware up to date	Not Set			



MOD 13- 19

Aruba Central eases Firmware management and upgrades. You can see the latest supported firmware version for your devices. You can see the upgrade status for all the devices and if there is a new version available, with an option to upgrade all devices, or choose specific devices to upgrade. Of course, you can choose from the available versions used for the upgrade.

You can update now or set a certain time for the update. You can also update per site: updating one site without affecting the other sites.

AP Firmware Upgrade and Compliance

Standard Upgrade

All selected APs upgrade simultaneously, may disrupt connectivity
Not recommended for production network during work hours
OK to use during a planned maintenance window

Live Upgrade

APs upgraded in sequence, based on AirMatch partitions
Minimized impact in a production environment
Required advanced licenses on APs

Firmware Compliance

Ensures that devices in a group all run a specific OS version
Choose APs (VCs), switches, or gateways
See list of supported firmware versions per in-group device

MOD 13- 20

When upgrading access points, there are two upgrade types:

- A standard upgrade will upgrade all the selected APs simultaneously, which may disrupt connectivity. Therefore, Aruba does not recommend a standard upgrade in a production network during the working hours. However, a good use case for standard upgrade is to upgrade APs in a maintenance window, since it will complete the upgrades faster than live upgrades.
- A live upgrade will upgrade APs based on AirMatch defined partitions, minimizing the impact in a production environment. The upcoming sections will explain the live upgrade in detail. Notice that live upgrade feature requires advanced licenses applied on the APs.

The firmware compliance ensures that devices in a group all run a specific OS version. Choose either the access points (virtual controllers for IAPs), switches, or gateways. Firmware compliance also allows you to view a list of supported firmware versions for each device in a group.


AP Live Upgrade

Seed AP

- Only seed APs download images from cloud
- Others download image from seed AP

Seed AP selection

- APs are L2 connected to its seed AP
- AP is same model as its seed AP
- Max 20 client APs assigned to a seed AP
- Seed AP randomly picked




MOD 13- 21

A live upgrade should not disrupt clients that are associated with the APs. The live upgrade interfaces with AirMatch service to partition all the APs into smaller sets that will upgrade in parallel. AirMatch partitions APs based on RF neighborhood data. APs partitions are used for clients to roam when the APs are undergoing an upgrade. A set of APs, known as seed APs, is selected. Only these APs download images from the cloud, the remaining APs download images from their designated seed AP. This saves time and bandwidth since the image is only downloaded to the company once, and then distributed by the seed AP(s).

With seed AP group selection, the client AP is layer 2 connected to its seed AP. The client AP is the same model as its seed AP and no more than 20 client APs are assigned to a seed AP. Non-seed APs use HTTPs to download the image from the seed APs.

Note: the seed APs are selected randomly.

AirMatch Partitions & Live Upgrades



AirMatch optimizes the RF coverage patterns and assigns unique RF channels to physically adjacent APs. This optimizes the RF coverage overlap between APs. In other words, APs set to the same channel will not be physically adjacent. This is especially true for the 5 GHz band, which has several non-overlapping channels. If the 5 GHz band is not available, then AirMatch uses the 2.4 GHz band.

AirMatch partitions APs into groups, based on their AirMatch-assigned 5 GHz channel. In the figure, AP1 and AP4 are both on channel 36, and assigned to Partition 1. AP2 and AP5 are on channel 52 and assigned to Partition 2. AP 3 is on channel 40 and is assigned to Partition 3.

If all APs in a partition reboot simultaneously, the surrounding APs (in other partitions) will provide the necessary RF coverage. Of course, this assumes a very healthy RF coverage pattern with ample overlap.

Sites can also make partitions. Therefore, APs in the same RF channel but in different sites would belong to separate partitions. For this reason, it is important to properly select your sites. If you have a campus with six buildings all in one site, then all APs in one RF channel would belong to one partition. If each building is its own site, then the AP partitions on the same RF channel would be limited to the building.

SNMP




Let's look at the open-standard Simple Network Management Protocol (SNMP).

CONFIDENTIAL - FOR TR

SNMP

SNMP agents running on devices




MOD 13-24

SNMP is an industry-standard protocol that allows you to manage and monitor a variety of network devices from a central location. Using an SNMP application, you can configure network devices and apply consistent security and management policies to all the devices on your network. You can also monitor SNMP compliant devices for conditions requiring administrative attention: for example, to determine if a server has gone down, a LAN or WAN link is no longer available, or a link has become congested. You can track device uptime, link states, and many other device information variables. When a problem occurs, SNMP compliant devices send an alert message to one or more designated SNMP servers, and you can then take steps to resolve the problem.

SNMP agents are processes that run on network devices such as switches or APs. These agents communicate with management software that is using SNMP for network device monitoring.

SNMP MIB



MOD 13- 25

SNMP management information bases (MIBs) are the heart of SNMP and allow for the storing and sharing of configuration and event-related information. This section introduces MIBs and how SNMP managers and devices interact with each other using the MIB information.

SNMP Manager

The SNMP manager is typically a dedicated server that communicates with the SNMP managed device (commonly called the agent). It often has a GUI interface, so administrators can view and change parameters for many network device from a single interface. The manager uses several message types to gather and modify network device parameters, Some of its key functions include:

- Queries agents
- Gets responses from agents
- Sets variables in agents
- Acknowledges asynchronous events from agents

Managed devices

The routers, switches, firewalls, wireless controllers, access points, and any other network devices that need to be centrally monitored, managed, and configured are referred to as managed devices (or agents). Internally, these managed devices have two main SNMP

components.

The SNMP Agent collects information and communicates with the SNMP Manager about device configuration and status. It stores and retrieves this information in a Management Information Base (MIB).

The MIB is a hierarchical database that stores device status, statistics, and configuration information. This is the information that is exchanged between the SNMP Manager and the managed device.

Communication and ports

Usually, SNMP Manager to Agent communication is a pull mechanism. When this device requires to get information from the managed device it requests the information using simple commands such as GET, GET NEXT, GET BULK. Changes in the device are done using the SET command. This communication uses UDP 161.

The exception to the pull mechanism is when the managed device has important information that cannot wait for the manager's period polling. "Newsflash!! My CPU utilization is dangerously high". In cases like this, the managed device pushes information to the manager without being asked, using an SNMP Trap. SNMP Traps use UDP 162.

CONFIDENTIAL - FOR INTERNAL PURPOSES ONLY

SNMP Versions

SNMPv1

- Community-based security
- Limited method to get managed device information

SNMPv2c

- Community-based security, that improves SNMPv1
- Protocol packet types, Transport mapping, MIB structure elements

SNMPv3

- Better security mechanisms
- Authentication, Message integrity, Encryption



MOD 13-26

The figure summarizes SNMP versions, as described below.

SNMP version 1. First version of the SNMP protocol and uses a community-based security mechanism. The community string is like a simple pre-shared passcode. If the managed device has a community string of public, then any SNMP manager that can reach the device can access the agent's MIB, as long as it sets its community string to public.

SNMP version 2c. This is the revised protocol, which includes enhancements of SNMPv1 in the areas of protocol packet types, transport mapping, MIB structure elements. This version also uses the community-based security mechanism.

SNMP version 3. Provides a higher level of security, including message integrity, authentication, and encryption. SNMPv3 is recommended.

You will learn how to configure SNMP during lab activities.

Note: Because SNMPv1 and v2 send the community strings in clear text, it is highly recommended not to use them; instead, SNMPv3 should be used.


AI Insights



You are going to be impressed with the AI Insights feature.

CONFIDENTIAL - FOR TRIAL

Aruba Central Cloud AIOps Portfolio



Aruba Central includes a full-service AIOps solution that automates common troubleshooting activities, reducing IT support tickets and associated costs. Aruba's AIOps solution is based on machine learning models that are consistently trained with network performance data collected from tens of thousands of Aruba customers across every vertical, market segment, and network type.

The AIOps solution's core components include:

- AI Insights which automatically diagnose an array of common network-impacting issues by using dynamic, per-site baselines that are continuously tuned as conditions change: no manual setup or service level threshold adjustments required. Built-in anomaly detection highlights the severity and impact of issues as they occur, helping IT pinpoint root cause and proper remediation steps with 95% accuracy.
- AI Search is a natural language processing (NLP) engine of the Aruba Central data lake that points to solution guides, troubleshooting tips, and more. Whether operators are looking for best practices on device configurations or need to isolate a problem impacting a specific user, AI Search provides fast, interactive responses that simplify day zero to day two operations.
- AI Assist uses event-driven automation to collect diagnostics, post them to a shared location, and even notify Aruba TAC for proactive customer support.

What AIOps Is and Isn't

It IS

Added to a management solution to:

- Automatically collect needed data to resolve persistent problems
- Establish normal behavior and deviations that impact users, devices, apps, sites
- Intelligently troubleshoot using Natural Language Processing Search results

It IS NOT

A legacy management or a way to:

- Find if one AP, switch port, or WAN link is down
- Or build a device configuration



MOD 13- 29

AIOps doesn't replace your management solution: it augments it. Aruba Central was designed with AIOps being part of it. Traditional configuration, management features, troubleshooting, and reporting tools are still there. The difference is that AIOps uses a wealth of telemetry to find ongoing patterns that point out potential problems.

AI Search

Built-in Natural Language Processing

Simple keyword queries

Interactive Responses

192.168

Refine search: Clients - Devices - Troubleshoot - UCC Calls - Threats

We found the following results:

Clients

IP ADDRESS: 192.168.10.7

TME-Latitude-E5420

MAC ADDRESS: a9:19:21:68:ef

Joe-iPhone

IP ADDRESS: 192.168.163.114

24:77:03:92:9c:98

IP ADDRESS: 16.192.168.13

Desktop-DQC9LUS

NAME: Android-E192168decd7e8...

Android-E192168decd7e8...

IP ADDRESS: 192.168.163.114

24:77:03:92:9c:98

LOCATION: Building B192.168

HarmonyHub

IP ADDRESS: 192.168.1.70

Galaxy-S5-Neo

CONNECTED SINCE

Apr 26, 4:59 PM

CONNECTED TO

Aruba7005_2F...

STATUS

Connected with fair performance

VIEW ALL 23 CLIENTS

Large number of neighbor APs

A large amount of neighbors APs have been detected on subnet 192.168.10.0/24. These neighbors may be creating high amount of traffic and interferences, which could lead to high RF airtime utilization in the area.

VIEW

MOD 13- 30

The AI search feature is an AI-based search engine that not only helps you find problems users are experiencing, but it also points you to Aruba content related to the problem, and where you can make the changes in a configuration. It's the only Neural Language Processing (NLP) solution that responds with day zero, day one, and post deployment tips in an interactive way.

AI Search - Creating a WLAN Example

Sample query & AI suggestions

Q How to configure a SSID?

AI Search found the following results:

SSID Configuration
Please click 'View' to navigate SSID configuration page.

Configuring an Employee Network
Please click 'Guide Me' to initiate the tutorial.

Support for Multiple PSK in WLAN SSID
Also, ensure that the RADIUS server configured for the wireless SSID profile is not an internal server. Click the Config icon. The tabs to configure the APs are displayed. Click WLANs tab. The WLANs detail page is displayed. Click + Add SSID to create a new SSID. To modify an existing SSID, click the Edit icon.

Configuring Cloud Auth Server in a WLAN SSID
The tabs to configure the APs are displayed. Click the WLANs tab. The WLANs details page is displayed. In the WLANs tab, click + Add SSID. The Create a New Network page is displayed. In the General tab, enter a name in the Name (SSID) text-box. Under Advanced Settings, click the Security tab. Select the security type and key length. Click the Save icon.

Guest Access Dashboard
Guest Count by Authentication Number of client devices based on the authentication type configured on the cloud guest SSIDs. Guest Count by SSID Number of guest connections per SSID. Client Type Type of the client devices connected on the guest SSIDs. Send Feedback ...

Manage

Overview

Devices

Clients

Guests

Applications

Security

Analyze

Alerts & Events

Audit Trail

Tools

Access Points

Switches

Gateways

WLANs

Access Points

Radios

Interfaces

Security

Services

Wireless SSIDs

NAME	SECURITY

+ Add SSID

READ

MOD 13- 31

AI Search provides fast, interactive responses that simplify operations from day zero (provisioning the network) to day two (maintaining the network). The example in the figure a WLAN creation search. Notice that the NLP engine suggests many topics related to the search, including a link to the WLAN configuration page, a guided configuration tutorial, and related articles.

AI Search - AP Example

APs

AI Search found the following results:

Access Points

- IP: 10.1.145.12
○ AP1-B2
- IP: 10.1.145.11
○ AP-1
- IP: 10.1.145.13
○ **AP-3**
- IP: 10.1.145.10
● AP-2

Client Details

STATUS	NAME
Down	AP-3
MODEL	MAC
305	c8:b5:ad:c7...
SERIAL	IP
CNDBJSS0FX	10.1.145.13
FIRMWARE VERSION	
10.3.0.0-10.....	

Insights

INSIGHTS
● 0 ● 0 ● 0

ADDITIONAL ACTIONS

Actions

Configure
Network Check Locate
Events Clients
Configure Group

Was this helpful?

MOD 13- 32

One example of the AI Search feature is to search managed devices. In this example, a search for APs returned all the APs the Aruba Central customer manages. The AP box expands by hovering the mouse over an AP, showing the AP status, name, model, and other device information. At the bottom, you will find additional actions like device configuration, other Central pages to perform additional actions, like the Network check button to perform troubleshooting or the Clients button to access clients page, and others.

AI Search - Client Search Example

The screenshot displays a web interface for an AI search tool. At the top, a search bar contains the text 'clients'. Below the search bar, a message states 'AI Search found the following results:'. The results are categorized under 'Clients' and are presented in three columns. The first column shows a detailed view for a client with IP 10.1.141.11, identified as a 'contractor'. This view includes fields for MAC (14:7d:da:49:44:xx), IP (10.1.141.11), SITE (Sao Paulo), HOST NAME (Henriques-MBP), and OS (Mac OS). It also shows a 'CONNECTED' status, 'WIRELESS' client type, and 'INSIGHTS' with three colored dots (red, yellow, green) and a '0' next to each. At the bottom of this card are buttons for 'Network Check', 'Live Events', 'Events', and 'Locate'. The second column shows a client with IP 10.1.141.12, identified as 'admin'. The third column shows a client with IP: null and MAC 00:0c:29:64:e9:6d. At the bottom left of the interface is a feedback section with the text 'Was this helpful?' and two thumbs-up/down icons. The bottom right corner of the interface displays 'MOD 13- 33'.

clients

AI Search found the following results:

Clients

IP: 10.1.141.11
contractor

MAC 14:7d:da:49:44:xx NAME contractor

IP 10.1.141.11 STATUS CONNECTED

SITE Sao Paulo CLIENT TYPE WIRELESS

HOST NAME OS

Henriques-MBP Mac OS

INSIGHTS
● 0 ● 0 ● 0

ADDITIONAL ACTIONS

Network Check Live Events Events

Locate

IP: 10.1.141.12
admin

IP: null
00:0c:29:64:e9:6d

Was this helpful?

MOD 13- 33

The figure shows results for an AI search.

AI Insights - AP to Stack Mapping Example



Stack 1 / Switch 1



Stack 1 / Switch 2




MOD 13- 34

One of the possible AI insights from Aruba Central is the AP-to-stacking map. The figure displays a site where the wireless coverage was done correctly, providing good coverage for the entire floor. However, Aruba Central AI noticed that neighbor APs are connected to the same switch in the switch stack. If one switch fails, an entire part of the floor may get low or no Wi-Fi coverage!

AI Insights - Per Stack RF Resiliency Example

Before

Resiliency 76%



After

Resiliency 84%



MOD 13- 35

Following the AI insight and physically moving APs connections from one switch to another will provide better coverage in case of failure or a reboot of one switch in the stack.

AI Assist

Problem

- Sporadic failures (e.g. port flaps) difficult to see and diagnose by humans
- Issue happens then “goes away”

Solution

- AI Assist’s event-driven automation
- Collects troubleshooting data and posts to shared location
- Alerts admin



The screenshot shows a 'Notifications' panel with two tabs: 'Alerts (3584)' and 'AI Assist (2)'. The 'AI Assist' tab is selected. It displays two alerts:

- IPSec tunnel GW-Oakmead-01:fast_inet::VPNC-SantaClara-01:datacenter_inet is DOWN
3 hours ago
- IPSec tunnel GW-Oakmead-01:fast_inet::VPNC-Amsterdam-01:public_inet is DOWN
3 hours ago

At the bottom of the panel is a 'View all' button.



MOD 13- 37

AI Assist uses event-driven automation to initiate automation that will automatically collect diagnostic information such as logs, packet captures, etc. that an administrator can use to troubleshoot an issue.

AI Assist

The screenshot shows the 'Alerts & Events' section of a network management interface. It features three tabs: 'ALERTS' with 3619 items, 'EVENTS' with 6264 items, and 'AI ASSIST' with 2 items. Below the tabs is a link for 'CLICK HERE FOR ADVANCED FILTERING'. A table titled 'EVENTS (2)' displays a single event: 'Tunnel State Change' occurring on 'Apr 30, 2020, 09:18'. The description is 'IPSec tunnel GW-Oakmead-01:fast_inet::VPNC-SantaClara-01:datacenter_inet is DOWN'. To the right of the table are filters for 'DYNAMIC' and 'ARUBA SUPP'. A callout box titled 'AI Assist: Alert Specific Description' points to the event description. A blue box at the top right lists features: 'Drill down to view, download, and email alert details', 'Can alert Aruba TAC: Proactively analyze, recommend', and 'Dramatically reduce time to identify and resolve issues'.

- Drill down to view, download, and email alert details
- Can alert Aruba TAC: Proactively analyze, recommend
- **Dramatically reduce time to identify and resolve issues**

AI Assist: Alert Specific Description

IPSec tunnel GW-Oakmead-01:fast_inet::VPNC-SantaClara-01:datacenter_inet is DOWN

Diag info feeds AI models for future insights, improves

self-optimization
healing
securing

MOD 13- 38

Here's how it works. When an issue, such as a port flap (port constantly goes up and down) occurs, AI Assist collects diagnostic information and alerts the administrator via the Alerts icon in the upper right. When the administrator selects this icon, they can see all the AI Assist alerts. In the figure, you can see an example of an IPsec tunnel that has an issue.

Aruba Central Licenses for AI Ops

Foundation License

AI Search

AI Insights

- Connectivity—Wi-Fi
- Wireless Quality
- Availability—Access Points
- Class and Company Baselines

AI Assist

- Dynamic logs

Advanced License

All the features in Foundation

AI Insights - Wireless Quality

- Outdoor clients impacting Wi-Fi performance
- Coverage Hole Detection
- Transmit power optimization

AI Assist

- Aruba support notification



MOD 13- 39

Aruba Central advanced licenses include all AI Ops features plus the features in the foundation licenses, and much more. Keep in mind that new features are added constantly, to check the latest license and feature support, navigate to the Aruba Central Documentation Center and search for licenses.

Alerts & Reports

This section covers the alerting and reporting features available in Aruba Central.

CONFIDENTIAL - FOR TR

Alerts & Events Page - Alerts

aruba Central View Upgrades Search or ask Aruba

Global Alerts & Events Summary List Config

Alerts & Events

ALERTS 3 **CRITICAL** 0 **MAJOR** 3 **MINOR** 0 **WARNING** 0 **EVENTS** 0

CLICK HERE FOR ADVANCED FILTERING

Show Acknowledged Alerts

Open Alerts (3) ACKNOWLEDGE ALL

Occurred On	Category	Severity	Description
Dec 10, 2021, 16:57:10	Virtual Controller Disconn...	major	Virtual controller VC1-B2 with version 8.9.0.0_81161 and IP address 10.1.145.12 disconnected, Group:IAP-TG...
Dec 9, 2021, 18:35:10	Switch Disconnected	major	Switch with serial SG04KWS165, MAC address b8:d4:e7:b4:1d:80 IP address 10.254.1.221 and Hostname 620...
Dec 9, 2021, 18:35:10	Switch Disconnected	major	Switch with serial SG03KN508L, MAC address b8:d4:e7:0e:a0:80 IP address 10.254.1.2 and Hostname 6300 d...



MOD 13- 41

A good starting point for any troubleshooting is the Alerts & Events dashboard. This dashboard displays a list of alerts and events generated for events about device provisioning, configuration, and user management. You can view the alerts and events in the List view and Summary view. You can use the Configuration view to see configuration alerts, but it is only available in the Global context. The List view components are different for the Alerts and Events tab, whereas the Summary view displays similar components.

You can view alert details and acknowledge alerts. Alerts are acknowledged automatically when the event count drops below the lowest severity threshold configured for the alert. Users with administrator 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.

Reports

24+ report types
Last day, Last week, Last month, Custom, NOW, Later
One time , Daily, Weekly, Monthly
View report, Download report, Email report

MOD 13- 42

The Reports page allows you to create various reports which you can configure to run on demand or periodically (scheduled). You must have read/write privileges or be an administrator user to create reports. The Reports page includes the browse reports option which displays the configured reports. The generated reports option displays the generated reports.

There are 24 reports in five categories from which you can choose, along with the time and date for the report. Note that new releases of Central can have new reports. You can generate the report immediately or schedule it to repeat daily, weekly or monthly. You can view the report from the Central GUI, download to a PDF file or email the report.

Create a Report

Reports Summary List Config

REPORT Which report would you like?

CONTENT What would you like in the report?

PERIOD How far back do you want to go?

RECURRENCE What is the frequency of the report?

CLIENTS

- Client Inventory
- Client Session
- Client Usage
- Guest
- Summary

INFRASTRUCTURE

- Capacity Planning
- Configuration & Audit
- Infra Inventory
- Network
- New Infra Inventory
- Resource Utilization
- RF Health
- Switch Capacity Planning
- WAN Availability
- WAN Inventory
- WAN Compliance
- WAN Transport Health
- WAN Utilization
- WAN Web Content Classification

SECURITY COMPLIANCE

- PCI Compliance
- RAPIDS
- Security Compliance

INSIGHTS

- Outdoor Clients

APPLICATIONS

- AppRF


CANCEL NEXT

MOD 13- 43

To create a report, select from one of the five categories, then click Next to navigate to the contexts. Here you can select groups, labels or a specific site. You can then select the report to be the last day, the last 7 or 30 days or customize the range.

From the Scheduled section, you can select either one time now to generate the report immediately or later to run it at a specific time. You can select how often you want to generate the report by choosing either one-time, daily, weekly or monthly intervals. Enter the report name and select if you want to email the report, then enter an email address or leave it blank. You can also pick the email format as PDF or CSV.

Report Example



The figure shows an example of a RAPIDs (rogue AP intrusion detection) report. This report displays the details of all rogue or interfering devices in Aruba Central. You can download or email the report while you’re viewing it.

UXI



This User eXperience Insights (UXI) feature set can really improve your situational awareness when managing and monitoring devices with Aruba Central.

CONFIDENTIAL - FOR TR

What is UXI

“The Wi-Fi is broken!”

“My call keeps dropping!”

“Is the issue Wi-Fi or elsewhere?”

Where to start...?



MOD 13-46

Where do IT teams start in a world of generic “the wifi is broken again” or “I’ve been struggling to get my app to work for the past week”? There could be a myriad of things that could create these kinds of problems. What makes it more difficult is that people do not complain as soon as they have a bad experience using a device, or application on the network. Instead, they tend to wait until multiple issues have happened, or they happen to bump into someone else having an issue. Or sometimes, as a result of a missed goal they will blame poor performance of the network or device they are using.

•

UXI components

Purpose-Built Sensors and Agents



- Actively test wired/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 13- 47

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 figure summarizes sensor and agent features and advantages.

The cloud-based systems provide real-time monitoring and analytic intelligence, as summarized in the figure.

Aruba User Experience Insight



Onboards like a user

- Sensors installed near users or on Zebra device
- Logs in testing applications and services like a user
- No cable run required



Onboard to ANY network

- Wireless/wired testing
- Measure DHCP, DNS, etc.
- Logs into captive portals



Test from the EDGE

- Are services & apps working?
- Zebra: Roaming, SIP call & low impact synthetic testing
- Sensor: Test throughput, and through the network



Actionable insights

- Realtime alerts; trend and baseline; Root cause analytics
- Integration to your NOC
- Stream to reporting tools

MOD 13-48

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 has sensor software which can be installed on a Zebra Android 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. In other words, Aruba can measure the end-to-end VoIP call.

These are the steps involved with UXI:

1. Sensors onboard via cellular or standard PoE. You just plug in the sensor, it looks for the network it's supposed to test. As soon as it finds that network cellular (if being used) turns off.
2. Just like a normal user the sensor associates to a network, goes through all of the same DHCP, DNS, 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. (note: the agent just onboards a wireless network not

wired).

3. Configure the sensor/agent to test services and applications which are critical to that location. First the sensor will test internal applications. First test internal applications that are inside of the firewall; for example, it might test printer or authentication server access. In the case of Zebra maybe test a basic ping access to a webserver or voice analysis of VoIP calls. (Note: the agent does not do web application tests, or throughput tests because it would impact the performance of the device itself).
4. All test results, incidents, alerts are sent to the UXI cloud dashboard real-time. The dashboard views give you the ability to trend, baseline, and troubleshoot in a very organized way. AIOps enables easy NOC notifications. People who tend to use this dashboard are commonly high-level are people who are responsible for a particular service. They just want to know is my application working in xx location. This dashboard also provides a simple way to see an external service like Facebook, for example, failing and see everything is working on the inside of my network, and only this application is having an outage. Alternately we stream all data, alerts, and/or 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, etc. Every step of the way every metric that's possible to capture is captured.

Because sensor sits outside the network onboarding just as a client would:

It can measure end-to-end app performance from client's perspective.

It is not subject to any of the same points of failure as the network.

It is vendor neutral.

The User Experience Insight (UXI) dashboard provides a single pane-of-glass: UXI and 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 UXI and AI Insights provide insights into how the network is performing and should perform. Collectively they provide holistic intelligence from the perspective of the customers and end-users and the network infrastructure. Note that AI Insight requires Aruba infrastructure, there are other requirements too (for

example, ideally over 2,000 access points).

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

At-A-Glance Status



MOD 13-49


The UXI 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 the problem. When green all is good and performing well. When red you know you have work to do to make your users happy.

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 13- 50

The figure shows a sample list of what is tested by UXI. Note that the application list on the right is only a sample; there's more and Aruba is constantly adding to the list. If there's something specific that you need to test, its possible to create your own test.

Not only can you test to a device or application, you can test through it. Here is a small list of things UXI can test:

- Log into website and applications
- Office 365, Gmail, AWS
- Pings, Webserver
- Teams, and other VoIP services
- Netflix, YouTube, Dropbox etc.
- Libraspeed, iPerf
- Voice Analysis, Roaming

Aruba Central Integration



MOD 13-51

UXI user experience status can be incorporated into the Aruba Central Network Health page.



Knowledge Check

Self-check on key learning points



Its time for a knowledge check.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Question #1

You can leverage Aruba Central's Topology pane to ease troubleshooting tasks, by quickly seeing how devices interconnect.

- True
- False



Knowledge Check

MOD 13- 53

T

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Question #2

When analyzing system status LEDs on an Aruba AP, you see a green light repeatedly blinking on for one second, then off for one second. What does this indicate?

- a. The device is fully functional.
- b. The device is operating in restricted power mode.
- c. There is a system error
- d. The device is booting.



Knowledge Check


MOD 13- 54

D

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Question #3

Match the Aruba Access Point (AP) firmware upgrade feature with its appropriate descriptions.



Knowledge Check ✓



MOD 13- 55

CONFIDENTIAL - FOR TRAINING PURPOSES

Question #4

A colleague asks you for advice about the Simple Network Management Protocol (SNMP). Which statements below should you use to help your colleague? (Select three.)

- a. An SNMP manager queries SNMP agents for information.
- b. The SNMP manager typically provides a GUI for network managers to review all gathered SNMP information
- c. The SNMP manager stores a hierarchical database – A Management Information Base (MIB) for each device.
- d. For the best security, your colleague should deploy SNMP version 3.



Knowledge Check


MOD 13- 56

A,B,D

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Question #5

Match each Aruba AIOps core component with an appropriate description.



Knowledge Check ✓



MOD 13- 57

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Question #6

Typically, you use Aruba Central's Alerts & Events page at the end of a troubleshooting session – to review your activities.

- True
- False



Knowledge Check

MOD 13- 58

F

Question #7

Which statements below are true about Aruba's User eXperience Insight (UXI) solution? (Select three.)

- a. The solution includes front-end sensors that request and generate traffic, much like an actual end user.
- b. The solution includes a cloud-based UI and intelligent analytics.
- c. One advantage of this solution is that it provides intelligent views from a network device perspective.
- d. The solution can test the entire client experience, from the Wi-Fi connection, through the network, and up to and including end-user applications.



Knowledge Check

MOD 13- 59

A,B,D

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY



Lab Activity

Consider the lab activity.

CONFIDENTIAL - FOR TR

Lab Tasks

13.01: Monitoring and Maintenance

Task 13.1

- Configure Alerts

Task 13.2

- Generating Reports in Aruba Central

Task 13.3

- Firmware Compliance

Module 14 - Troubleshooting



MOD 13- 61

The figure provides a brief review of lab tasks. Please see your lab guide for details. When complete, you will continue with Module 14 – Troubleshooting.


Rev.22.41

aruba
a Hewlett Packard
Enterprise company

Troubleshooting

Aruba Campus Access
Fundamentals

It's the final module of our course! Module 14: Troubleshooting.



Objectives

- Perform password recovery and factory reset procedures
- Troubleshoot Central connectivity
- Enable spectrum analysis
- Explore Central troubleshooting tools

When things go astray, your ability to troubleshoot is vital. After completing this module, you will be able to perform password recovery and factory reset procedures, troubleshoot Aruba Central connectivity issues, and enable spectrum analysis.

Password Recovery & Factory Reset

You begin with password recovery and factory reset procedures.

CONFIDENTIAL - FOR TRAINING

AOS-CX Password Recovery Process

1. Connect to the switch via console.
2. Power cycle the switch.
3. When the system prompts, select the **Service OS console** option by typing **0**.

Boot Profiles:

```
0. Service OS Console
1. Primary Software Image [FL.10.09.1031]
2. Secondary Software Image [FL.10.09.1031]
```

Select profile(primary): **0**

4. Login with user **admin**. No password is set for this account.
5. Enter **password** keyword and type the new password for the admin account.

```
ServiceOS login: admin
SVOS> password
Enter password: *****
Confirm password: *****
```

6. Enter **boot**.
7. Login using the admin username and the password that was set on step 5.



MOD 14- 4

In case that the credentials do not work to login into the switch, a password recovery process has to be performed. Here are the basic steps to perform a password recovery:

1. Connect to the switch using the console port.
2. Power cycle the switch.
3. When the system prompts, select the Service OS console option by typing 0.
4. Login with the user admin. No password is set for this account.
5. Enter password keyword and type the new password for the admin account.
6. Enter boot command.
7. Once the switch boots completely, login using the admin username and the password that was set on step 5.

AOS-CX Factory Default

Zeroize

- Deletes all checkpoints in the system
- Returns the system to factory default

Erase-startup

- Deletes only the startup checkpoint in the system
- No factory default

```
switch(config)# erase all zeroize
```

This will securely erase all customer data and reset the switch to factory defaults. This will initiate a reboot and render the switch unavailable until the zeroization is complete.

This should take several minutes to one hour to complete. Continue (y/n)? **y**

The system is going down for zeroization.



MOD 14- 5

The command `erase all zeroize` restores the switch to its factory default configuration. You will be prompted before the procedure starts. Once complete, the switch will restart from the primary image with factory default settings.

The command `erase-startup` erases the startup checkpoint only: all other checkpoints will still remain on the switch. Using the `erase all zeroize` command, however, will erase all checkpoints.

AP CLI Password Reset Via Central

The screenshot shows the Aruba Central web interface. The left sidebar contains a navigation menu with sections: Manage (Overview, Devices, Clients, Guests, Applications, Security) and Analyze (Alerts & Events). The 'Devices' section is expanded, showing a list of devices. The main content area is titled 'Access Points' and shows a breadcrumb trail: WLANs > Access Points > Radios > Interfaces > Security > Services > **System** > IoT > Configuration Audit. Under the 'System' tab, there are sections for 'General', 'Administrator', and 'Mesh'. The 'Administrator' section contains a table with columns: Username, Password, and Type. The table has one row with the username 'admin', a masked password '*****', and the type 'Admin'. A blue dashed box highlights the edit and delete icons in the rightmost column of the 'admin' row.

Username	Password	Type
admin	*****	Admin

MOD 14- 6

Now that you know how to perform the password recovery procedure and factory reset the AOS-CX switches, let's move on to the Aruba APs. The next two sections show you how to perform the password recovery process as well as how to set the APs to their factory default settings.

To reset AP password, navigate to AP Group > Devices > System. Scroll down to Administrator and change the admin password.

AP Factory Reset

1. Connect to AP via console.
2. Power cycle the AP.
3. Interrupt boot process by hit **Enter**

```
Net:    eth0
Radio:  bcm47622#0,bcm47622#1
Reset:  warm
FIPS:   passed
Hit <Enter> to stop autoboot:  0
```

4. Type **factory** and hit **enter**

```
apboot> factory
Clearing state...
Checking OS image and flags
...
done
```

5. Type **reset**



MOD 14- 7

The figure shows the AP factory reset procedure.

Troubleshooting Central Connectivity



And now, on to troubleshooting Aruba Central connectivity.

CONFIDENTIAL - FOR TRAINING

Central Connection Troubleshooting - AP

```
AP-50-a8# show activate
```

```
IAP MAC Address      :20:4c:03:32:aa:a8
IAP Serial Number    :CNG9K2R64T
Cloud Activation Key  :GHGMVKXI
Activate Server       :device.arubanetworks.com
Activate Status       :fail-prov-no-shipped
Provision interval    :5 minutes20:4c:03:32:aa:a8#
```

No entry for AP on Aruba Activate or no internet access

```
AP-50-a8# show ap debug cloud-server
```

```
AP mgmt mode         :local-mgmt
cloud config recved   :FALSE
autojoin mode         :disable
state diff            :disable
Device Cert status    :SUCCESS
```

Check Aruba Central AP License Assignment



MOD 14- 9

If an Aruba AP is having issues contacting Central, you can use the AP's console and CLI to perform some basic troubleshooting.

On the AP console, use `show activate` to see Aruba Activate status. In the figure, there is a problem. The status is `fail-prov-no-shipped`. This is an indication that Activate does not know about the AP or the AP does not have internet access.

Another useful command is the `show ap debug cloud-server`. In this example, `cloud config recved` is `FALSE`. If the AP has internet access it's time to verify in Central the AP was assigned a license.

Central Connection Troubleshooting - Switch

```
6300(config)# show aruba-central
```

```
Central admin state : enabled
```

```
Central location : N/A
```

```
VRF for connection : N/A
```

```
Central connection status : N/A
```

```
Central source : activate
```

```
Central source connection status : connection failure
```

```
Central source last connected on : N/A
```

```
System time synchronized from Activate : True
```

```
Activate Server URL : devices-v2.arubanetworks.com
```

```
CLI location : N/A
```

```
CLI VRF : N/A
```

```
Source IP : N/A
```

```
Source IP Overridden : False
```

```
Central support mode : disabled
```

No entry for the switch
on Aruba Activate
or no internet access



MOD 14- 10

The figure shows an example of how to troubleshoot an AOS-CX switch and its connection to Central. The show aruba-central command output shows that you don't have a Central location (URL) and the switch is not connected. Another problem could be that Activate has no provisioning rule for this switch. Alternatively, the problem could be that the Central account has not given the switch a subscription key.

AOS-CX Switch Logs

```
6200# debug ztp all
6200# show logging -r
-----
Event logs from current boot
-----
2021-10-01T13:18:38.386195+00:00 6200 hpe-restd[843]:
Event|4650|LOG_INFO|AMM|-|Unable to fetch Aruba Central location from
activate via VRF .
2021-10-01T13:18:38.123971+00:00 6200 hpe-restd[843]:
Event|7708|LOG_INFO|UKWN|1|Certificate devices-v2.arubanetworks.com verified
and accepted
2021-10-01T13:18:34.011232+00:00 6200 hpe-restd[843]:
Event|4648|LOG_INFO|AMM|-|Switch time is synced with Aruba Activate Server
https://devices-v2.arubanetworks.com.
```



MOD 14- 11

Debugging ZTP offers more information on why a switch does not connect to Aruba Central. Enable debug aruba-central and debug ztp to dig deeper into the connection issue.

Here the logs confirm that the switch has access to Activate but Activate has no rules for this switch. This could be an issue with Activate or Central not having a subscription key to this device. Once you correct the problem, issue the show aruba-central command to verify if the problem has been fixed. This should show you that the switch is connected to the Central account.

Note: the command to debug ZTP events on AOS-CX switches is debug ztp. Also, the “-r” parameter in the show logging command displays the logging events from newest to oldest (reverse chronological order.)

AOS-CX Aruba Central Disable

Local config not possible when managed by Central

```
sw-edge1(config)#  
  aruba-central  Configure Aruba-Central  
  debug          Configure debug logging  
  end            End current mode and change to enable mode  
  exit           Exit current mode and change to previous mode  
  list           Print command list  
  no             Negate a command or set its defaults
```

Disabling Aruba Central

```
sw-edge1(config)# aruba-central  
sw-edge1(config-aruba-central)# disable
```



MOD 14-12

By default, the AOS-CX switch cannot be locally configured when it is managed by Central. The command `aruba-central disable` disables the connection to the Aruba Central server. When the connection is disabled, the switch does not attempt to connect to the Aruba Central server or fetch central location from any of the three sources (CLI/Aruba Activate/DHCP). It also disconnects any active connection to the Aruba Central server. Disabling and enabling Aruba Central will trigger configuration refresh, where the current configuration will be pushed from Central. Using this command is useful to kickstart the switch to reconnect to Central when troubleshooting connection issues. The `aruba-central enable` command re-enables the Central access.

Central Configuration Lockout Mode/Support Mode

Set device to only be writable from Aruba Central

```
sw-edge1(config)# configuration-lockout central managed
sw-edge1(config)# show configuration-lockout
Configuration Lockout
-----
central: managed
```

Enable Central Support Mode

```
sw-edge1(config)# aruba-central support-mode
sw-edge1(config)# show aruba-central
<omitted>
Central support mode : enabled
```



MOD 14- 13

The AOS-CX switch connects to Aruba Central in either monitor or managed mode. When the device connects in monitor mode, Aruba Central monitors the switch configurations. When it connects in managed mode, the configuration-lockout central managed command does not allow configuration changes from other switch configuration methods, like the switch as CLI or the switch's web UI.

The aruba-central support-mode allows the device to be writable for all operations in Aruba Central lockout mode for troubleshooting. The no form of this command disables this activity.




Spectrum Analysis

Did you ever wish you could see outside of the human range, like x-ray vision or perhaps into the microwave range? You should check out spectrum analysis.

CONFIDENTIAL - INTERNAL

Spectrum Analysis Overview



MOD 14-15

Wireless networks operate in environments with electrical and RF devices that can interfere with network communications. Microwave ovens, cordless phones, and even adjacent Wi-Fi networks are all potential sources of continuous or intermittent interference. The spectrum monitor software modules on IAPs and AOS 10 APs can examine the RF environment in which the Wi-Fi network is operating, identify interference, and classify its sources. An analysis of the results can then be used to quickly isolate issues associated with packet transmission, channel quality, and traffic congestion caused by contention with other devices operating in the same band or channel. Spectrum monitors are Instant AP radios that gather spectrum data but do not service clients. Each SM scans and analyzes the spectrum band used by the AP's radio (2.4 GHz, 5 GHz, or 6 GHz).


By default, an AP's radio accepts device associations. When you switch a radio to spectrum analysis mode, you can use it for troubleshooting RF issues. However, when it is in this mode, it will disconnect any associated users and deny user's access to the radio. When done troubleshooting, remember to disable spectrum analysis for the radio so that it can, again, allow users to associate to it. Alternatively, you can dedicate an AP(s) in an area that perform spectrum analysis full time. When doing this full time, these APs are commonly called spectrum monitors or SMs.

Central Troubleshooting Tools

This section introduces you to some tools that Aruba Central can use to troubleshoot connection and device issues.

CONFIDENTIAL - FOR TRAINING

Network Check



Access-Point

- Ping Test
- Traceroute
- Speed Test (iPerf)
- HTTP Test
- HTTPS Test
- TCP Test

Gateway

- Ping Test
- Traceroute

Switch

- Ping Test
- Traceroute
- Ping Sweep Test
- Speed Test (iPerf)
- HTTP Test
- HTTPS Test



MOD 14- 17

The Network Check page provides information on network issues faced with devices Aruba Central manages. You can diagnose network performance from the Tools tab using various test types such as, ping test, speed test, TCP test, and Traceroute.

Network Check – Ping Test

The screenshot displays the 'Network Check' interface with the following configuration:

- Device Type:** Access Point
- Sources:** AP-1
- Test:** Ping Test
- Destination Type:** Hostname/IP Address
- Hostname/IP Address:** google.com

The 'RUN' button is visible at the bottom of the configuration section.

The 'DEVICE OUTPUT' section shows the results for device AP-1:

```
Output for the device: AP-1
Test Type: PING
Source: [Access Point] AP-1
Target: [EXTERNAL] google.com

=====
Output Time: 2021-12-09 18:35:18 UTC

COMMAND=ping google.com
PING 172.217.173.78 (172.217.173.78): 56 data bytes
64 bytes from 172.217.173.78: icmp_seq=0 ttl=115 time=7.3 ms
64 bytes from 172.217.173.78: icmp_seq=1 ttl=115 time=3.9 ms
64 bytes from 172.217.173.78: icmp_seq=2 ttl=115 time=7.9 ms
64 bytes from 172.217.173.78: icmp_seq=3 ttl=115 time=4.7 ms
64 bytes from 172.217.173.78: icmp_seq=4 ttl=115 time=8.3 ms

--- 172.217.173.78 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max = 3.9/6.4/8.3 ms

=== Troubleshooting session completed ===
```

MOD 14- 18

The example in the figure shows a ping test executed from one managed AP to an external host/address. To run a ping test, select the device type, the source devices, the test type (ping in this example), the destination type, and the destination address, then click Run. A device output section will populate with the device's CLI output.

Device Check

Select Device, Select a Type of Test and Run the test

Device Type ▾

Test ▾

Devices which are already running commands shall not execute newly added commands
Output history of device with buffer space issues shall be automatically cleared

RUN

Gateway

- Interface Bounce
- PoE Bounce

Switch

- Interface Bounce
- PoE Bounce
- Chassis Locate



MOD 14- 19

The Device Check page allows you to run diagnostic checks and troubleshoot switches and gateways. The following tests are available:

Gateways:

- Interface Bounce restarts the port interface and forces a client to re-initiate a DHCP request.
- PoE Bounce restarts the PoE port and the device that is either connected to the PoE port or powered by it.

Switches:

- Interface Bounce restarts the port interface and forces a client to re-initiate a DHCP request.
- PoE Bounce restarts the PoE port and the device that is either connected to the PoE port or powered by it.
- Chassis Locate activates the switch locator LED. The locator LED indicates the physical location where an Aruba switch is currently installed.

Commands

Network Check | Device Check | **Commands** | Console | Proximity Tracing

Select a Device Type, Select Devices, Add the Commands and Run them

Device Type ▼

Select commands from one or more categories

Categories | **Commands** | Selected Commands

No data to display | No data to display | No data to display

Add > | < Remove | < Remove All

* - Supports additional mandatory filters
+ - Supports additional optional filters
^ - Cannot run with other commands

This list of commands will run in the order of its sequence (Maximum is 20)

Devices which are already running commands shall not execute newly added commands
Output history of device with buffer space issues shall be automatically cleared
Few commands require the log level to be set as debug to see the output

RUN | RESET

MOD 14- 20

The Commands tab allows you to perform a network health check on devices at an advanced level using command categories. Read-only users can also perform advance checks. To run a command, select the device type you want to run the command then select the device(s). The category box will show the command category applicable to the selected device(s). Once you chose a category, the available commands will populate. Choose the command you want to execute and click Add. When all the commands you want to run are in the selected commands box, click Run. A device output section will display the CLI output from each selected device.

Commands - Example

The screenshot displays the 'Commands' tab in a network management interface. At the top, there are navigation tabs: 'Network Check', 'Device Check', 'Commands' (active), 'Console', and 'Proximity Tracing'. Below these, a 'Device Type' dropdown is set to 'Access Point' and an 'Available Devices' dropdown is set to 'AP-2'. A dashed orange box highlights the 'Categories' and 'Commands' sections. The 'Categories' list includes 'All Category', 'Wireless' (selected), 'Security', 'Network', 'Airgroup', and 'Custom'. The 'Commands' list includes 'AP Client Match History', 'AP Client Match Live', 'AP Client View', 'AP Driver Configuration', 'AP ESSID Table', and 'AP Management Frames*'. A 'Repeat' checkbox is present. To the right, the 'Selected Commands' section shows 'AP Client Table' as the selected command. Below this, a 'RUN' button is highlighted with a dashed orange box, and a 'RESET' button is also visible. A large diagonal watermark 'CONFIDENTIAL - FOR TRAINING PURPOSES ONLY' is overlaid on the image.

Network Check | Device Check | **Commands** | Console | Proximity Tracing

Device Type: Access Point | Available Devices: AP-2

Select commands from one or more categories

Categories

- All Category
- Wireless
- Security
- Network
- Airgroup
- Custom

Commands

- AP Client Match History
- AP Client Match Live
- AP Client View
- AP Driver Configuration
- AP ESSID Table
- AP Management Frames*

47 Commands | 0 selected

* - Supports additional mandatory filters
+ - Supports additional optional filters
^ - Cannot run with other commands

☐ Repeat

Selected Commands

AP Client Table

1 Commands | 0 selected
This list of commands will run in the order of its sequence (Maximum is 20)

ADD > | < REMOVE | < REMOVE ALL

Devices which are already running commands shall not execute newly added commands
Output history of device with buffer space issues shall be automatically cleared
Few commands require the log level to be set as debug to see the output

RUN | RESET

MOD 14- 21

The figure shows how to run a command about the wireless clients on the selected AP: AP-2.

Remote Console

Remote Console Session

New Session

Saved Sessions

Device Type

Access Point

Access Point

Access Point

Access Point

AP-2

AP-2

AP-2

Username

admin

admin

admin

Password

Create New Session

REMOTE CONSOLE

Console session for the device: AP-2

Current Session

admin@AP-2 [03:42:26 PM] x +

dbclient:
Host 'localhost' key accepted unconditionally.
(ssh-rsa fingerprint md5 05:d7:9b:14:9b:78:fe:9d:15:8d:e5:b1:42:23:21:16)

show tech-support and show tech-support supplemental are the two most useful outputs to collect for any kind of troubleshooting session.

AP-2#
AP-2#

[2021 Dec 09 03:42:27 PM] Ssh session started

MOD 14- 22

The Remote Console enables you to connect to the device directly through SSH sessions to troubleshoot any device issues in the network. You can either view saved sessions or create new sessions.



Knowledge Check

Self-check on key learning points



Its time for a knowledge check.

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Question #1

It is time to play **True** or **False**!

- To perform the password recovery process on an AOS-CX device, you must connect via the console port and power cycle the switch. **TRUE**
- Both the **erase all zeroize** and **erase startup** commands delete the startup checkpoint. **TRUE**
- You can reset an AP's CLI password from Aruba Central. **TRUE**
- When using the command **show logging -r**, the “-r” means to list log entries in reverse chronological order. **TRUE**

a

Knowledge Check

MOD 14- 24

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY

Question #2

See the figure. What could the highlighted output indicate? (Select two.)

- a. There is no entry for this AP on Aruba Activate.
- b. The AP has no internet access.
- c. Aruba's Activate service is offline.
- d. The AP is currently being provisioned by Aruba Central.

AP-50-a8# **show activate**

IAP MAC Address	:20:4c:03:32:aa:a8
IAP Serial Number	:CNG9K2R64T
Cloud Activation Key	:GHGMVKXI
Activate Server	:device.arubanetworks.com
Activate Status	:fail-prov-no-shipped
Provision interval	:5 minutes20:4c:03:32:aa:a8#

MOD 14-25

A,B

Question #3

Aruba AP radios can act as spectrum monitors, to gather spectrum data while servicing clients.

- True
- False

a

Knowledge Check

MOD 14- 26

F

Question #4

You are using Aruba Central's troubleshooting tools to analyze some network issues. Which actions can you perform using these tools? (Select all that apply.)

- a. Ping test
- b. Bounce an interface
- c. Blink the switch locator LED
- d. Start a remote console session to a device

a

Knowledge Check

MOD 14- 27

A,B,C,D

CONFIDENTIAL - FOR TRAINING PURPOSES ONLY



Lab Activity

Consider the lab activity.

CONFIDENTIAL - FOR TRAINING

Lab Tasks

14.01: Troubleshooting Overview

Task 14.1

- Collection Log Information for Senior Support or TAC

Task 14.2

- Troubleshoot Switch to Aruba Central connection

Task 14.3

- Site Level Troubleshooting

Congratulations!
You have completed this course!



MOD 14- 29

The figure provides a brief review of lab tasks. Please see your lab guide for details. When complete, congratulations! You have completed this course.

Next Steps



Certification

How to Study for Aruba Certification Exams – Free!
Sign up, take exam

Advanced knowledge

The next course in this series is:
Implementing Aruba Campus Networks, rev 23.11



MOD 14- 30

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 the next course in this series: *Implementing Aruba Campus Networks, rev 23.11*

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-ACAF-ILT-v22.41