aruba®
NETWORKS

Aruba Campus Access Fundamentals

LAB GUIDE

EDU-ACAF-RLABS-v22.41

Aruba Campus Access Fundamentals

22.41

Lab Guide

December 2022

Aruba Campus Access Fundamentals

SKU: EDU-ACAF-RLABS-rev.22.41

DECEMBER 2022

Aruba Campus Access Fundamentals

LAB GUIDE TABLE OF CONTENTS

Lab 01.01 Testing Remote Lab Connectivity

Overview

The Aruba Training Lab provides the equipment you need to complete several lab activities. You should know the purpose and access procedures for this equipment.

- **MGMT PC**: This client is used for remote lab management access, traffic analysis, and to access the switch's CLI via SSH.
- **PC-1:** This client is used for connectivity testing. It has a wired and wireless NIC.
- **PC-4**: This client is used for connectivity testing. It has a wired and wireless NIC.
- **Edge-1 switch**: This is one of your Access-Switches; it will be named sw-edge1.
- **Edge-2 switch**: This is one of your Access-Switches, named sw-edge2.
- **Agg-1 switch**: This is the primary aggregation switch, named sw-agg1.
- **Agg-2 switch**: This is the secondary aggregation switch, named sw-agg2.
- **AP1**: This is the first AP, connected to your sw-edge1 on port 2.
- **AP2**: This is the second AP, connected to your sw-edge2 on port 2.
- **OOBM switch**: You have NO access to this switch.
- **Core router**: Your aggregation switches connect with a routed connection to the Core router.
- **Windows Server**: This system will provide DHCP, DNS and NTP services for the lab devices. You have NO access to this server.
- **ClearPass server**: It is used as an authentication, authorization, and accounting (AAA) RADIUS server for your network environment.

Objectives

After completing this lab, you will have all needed information to support the hands-on labs in this course.

Task 1: Aruba Training Remote Lab Access

Objectives

- Validate remote lab connectivity and ability to log in.
- Ensure that you have remote lab access during this training.

Steps

1. On your local computer, launch a web browser, and access the Aruba Training Lab web portal at the URL:https://arubatraininglab.computerdata.com.
2. Enter your **username** and **password** (if you do not have one, ask your instructor for the credentials) and click the **Sign in** button.

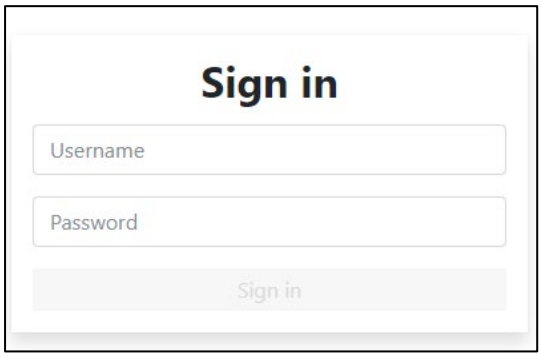

Task 2: Testing Connectivity

Objectives

- Test connectivity and authentication credentials for each of the devices.
- Working from the Aruba Training Lab diagram, connect and log into the lab devices and your client PCs.

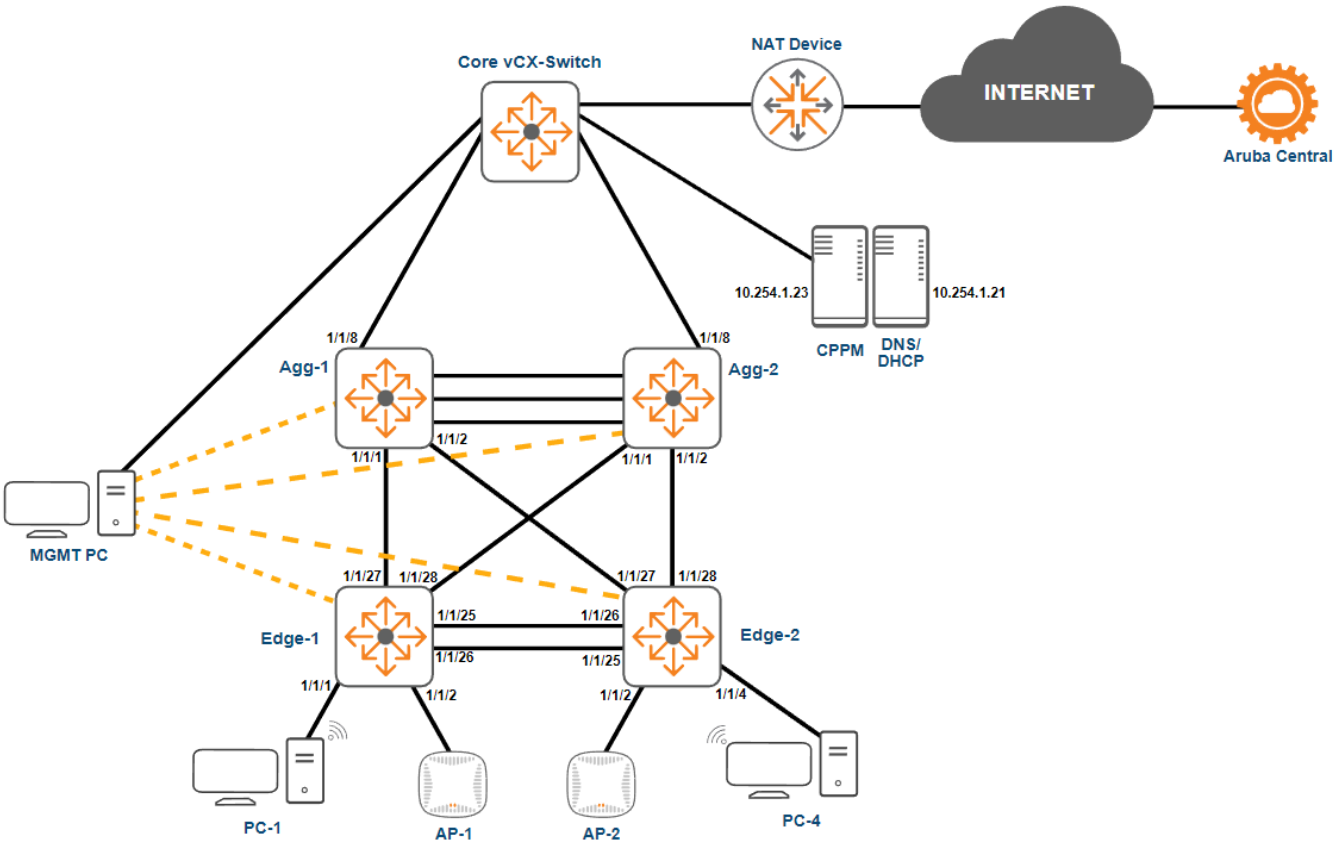

sw-edge1 (Edge-1)

This device will be the first edge switch in your lab environment, to be named **sw-edge1**.

1. To connect to the console of the Edge-1 switch, right-click on the icon in the lab diagram and select **Open Console.** A new browser tab should open with a blank, black screen.
2. Press **[enter]** a couple times, and you will see a user prompt. Login using admin and no password (leave it blank).
3. The switch will ask you to define a new password; hit **[Enter]** twice.

NOTE: This switch is factory default at the start of the labs. A factory default switch prompts the administrator to change the password after the first login.

```
6300 login: admin
Password:


Please configure the 'admin' user account password.
Enter new password:
Confirm new password:
6300#
```

4. Take note of the switch **sw-edge1** *serial number* and *MAC address*. You may also save this in a text file on your local system. When using Aruba Central, this information will be used to identify the devices.

```
show system
```

Sw-edge1	
Serial Number	
MAC Address	

```
6300# show system
Hostname           : 6300
System Description : FL.10.09.1040
System Contact     :
System Location    :

Vendor             : Aruba
Product Name       : JL666A 6300F 24G CL4 PoE 4SFP56 Sw
Chassis Serial Nbr : SG00KN500Z
Base MAC Address   : 64e881-3f6540
ArubaOS-CX Version : FL.10.09.1040

Time Zone          : UTC

Up Time            : 4 days, 20 hours, 29 minutes
CPU Util (%)       : 10
Memory Usage (%)   : 23
```

sw-edge2 (Edge-2)

This device will be the second edge switch in your lab environment, to be named **sw-edge2**.

5. To connect to the console of the Edge-2 switch, right-click on the icon in the lab diagram and select "**Open Console.**" A new browser tab should open with a blank, black screen.
6. Press **[enter]** a couple times, and you will see a user prompt.

> **NOTE:** This switch has been pre-configured in the remote lab environment; you must connect using the correct credentials – see next step.

7. Login using **admin** and password **Aruba123!**

> **IMPORTANT:** For ease of use we use a simple password in the lab (Aruba123!), please never use a simple password in real life!

```
login: admin
Password:
```

8. Take note of the switch sw-edge2 *serial number* and *MAC address*. You may also save this in a text file on your local system.

```
show system
```

Sw-edge2	
Serial Number	
MAC Address	

Aggregation Switches

These devices will be the primary (A) and secondary (B) aggregation switches in your campus lab environment. They will be named **sw-agg1** and **sw-agg2**.

9. To connect to the console of the **Agg-1,** right-click on the icon in the lab diagram and select "**Open Console.**" Press **[enter]** a couple times, login with username **admin**, password **Aruba123!**
10. This device has been pre-configured. Verify the prompt shows **sw-agg1**.

```
sw-agg1 login: admin
Password:

sw-agg1#
```

11. Repeat the previous 2 steps for the **Agg-2,** the hostname should be **sw-agg2**.

```
sw-agg2 login: admin
Password:

sw-agg2#
```

NOTE: You don't need to record the serial and MAC address of the 8325 switches: they will be automatically provisioned with the correct configuration.

AP1 and AP2

There are two APs in your lab setup. These APs are factory default at the start of the training labs. In the next steps you will make an inventory with the MAC and serial number of both APs. This will make it easier in later labs to identify each AP either on the switch or in Aruba Central.

You will take note of the AP MAC address using the AP console connection. Right after an AP starts to boot, you can press [**Enter**] to access the *apboot* environment. In this apboot mode you can execute the **mfginfo** command (manufacturing information) to see the AP MAC address and serial number.

In the next procedure you will reboot the AP. Make sure to switch quickly to the AP console upon rebooting: you will need to press [**Enter**] to access the apboot mode within a few seconds after the AP starts to boot.

12. To connect to the console of the **AP1,** right-click on the icon in the lab diagram and select "**Open Console.**" Press [**Enter**] a couple times, you should see a login prompt.
13. Return to the lab dashboard, right-click again on AP1 again and select **reboot**.
14. Quickly switch to the web page with the AP1 console access; press [**Enter**] when you see the option to access the apboot mode.

```
APBoot 2.4.0.8 (build 64221)
Built: 2018-03-28 at 20:30:14

Model: AP-303H
DRAM:  512 MiB
Flash: Detected W25Q32FV_SPI: total 4 MiB
NAND:  Detected MX35LFxGE4AB: total 128 MiB
Power: 802.3af POE
Net:   eth0
Radio: ipq4029#0, ipq4029#1
Reset: warm
FIPS:  passed

Hit <Enter> to stop autoboot:  0
apboot>
```

15. In the apboot context, execute the **mfginfo** command and press [**Enter**].

```
mfginfo
```

Example output, your MAC and serial will likely be different from this output.

```
apboot> mfginfo
Inventory:
Card 0: System
        Wired MAC           : 20:4c:03:c5:fc:34
        Wired MAC Count     : 4
        Date Code           : 052520
        Serial              : CNKCK2R7R0
        Wireless MAC        : 24:62:ce:c5:b4:70
        Wireless MAC Count  : 2
        Country             : CCODE-US-b69c719895e67525a096729da53abcb37ee4b837
Card 1: CPU
        Assembly            : 2010258C
        Serial              : Y105810DA
        Date Code           : 051320
        Major Rev           : 02
        Minor Rev/Variant   : 00
Card 2: Power
        Assembly            : 2010259C
        Serial              : Y105803B8
        Date Code           : 051320
        Major Rev           : 02
        Minor Rev/Variant   : 00
apboot>
```

16. Take note of the Serial number and MAC address:

AP1	
Serial Number	
MAC Address	

17. After you have noted the serial and MAC address, you can enter the **reset** command to reboot the AP.

```
reset
```

```
apboot> reset
resetting ...
```

NOTE: Even if you don't run any command, the AP will reboot automatically after a few minutes.

18. Repeat the previous procedure for AP2 and take note of the serial number and MAC address:

AP2	
Serial Number	
MAC Address	

MGMT PC, PC-1 and PC-4

The MGMT PC is used for device management access. The PC1 and PC4 will be used as lab test-hosts for wired and wireless access.

19. To access the desktop MGMT PC, right-click on the icon in the lab diagram and select "**Open Desktop**." A new browser tab will open with the remote desktop.
20. Repeat the previous step on **PC-1** and **PC-4**.

Core-router (via MGMT PC)

Unlike the access and aggregation switches, the core router in the remote lab is not a physical switch, it is running the AOS-CX simulator software developed by Aruba.

21. Switch back to the **MGMT PC** desktop.
22. On the desktop, you can open Putty or MTPutty (multi-tab Putty). Either application is fine, use what works best for you (separate Putty windows or multiple tabs in MTPutty).

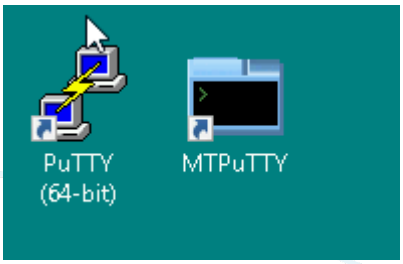

23. You will find saved sessions to **rtr-core1**

Putty example:

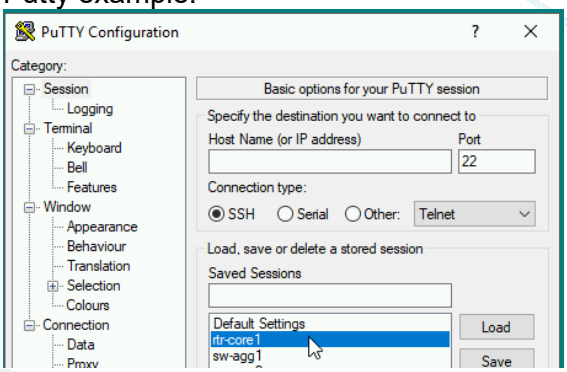

MTPutty example: expand the Putty Sessions.

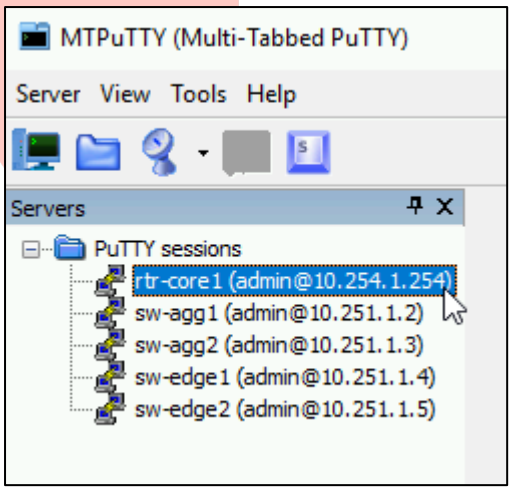

24. Double click **rtr-core1** saved session to open the connection.
25. Login using username **admin** and password Aruba123!

ClearPass (via MGMT PC)

ClearPass is the AAA policy management solution used in the Aruba lab environment.
The ClearPass admin web user interface (UI) can be accessed using the MGMT PC.

26. Switch back to the **MGMT PC** desktop. On MGMT PC, open a web browser, Google Chrome for example, and navigate to: **https://10.254.1.23**
27. You will be presented a security certificate warning. Accept the warning. You will see the login page.
28. Click on the **Policy Manager** icon.
29. Login using username **admin** and password **Aruba123!**

Aruba Central Access

Aruba Central is a cloud-based solution. It can be accessed with any system that has internet access.
You can access Aruba Central using your local PC internet connection.

30. On your local PC, open a web browser and navigate to

https://common.cloud.hpe.com

31. Click on **Sign in with SSO.**

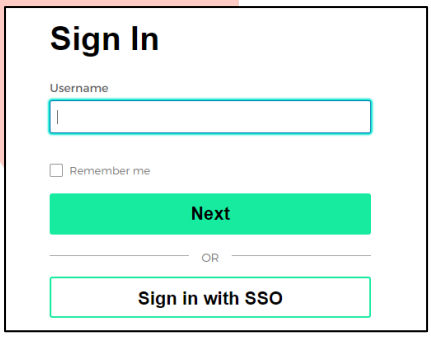

32. In the Sign In With Single Sign-On box, enter the **username** for HPGLCP (with the @arubatraininglabs.net suffix) and click **Next**. The username is given to you with your lab credentials when you purchase the lab access separately.

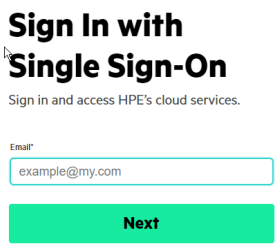

You will now be redirected to the Remote lab login page based on the arubatraininglabs.net domain name.

33. Enter the username and the password for HPE GLCP (@arubatraininglabs.net) as provided by your instructor and click **Log In**.

34. After the login has completed successfully, the GLCP will present the Aruba Central application in the Featured Applications list. Click **Launch** for the Aruba Central tile.

35. This will take you to the Aruba Central **Global** > **Overview** page.

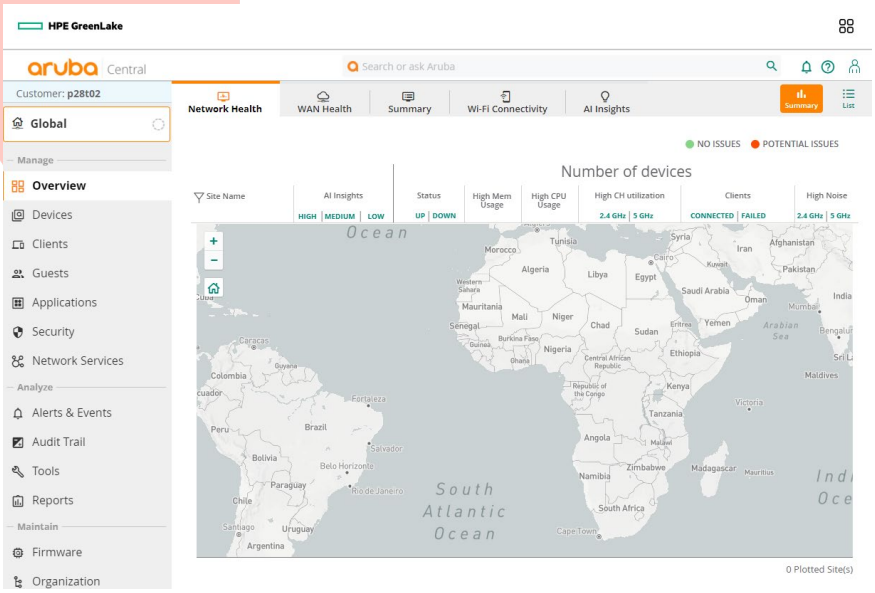

NOTE: In these steps you have seen how to access Aruba Central. In the remainder of the lab guide these steps will not be repeated. You will simply be instructed to access Aruba Central. Please refer to this task if you would need assistance with the login steps.

You have completed this lab!

Lab 01.02 Networking Fundamentals

Overview

In these lab activities you will have the role of an associate network engineer.
You will assist your senior colleague during the setup of a new customer network, and you will explore several technologies and features during this setup.
The new customer network is a campus that requires wired and wireless access.

You must learn how to configure these features and respond to customer questions about your setup.

First, you prepare and build the wired network. Once this infrastructure is in place, you add the wireless access layer to the campus.

In the current lab activity, your senior colleague wants to prepare your basic network knowledge before you go to the customer. You will explore the structure of network packets and the role of layer 2, layer 3, and layer 4 TCP and UDP in these packets.

Objectives

- Understand headers and encapsulation
- Understand a TCP header
- Understand a UDP header

Task 1: Discover Headers and Encapsulation

A key step to more fully understand data forwarding and networking protocols is the ability to analyze packets and identify their OSI model headers and header content.

In this task you will explore Ethernet, IP, UDP and TCP headers.

Objectives

- Understand the principle of encapsulation in an Ethernet frame
- Understand the Ethernet packet structure
- Understand the IP header
- Understand a TCP header and session
- Identify the Application layer

Steps

NOTE: In the next steps you will make a network trace yourself using Wireshark. Example traces can also be found in the ACAF Student Files folder on the desktop of your MGMT PC.

Example files:
lab01.02-ip-tcp-http-10.251.1.2.pcapng
lab01.02-ip-udp-dhcp.pcapng

Start the Wireshark packet analyzer

Wireshark is a well-known, open-source packet analyzer tool. It can capture traffic in different media types - Ethernet, 802.11, Bluetooth, USB, and more. It is supported on main desktop operating systems like Microsoft Windows, MacOS, and many Linux distributions. For more information, go to:
www.wireshark.org
https://wikipedia.org/wiki/Wireshark

In the next steps you will start a network trace and capture the traffic generated when accessing a Web Page.
You will use MGMT PC to open an HTTP connection to sw-agg1, on the configured Out of Band IP address 10.251.1.2.

1. Use the lab dashboard to open a session to **MGMT PC**
2. On the desktop, double-click the **Wireshark** icon.

By default, Wireshark will also show the raw, hexadecimal packet contents. While this may be useful in some troubleshooting, you may disable that view.

3. Expand the "**View**" menu and **uncheck** the "**Packet Bytes**" option.

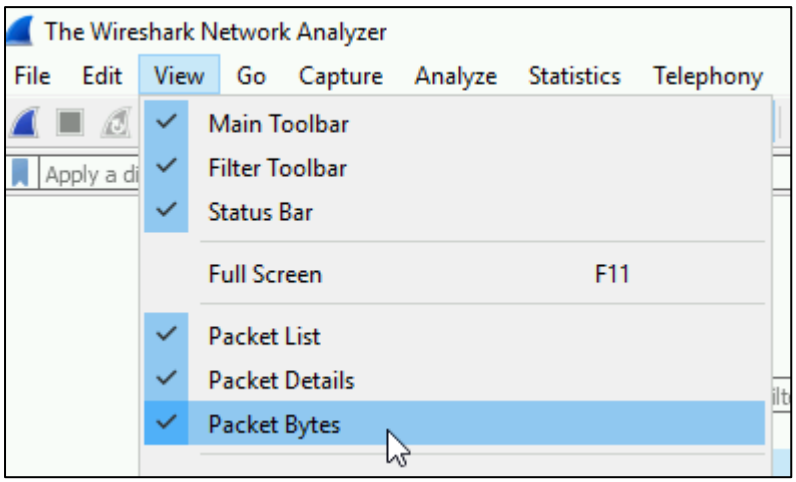

4. Double click the **OOBM** entry. This starts a packet capture session on that interface.

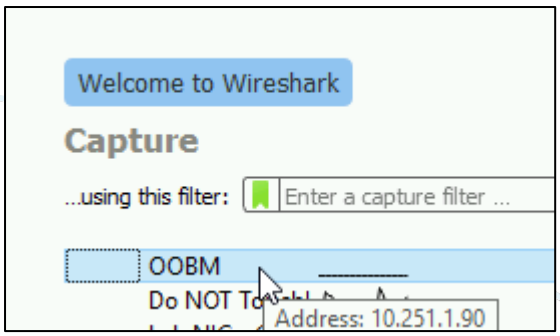

5. Identify the following components in the Wireshark UI.

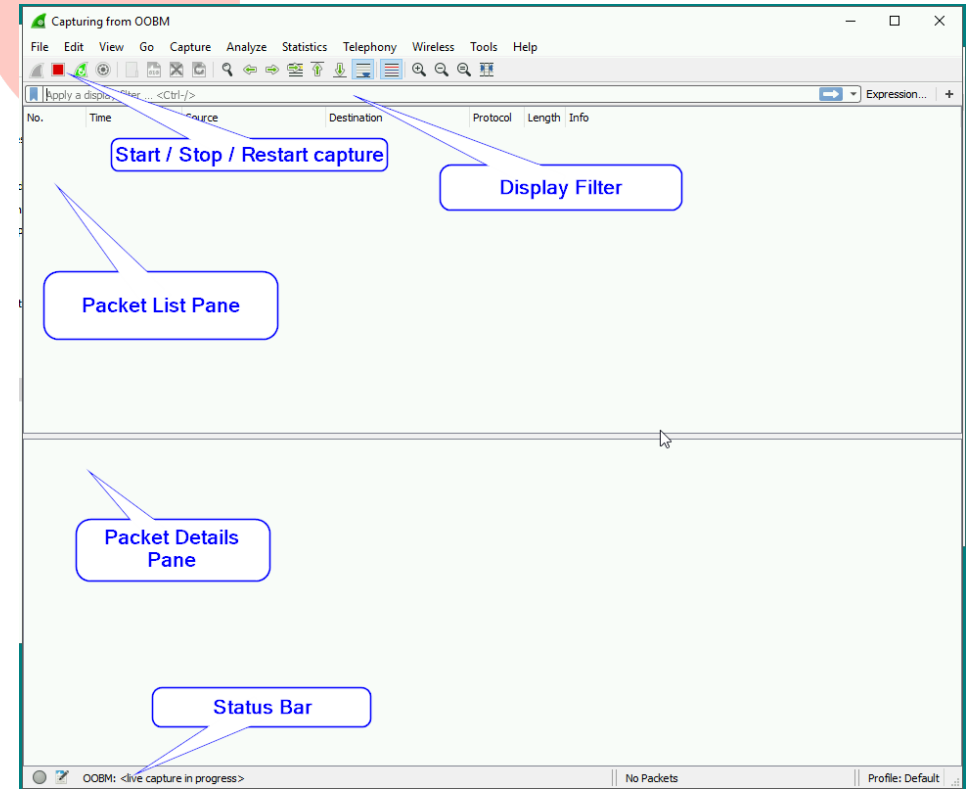

6. In the **Display Filter** field, type "**ip.addr == 10.251.1.2**" with no quotes and hit **[Enter]**. This tells Wireshark to only display packets to and from that IP address.

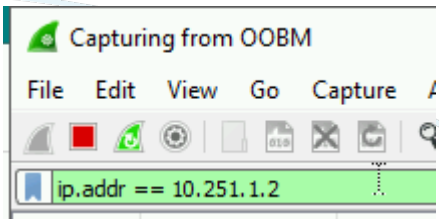

> **TIP:** In the example above, the background color of the display filter is green, indicating that syntax is correct. A red or yellow background indicates that Wireshark has a problem with how you entered the filter text.

7. Open a browser and type "**10.251.1.2"** IP address in the URL field and hit **[Enter]**. A page will pop up. There is no need to login or accept the security warning.
8. Move back to Wireshark. You should see a long list of entries that represent every single Data Unit exchanged with the server while downloading the page.
9. Stop the network capture by clicking on the red stop button in the Toolbar.
10. Scroll all the way up.

11. You will first see three packets listed as “SYN”, “SYN, ACK” and “ACK” under the Info column.

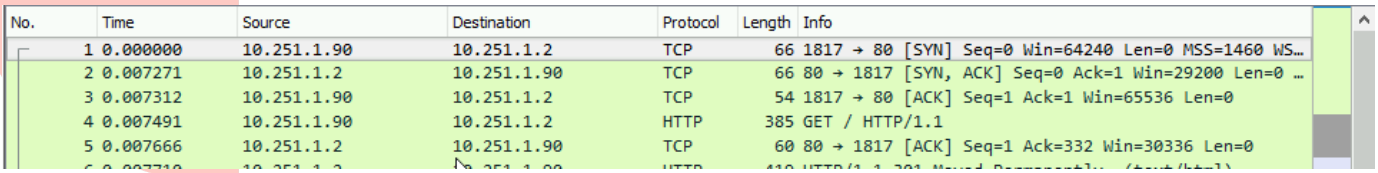

No.	Time	Source	Destination	Protocol	Length	Info
1	0.000000	10.251.1.90	10.251.1.2	TCP	66	1817 → 80 [SYN] Seq=0 Win=64240 Len=0 MSS=1460 WS…
2	0.007271	10.251.1.2	10.251.1.90	TCP	66	80 → 1817 [SYN, ACK] Seq=0 Ack=1 Win=29200 Len=0 …
3	0.007312	10.251.1.90	10.251.1.2	TCP	54	1817 → 80 [ACK] Seq=1 Ack=1 Win=65536 Len=0
4	0.007491	10.251.1.90	10.251.1.2	HTTP	385	GET / HTTP/1.1
5	0.007666	10.251.1.2	10.251.1.90	TCP	60	80 → 1817 [ACK] Seq=1 Ack=332 Win=30336 Len=0

Question: What do they mean?

_____

_____

Question: What are these three packets for?

_____

_____

12. Select the entry that lists “**GET / HTTP/1.1**” in the Info column. Five entries will appear in the “Packet Details” section including Frame details and Data Link, Network, Transport and Application headers.

Question: What protocols are listed in “Frame details” section and what OSI model layers do they belong to?

Data Link header:

Network header:

Transport header:

Application header:

13. Click, then expand the “**Ethernet II**” entry.

Question: What is the length of the header?

TIP: You can see the header length at the very bottom of the window in the status bar.

Example:

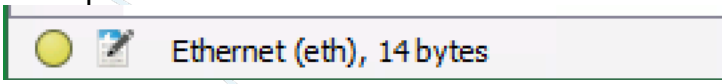

Question: What are the values of Destination and Source fields?

Question: What is the Type value (also known as Ethertype)?

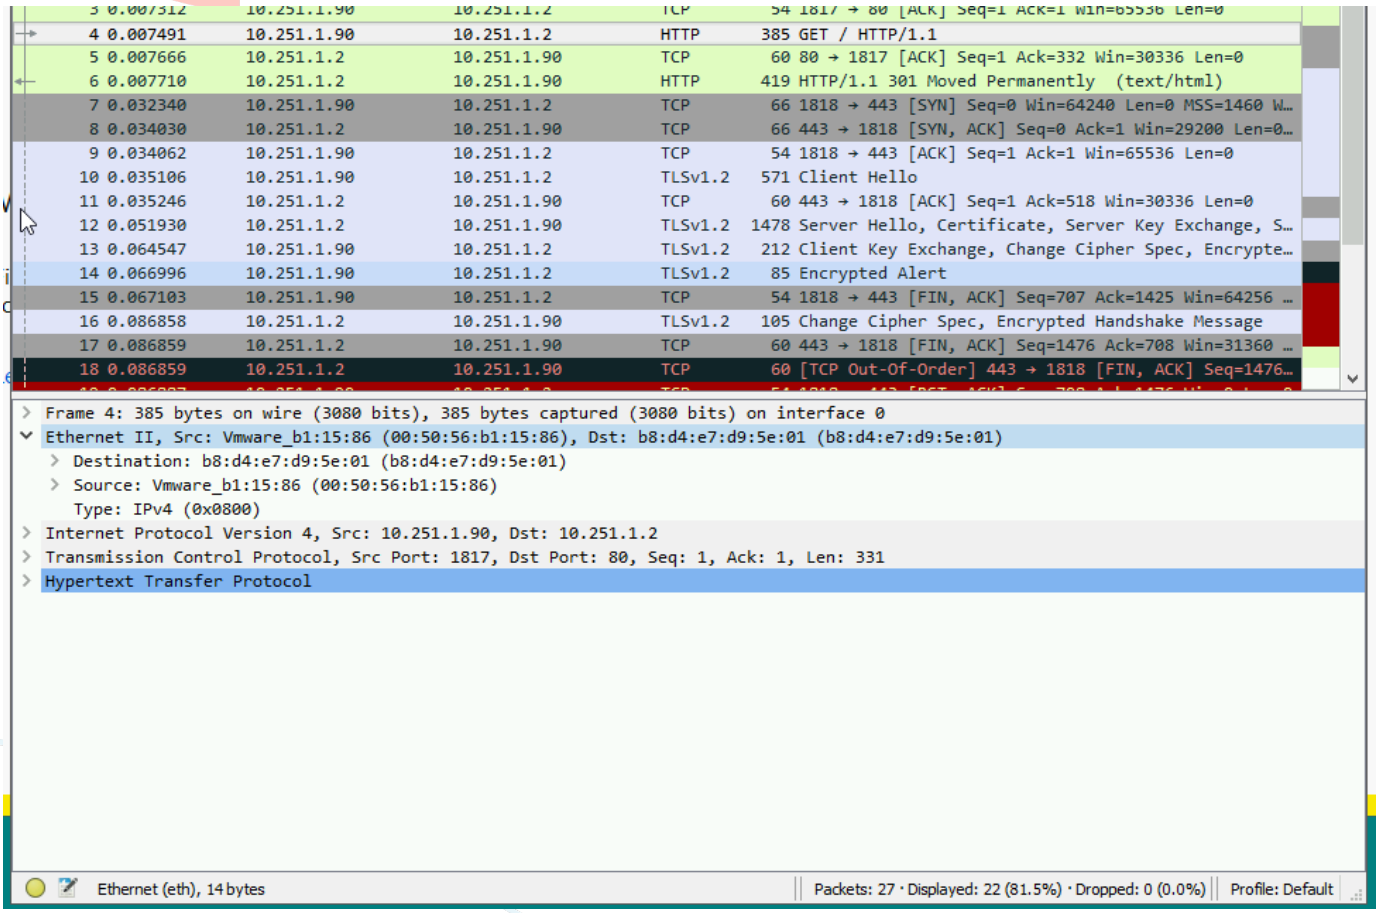

14. Click, then expand the "**Internet Protocol Version 4**" entry.

Question: What is the length of the header?

Question: What is the protocol version?

Question: What is the Time to live value?

Question: What is the Protocol number?

Question: What does the IP protocol number represent and what is its main purpose of this field?

Answer: IP protocol number or Protocol for short, is a numeric identification of the upper layer protocol contained in the packet's payload. The IANA has assigned unique values to each IP protocol, e.g. ICMP is IP protocol 1, TCP is 6, UDP is 17 and GRE is 47.

Question: What are the values of the Destination and Source fields?

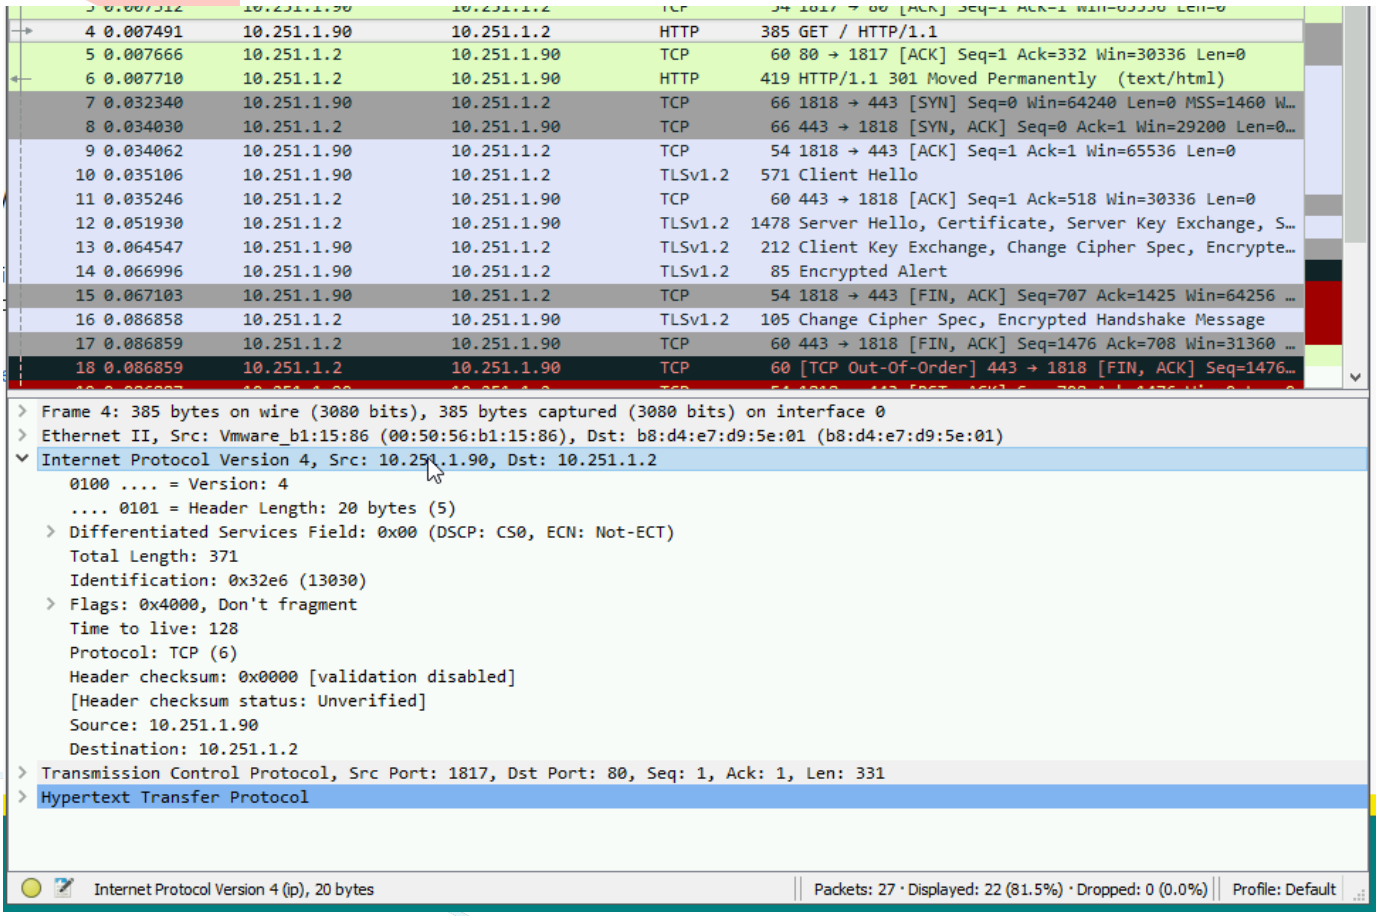

15. Click, then expand the "**Transmission Control Protocol**" entry.

Question: What is the length of the header?

Question: What are the first two fields?

Question: What are they for?

Question: What is the sequence number for?

16. Expand "**Flags**". Examine the flags while reviewing their meaning/usage below:

- Acknowledgement: Indicates that the acknowledgement field is significant. All packets after the initial SYN packet sent by the client should have this flag set.
- Reset: Reset the connection. Seen on rejected connections.
- Syn: Synchronize the sequence numbers. Seen on new connections.
- Fin: No more data from sender. Seen after a connection is closed.

For more detailed information on flags do an internet search on “TCP Flags”, or search for “RFC 9293”

Question: What is the Window size?

Question: What is the Window size for?

Answer: The window size field is the number of bytes the sender will buffer for the response. During 3-way handshake both sender and receiver will say how large their receive window is.

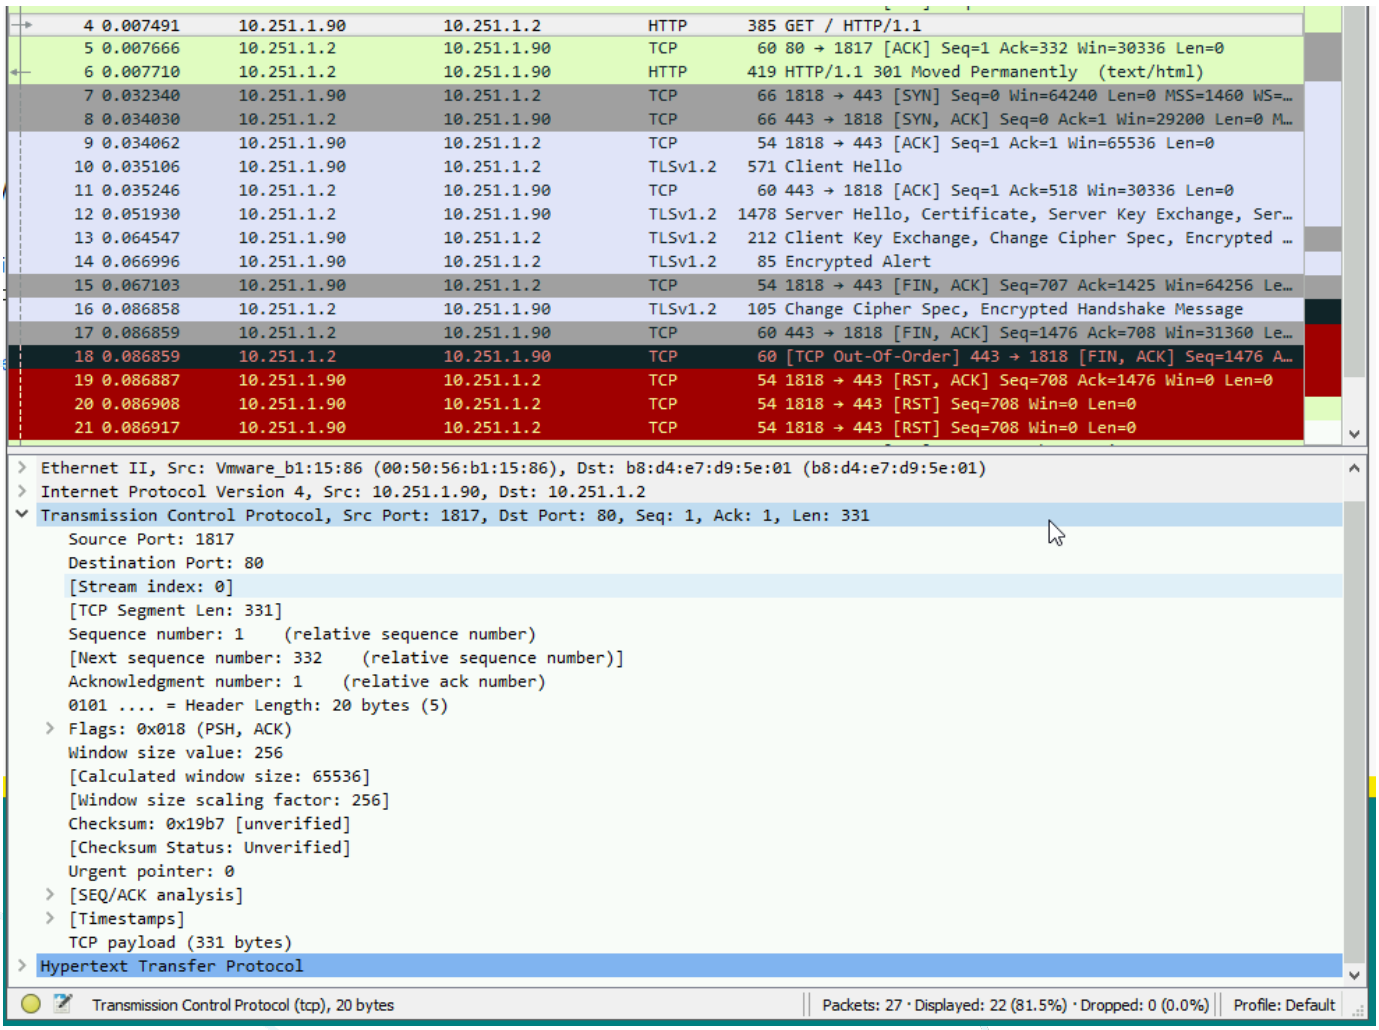

17. Expand the “Hypertext Transfer Protocol” entry.

The Hyper Text Transfer Protocol (HTTP) header accommodates four main commands: GET, PUSH, PUT and DELETE. Usually after the 3-way handshake, the first HTTP payload has a GET instruction, to download the web page.

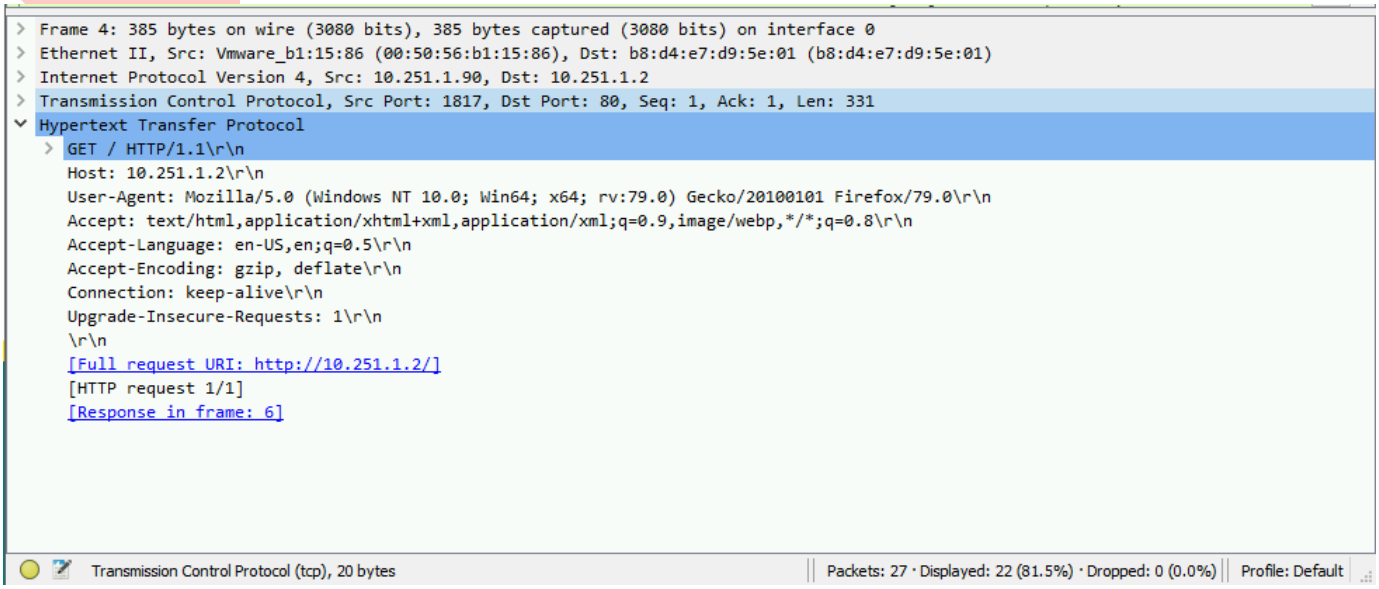

18. Scroll to the end of the session.

After the page request, the server sends many packets. These are acknowledged by the client and displayed as the black with red entries (image below). They contain the actual web page content. Once the page is fully loaded in the browser there is a FIN segment coming from the client signaling the end of the session. It is followed by a similar one from the server. Finally, the client sends that last ACK packet.

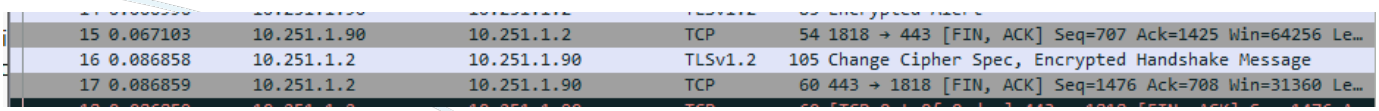

This concludes the task related to TCP HTTP connections.

Task 2: UDP header – DHCP Example

Now you will investigate a UDP header and compare it with the TCP header.

In this investigation you will explore the Dynamic Host Configuration Protocol (DHCP), which is used to assign IP addresses to devices.
In the remote lab environment, the MGMT PC is connected to the Out-Of Band Management (OOBM) port of the lab switches.
MGMT PC has been configured with a DHCP server to dynamically provide your lab switches with an IP address on their mgmt (Management) interface.

Objectives

- Understand the UDP Header

Steps

First you will start the Wireshark capture on the MGMT PC. Then you will stop and start the DHCP client on your sw-edge2, to generate DHCP traffic. This traffic will be captured by the Wireshark on MGMT PC.

1. On MGMT PC, in the Wireshark application, click the **start** button.

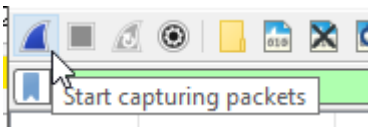

2. Click "**Continue without Saving**" button. This will discard the current packet capture list.
3. In the Wireshark display filter, enter **bootp** and press ENTER. This will ensure that only DHCP or Bootp packets are displayed in the packet list.
4. Use the remote lab interface to open a console connection to sw-edge1.

> **NOTE:** You must use sw-edge1 for this step, the other switches have a static IP address on their management interface, while sw-edge1 still has the default DHCP client active on the management interface.

5. You can login with username **admin** and blank password (just press ENTER).
6. Access the configuration, disable and enable the mgmt interface. This is commonly referred to as bouncing an interface (disable and enable).

```
configure
 interface mgmt
  shutdown
  no shutdown
  end
```

```
6300# configure
```

```
6300(config)# interface mgmt
6300(config-if-mgmt)# shutdown
6300(config-if-mgmt)# no shutdown
6300(config-if-mgmt)# end
6300#
```

7. On **MGMT PC**, wait until 4 DHCP packets are shown in the packet list.
8. Stop the Wireshark capture by clicking on the stop button in the toolbar.
9. Click on the first packet (DHCP Discover).
10. Select and expand the "**User Datagram Protocol**" entry in the Packet Details section.

Question: What is the length of the header?

Question: What is the first impression when comparing with the TCP header (Task 1 step 15)?

Question: What fields do they have in common?

Question: Can you see any Acknowledgment flag embedded in the header?

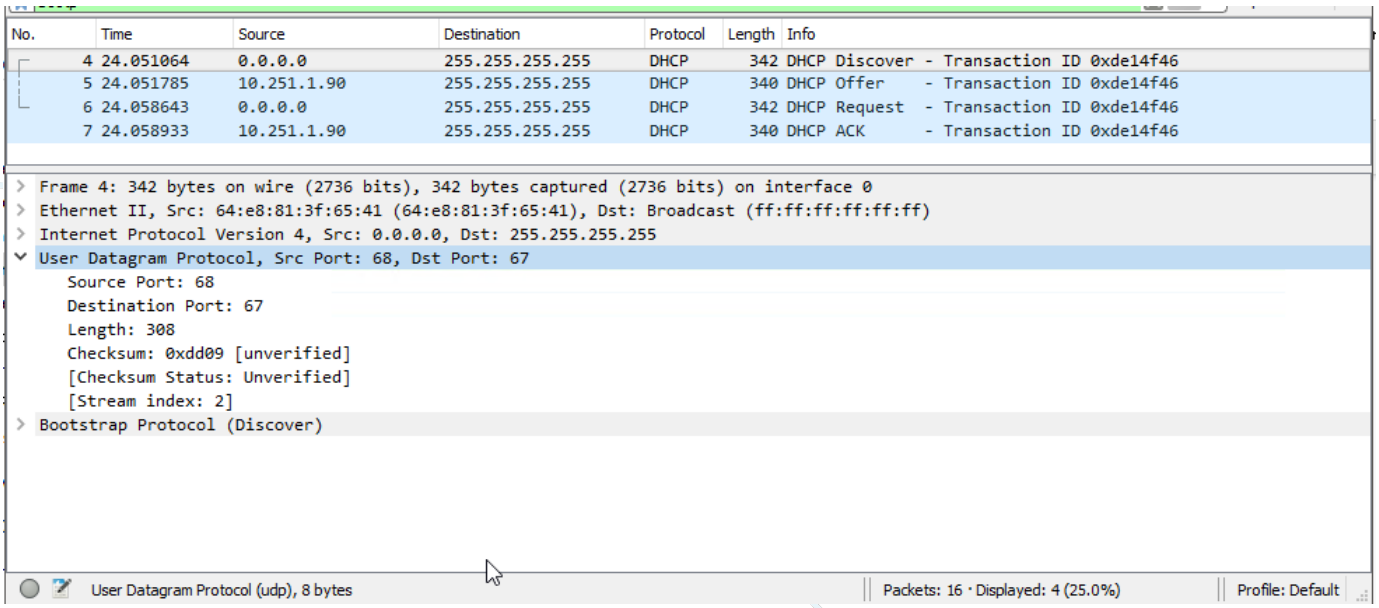

11. Click and expand the "**Bootstrap Protocol (Discover)**" entry.

Question: What is the DHCP Message type?

Answer: Discover.

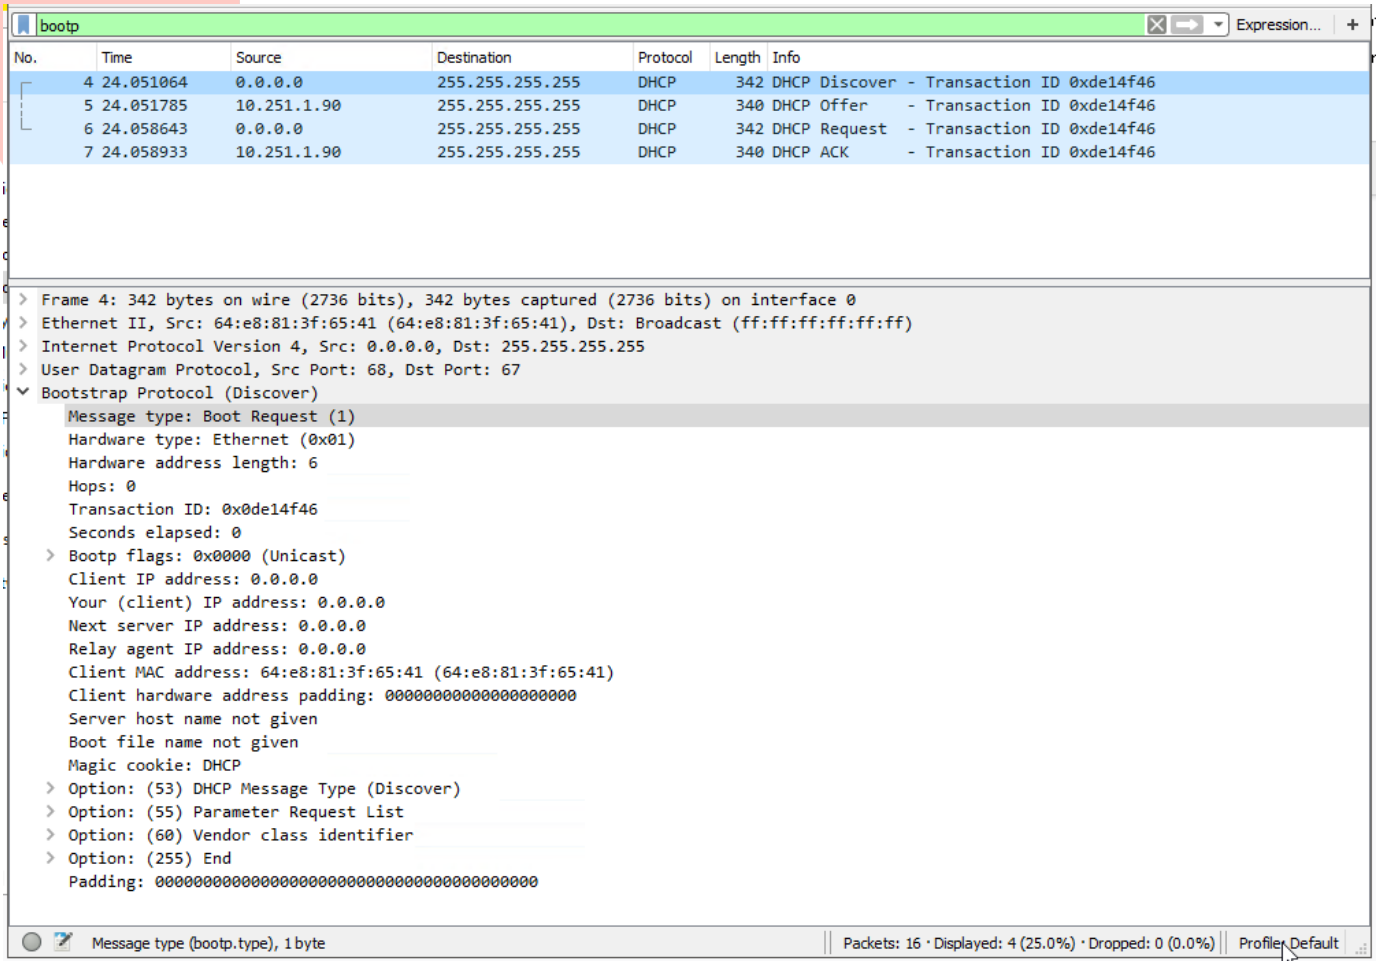

12. Click on each of the following 3 DHCP packets.

Question: What are the other 3 DHCP packet types in the trace?

Answer: After the Discover packet, the next packets are Offer, Request, Acknowledge.

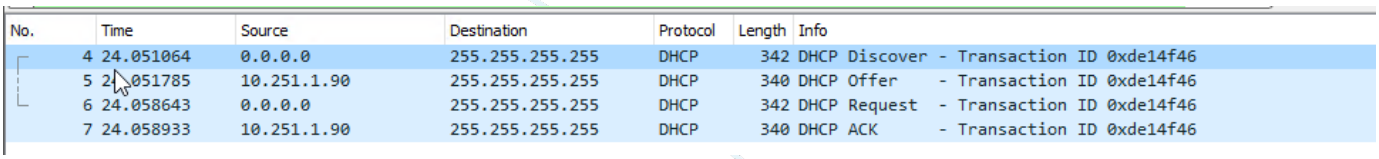

Due the lack of acknowledgement at the transport level, some UDP based applications do support the feature at Layer 7 level.
When the DHCP clients sends a SELECT packet to the DHCP server, the server confirms the selection by sending an ACKNOWLEDGE to the client.

Also notice how, unlike TCP, the transmission suddenly stops without any FIN signaling at the transport layer. This is because at the application layer level the DHCP process knows that the packet exchange is complete after the ACK.

You have completed this lab!

Lab 02.01 – Switching Fundamentals

Overview

In this lab activity you will start to work on the new customer environment. Your senior colleague has already configured the aggregation switches in the new setup. You must learn how to add a factory default switch to this network.

The sw-edge1 is currently at factory default state. Your colleague wants you to understand how you must perform the initial configuration and VLAN configuration on an Aruba AOS-CX access switch.

In the current stage of the labs, there is no need to consider redundancy. The new switch will be connected with a single link in this lab.

Objectives

After completing this lab, you will be able to:

- Gain initial access to a lab switch
- Perform basic switch configuration and validation
- Configure and verify an uplink connection to an aggregation switch
- Configure a switch port as a VLAN access port
- Configure a VLAN trunk on an uplink port
- Get ample practice with VLAN port configurations

Task 1: Initial Access to Edge switch

Your colleague has asked you to add a new switch to the network. In this task, the initial configuration will be applied to a switch, sw-edge1, in its factory default state. The second edge switch (sw-edge2) and the aggregation switches have already been configured for you, so this task only applies to sw-edge1.

Objectives

- Set the admin password.
- Configure Out-of-Band Management interface.
- Use the OOBM IP address to make an SSH connection to the switch.

Steps

1. Use the lab dashboard to open the console of **sw-edge1** and press <ENTER>.
2. Login with username **admin**, blank password (no password).

> **IMPORTANT:** Make sure to follow the requested passwords of the lab guide, some configurations and scripts depend on the correct password. For ease of use we use a simple password in the lab (Aruba123!), please never use a simple password in real life!

3. Review the default running configuration with the **show running-configuration** command.

```
show running-config
```

```
6300# show running-config
Current configuration:
!
!Version ArubaOS-CX FL.10.09.1040
!export-password: default
user admin group administrators password ciphertext
AQBapUO8pKgfd1UOrOKxkHe65hSaTMZPhmtPrSdYPwREeA6UYgAAAHwShgpkD5jJYm9p1nV5y1aJBwiVIm18U
sTiISeScMviUSwi8ON2rDxfNm+Gb+wWcL6nk+Qara/ylFnO/K/2Bgvqvq8QRkzzXNpnSAQaEWJ1M4BsHaSAI8
krMucrAIPa/89F
ntp server pool.ntp.org minpoll 4 maxpoll 4 iburst
ntp enable
!
!
...
```

4. Access the configuration mode.

```
config
```

```
6300# configure
6300(config)#
```

5. Change the admin password to **Aruba123!**

```
user admin password plaintext Aruba123!
```

```
6300(config)# user admin password plaintext Aruba123!
```

6. Change the hostname to sw-edge1.

```
hostname sw-edge1
```

```
6300(config)# hostname sw-edge1
sw-edge1(config)#
```

7. Access the management interface context (OOBM). The MGMT PC system has an IP interface in this subnet. Enable the interface and configure static IP address 10.251.1.4/24.

```
interface mgmt
    no shutdown
    ip static 10.251.1.4/24
end
```

```
sw-edge1(config)# interface mgmt
sw-edge1(config-if-mgmt)#     no shutdown
sw-edge1(config-if-mgmt)#     ip static 10.251.1.4/24
sw-edge1(config-if-mgmt)# end
sw-edge1#
```

NOTE: The OOBM port is not a commonly used feature in a Campus deployment, but it is convenient in a training remote lab environment since you can break the 'normal' network and still have OOBM SSH access to the switches from the MGMT PC desktop in the lab.

8. Review the interface mgmt configuration.

```
show interface mgmt
```

Example output, your output may be different.

```
sw-edge1# show interface mgmt
  Address Mode                  : static
  Admin State                   : up
  Link State                    : up
  Mac Address                   : 64:e8:81:dd:81:81
  IPv4 address/subnet-mask      : 10.251.1.4/24
  Default gateway IPv4          :
  IPv6 address/prefix           :
  IPv6 link local address/prefix: fe80::66e8:81ff:fedd:8181/64
  Default gateway IPv6          :
  Primary Nameserver            :
  Secondary Nameserver          :
  Tertiary Nameserver           :
```

9. Verify connectivity using a ping to MGMT PC (10.251.1.90) in the mgmt VRF.

```
ping 10.251.1.90 vrf mgmt
```

```
sw-edge1# ping 10.251.1.90 vrf mgmt
PING 10.251.1.90 (10.251.1.90) 100(128) bytes of data.
108 bytes from 10.251.1.90: icmp_seq=1 ttl=128 time=0.582 ms
108 bytes from 10.251.1.90: icmp_seq=2 ttl=128 time=0.639 ms
108 bytes from 10.251.1.90: icmp_seq=3 ttl=128 time=0.560 ms
108 bytes from 10.251.1.90: icmp_seq=4 ttl=128 time=0.465 ms
108 bytes from 10.251.1.90: icmp_seq=5 ttl=128 time=0.550 ms

--- 10.251.1.90 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4094ms
rtt min/avg/max/mdev = 0.465/0.559/0.639/0.056 ms
sw-edge1#
```

NOTE: The mgmt interface is automatically bound to a dedicated VRF (routing table) so it is completely isolated from the front-facing network ports of the switch.

10. Make a local backup (checkpoint) of the current configuration.

```
copy running-config checkpoint oobm
```

```
sw-edge1# copy running-config checkpoint oobm
Copying configuration: [Success]
```

11. Review the list of checkpoints (local configurations); make sure there is a checkpoint named *oobm.*

```
show checkpoint
```

Example, your output may be different.

```
sw-edge1# show checkpoint
NAME                        TYPE        WRITER  DATE(YYYY/MM/DD)      IMAGE VERSION
oobm                        latest      User    2022-10-15T17:52:30Z  FL.10.09.1040
```

Verify MGMT PC IP access to sw-edge1 using OOBM Network

In the next steps you will review the configured IP address on MGMT PC. You then connect to sw-edge1, via SSH, to the management IP address you configured in the previous steps.

12. Open a connection to MGMT PC.
13. On MGMT PC, open a command prompt (**cmd**).
14. Run **ipconfig** to review the current interfaces and IP address assignments.

You should see these 2 interfaces:

OOBM 10.**251**.1.90 (only for OOBM access)

	no default gateway, isolated subnet to the management ports of the switches.
Lab NIC Core	10.**254**.1.90 In-band access to the devices via the core router.

```
Ethernet adapter OOBM:

   Connection-specific DNS Suffix  . :
   IPv4 Address. . . . . . . . . . . : 10.251.1.90
   Subnet Mask . . . . . . . . . . . : 255.255.255.0
   Default Gateway . . . . . . . . . : 10.251.1.254

Ethernet adapter Lab NIC - Core:

   Connection-specific DNS Suffix  . :
   IPv4 Address. . . . . . . . . . . : 10.254.1.90
   Subnet Mask . . . . . . . . . . . : 255.255.255.0
   Default Gateway . . . . . . . . . :
```

15. Use Putty or MTPutty to open an SSH connection to the OOBM IP address of sw-edge1 (**10.251.1.4**). If prompted, **accept** the SSH key warning when .

> **IMPORTANT:** Be careful with the SSH key warning. You should only get this warning when the remote system is new or when it has re-generated its SSH host keys. Otherwise, some intermediate system may be trying to attack your systems.

16. Login with **admin / Aruba123!** credentials.

Task 2: Configure Edge switch uplink to Aggregation switch

The sw-edge1 switch currently does not have any active uplink interfaces. In this task you will enable the uplink port to the aggregation switch sw-agg1.

Objectives

- Configure and verify the uplink connection.
- Use LLDP to verify the connections.

Steps

1. Use MGMT PC to open an SSH connection to sw-edge1.
2. On sw-edge1, access the configuration mode and disable ports that connect to other switches (1/1/25-1/1/28) using the **shutdown** command.

```
config
interface 1/1/25-1/1/28
 shutdown
 exit
```

```
sw-edge1# config
sw-edge1(config)# interface 1/1/25-1/1/28
sw-edge1(config-if-<1/1/25-1/1/28>)#  shutdown
sw-edge1(config-if-<1/1/25-1/1/28>)#  exit
```

3. On sw-edge1, enable port 1/1/27, which connects to sw-agg1.

```
interface 1/1/27
 no shutdown
```

```
sw-edge1(config)# interface 1/1/27
sw-edge1(config-if)#  no shutdown
```

4. Configure port 1/1/27 with description **sw-agg1**

```
 description sw-agg1
```

```
sw-edge1(config-if)#  description sw-agg1
```

5. Review the updated configuration of the current switch port and exit the interface context. Exiting the interface context brings you back to the global configuration context.

```
 show running-config current-context
 exit
```

```
sw-edge1(config-if)#  show running-config current-context
interface 1/1/27
    no shutdown
    description sw-agg1
    no routing
```

```
    vlan access 1
    exit
sw-edge1(config-if)#  exit
```

6. Review the brief interface status. Port 1/1/27 should now have the “up” status and the updated description.

```
show interface brief
```

```
...
1/1/26    1       access SFP+DAC1       no      down    Administratively down  --
--
1/1/27    1       access SFP+DAC1       yes     up                             10000
sw-agg1
1/1/28    1       access SFP+DAC1       no      down    Administratively down  --
--
vlan1     --              --            yes     up                             --
--
...
```

7. Verify sw-agg1 is now listed as an LLDP neighbor.

```
show lldp neighbor-info
```

```
sw-edge1(config)# show lldp neighbor-info

LLDP Neighbor Information
=========================

Total Neighbor Entries          : 6
Total Neighbor Entries Deleted  : 2
Total Neighbor Entries Dropped  : 0
Total Neighbor Entries Aged-Out : 2

LOCAL-PORT  CHASSIS-ID         PORT-ID                    PORT-DESC
TTL      SYS-NAME
-------------------------------------------------------------------------------------
----------------------
1/1/1       68:b5:99:a3:a1:c0  11                         11
120      P28-PC1-Fanout
1/1/3       ec:b1:d7:1b:07:00  ec:b1:d7:1b:07:02          25
120
1/1/27      b8:d4:e7:da:80:00  1/1/1                      1/1/1
120      sw-agg1
mgmt        00:23:89:d6:7d:e7  GigabitEthernet2/0/21      T11-Edge-1-OOBM
120      P28-OOBM-Fanout
```

Question: On what port is sw-edge1 connected on sw-agg1?

Answer: Port 1/1/1

> **NOTE:** There may be LLDP neighbors with sys-name that contains ‘Fanout’. These are switches that connect the lab VMs to the switch ports of your lab devices.

NOTE: Since there is no direct connection between the lab PC systems and the switch ports, shutting down the switch port will not be visible to the Lab VM PC, so you may need to manually refresh the IP configuration.

8. Verify the LLDP neighbor details for port 1/1/27.

```
show lldp neighbor-info 1/1/27
```

```
sw-edge1(config)# show lldp neighbor-info 1/1/27

Port                            : 1/1/27
Neighbor Entries                : 1
Neighbor Entries Deleted        : 0
Neighbor Entries Dropped        : 0
Neighbor Entries Aged-Out       : 0
Neighbor System-Name            : sw-agg1
Neighbor System-Description     : Aruba JL635A  GL.10.09.1040
Neighbor Chassis-ID             : b8:d4:e7:da:80:00
Neighbor Management-Address     : 10.1.1.1
Chassis Capabilities Available  : Bridge, Router
Chassis Capabilities Enabled    : Bridge, Router
Neighbor Port-ID                : 1/1/1
Neighbor Port-Desc              : 1/1/1
Neighbor Port VLAN ID           : 1
TTL                             : 120

Neighbor EEE information        : DOT3
Neighbor TX Wake time           : 0 us
Neighbor RX Wake time           : 0 us
Neighbor Fallback time          : 0 us
Neighbor TX Echo time           : 0 us
Neighbor RX Echo time           : 0 us
```

Question: What is the system description for the sw-agg1?

Answer: Aruba switches report their product code as part of the system description. In this example JL635A (8325).

9. PC1 is connected to port 1/1/1, configure description ‘pc1’ on port 1/1/1.

```
interface 1/1/1
 description pc1
 exit
```

```
sw-edge1(config)# interface 1/1/1
sw-edge1(config-if)#  description pc1
sw-edge1(config-if)#  exit
```

10. AP1 is connected to port 1/1/2, configure description ‘ap1’ on port 1/1/2.

```
interface 1/1/2
 description ap1
 exit
```

```
sw-edge1(config)# interface 1/1/2
sw-edge1(config-if)#  description ap1
sw-edge1(config-if)#  exit
```

Task 3: Configure PC port as VLAN Access port

In this task you will create a new VLAN on sw-edge1 and assign the PC connected to port 1/1/1 to this new VLAN.

Objectives

- Create a VLAN
- Assign an access port to a VLAN
- Verify the configuration

Steps

1. Using MGMT PC, open an SSH connection to sw-edge1 and access the configuration mode.

> **NOTE:** Use the command **config** to access the configuration mode; this will not be repeated for every task in the guide.

2. On sw-edge1, review the default list of VLANs.

```
show vlan
```

```
sw-edge1(config)# show vlan

--------------------------------------------------------------------------------
-----
VLAN  Name                           Status  Reason                 Type
Interfaces
--------------------------------------------------------------------------------
-----
1     DEFAULT_VLAN_1                 up      ok                     default     1/1/1-
1/1/28
```

Question: What VLANs exist by default on the switch?

Answer: Only VLAN 1 exists by default on the switch. It is the default VLAN and it cannot be removed.

3. Use the **show interface brief** output to review the default port status for port 1/1/1, which connects to PC1.

```
show interface brief
```

```
sw-edge1(config)# show interface brief
--------------------------------------------------------------------------------
------------------------
Port      Native  Mode   Type           Enabled Status  Reason                 Speed
Description
```

```
          VLAN                                                                  (Mb/s)
--------------------------------------------------------------------------------
------------------------
1/1/1     1       access 1GbT           yes     up                              1000
pc1
1/1/2     1       access 1GbT           yes     up                              1000
ap1
```

Question: What is the VLAN mode on a port by default?

Answer: Access

Question: What is the Native VLAN on a port by default?

Answer: VLAN 1

4. Create VLAN 11, assign the VLAN name **v11-employee** and revert to the global context.

```
vlan 11
 name v11-employee
 exit
```

```
sw-edge1(config)# vlan 11
sw-edge1(config-vlan-11)#  name v11-employee
sw-edge1(config-vlan-11)#  exit
```

5. List the VLANs to verify that the new VLAN was created.

```
show vlan
```

```
sw-edge1(config)# show vlan

--------------------------------------------------------------------------------
-----
VLAN  Name                          Status  Reason                  Type
Interfaces
--------------------------------------------------------------------------------
-----
1     DEFAULT_VLAN_1                up      ok                      default     1/1/1-
1/1/28
11    v11-employee                  down    no_member_port          static
```

6. Review member ports of VLAN 11.

```
show vlan 11
```

```
sw-edge1(config)# show vlan 11

--------------------------------------------------------------------------------
-----
VLAN  Name                          Status  Reason                  Type
Interfaces
```

```
--------------------------------------------------------------------------------
-----
11    v11-employee              down    no_member_port          static
```

7. Review the MAC-address table for VLAN 11.

```
show mac-address-table vlan 11
```

```
sw-edge1(config)# show mac-address-table vlan 11
No MAC entries found.
```

8. Assign the port connected to PC1 (1/1/1) to VLAN 11 as an access port.

```
interface 1/1/1
 vlan access 11
 exit
```

```
sw-edge1(config)# interface 1/1/1
sw-edge1(config-if)#  vlan access 11
sw-edge1(config-if)#  exit
```

9. Verify your configuration by reviewing the port membership of VLAN 11. Port 1/1/1 should now be member of VLAN 11.

```
show vlan 11
```

```
sw-edge1(config)# show vlan 11

--------------------------------------------------------------------------------
-----
VLAN  Name                      Status  Reason                  Type
Interfaces
--------------------------------------------------------------------------------
-----
11    v11-employee              up      ok                      static      1/1/1
```

10. Verify that port 1/1/1 is no longer member of VLAN 1.

```
show vlan 1
```

```
sw-edge1(config)# show vlan 1

--------------------------------------------------------------------------------
-----
VLAN  Name                      Status  Reason                  Type
Interfaces
--------------------------------------------------------------------------------
-----
1     DEFAULT_VLAN_1            up      ok                      default     1/1/2-
1/1/28
```

11. Alternatively, you can also review the VLAN port configuration.

```
show vlan port 1/1/1
```

```
sw-edge1(config)# show vlan port 1/1/1

--------------------------------------------------------------------------------
VLAN  Name                             Mode           Mapping
--------------------------------------------------------------------------------
11    v11-employee                     access         port
```

Task 4: Configure Uplink port as VLAN Trunk Port

In this task the sw-edge1 uplink port to sw-agg1 will be configured to support multiple VLANs. This is known as a VLAN trunk. By default, VLAN 1 is the native (untagged) VLAN. You will add VLAN 11 as a tagged VLAN.

Objectives

- Configure a VLAN trunk
- Allow VLANs on a VLAN trunk
- Verify the configuration

Steps

1. On sw-edge1, verify you are in the configuration mode. Review the current configuration of uplink port 1/1/27.

```
interface 1/1/27
 show running-config current-context
```

```
sw-edge1(config)# interface 1/1/27
sw-edge1(config-if)#  show running-config current-context
interface 1/1/27
    no shutdown
    description sw-agg1
    no routing
    vlan access 1
    exit
```

2. Configure the uplink as VLAN Trunk, allow VLAN 11.

```
 vlan trunk allowed 11
```

```
sw-edge1(config-if)#  vlan trunk allowed 11
```

3. Review the updated port configuration.

```
show running-config current-context
```

```
sw-edge1(config-if)# show running-config current-context
interface 1/1/27
    no shutdown
    description sw-agg1
    no routing
    vlan trunk native 1
    vlan trunk allowed 11
    exit
```

NOTE: You may also check the interface configuration from any other context using the command: **show running-config interface** <member/slot/port>.

```
show running-config interface 1/1/27
```

4. Review the list of allowed VLANs on the trunk port.

```
show vlan port 1/1/27
```

```
sw-edge1(config-if)# show vlan port 1/1/27

---------------------------------------------------------------------------
VLAN  Name                            Mode            Mapping
---------------------------------------------------------------------------
11    v11-employee                    trunk           port
```

Question: What VLANs are currently allowed on port 1/1/27?

Answer: Only VLAN 11 is allowed on the port.
Note
When converting a port from access to trunk using the **vlan trunk allow** command, only the VLANs in the command are allowed on the port.

5. Allow VLAN 1 on the trunk.

```
 vlan trunk allow 1
```

```
sw-edge1(config-if)#  vlan trunk allow 1
```

6. Review the list of allowed VLANs on the trunk port.

```
show vlan port 1/1/27
```

```
sw-edge1(config-if)# show vlan port 1/1/27

---------------------------------------------------------------------------
VLAN  Name                            Mode            Mapping
---------------------------------------------------------------------------
1     DEFAULT_VLAN_1                  native-untagged port
11    v11-employee                    trunk           port
```

Question: Did the allow VLAN 1 command overwrite the existing list of allowed VLANs?

Answer: No

7. Simplify the allowed VLAN configuration by allowing all VLANS.

```
vlan trunk allowed all
```

```
sw-edge1(config-if)# vlan trunk allowed all
```

Question: Does the **allowed all** command imply that any VLAN will be passing on this VLAN trunk?

Answer: No, it means that all VLANs that have been defined on this switch will be allowed. When a VLAN doesn't exist on the switch global configuration, traffic marked with this VLAN ID will not be accepted on the VLAN trunk.

Question: Will this apply to future VLANs?

Answer: Yes, any VLAN that would be created on this switch will automatically be allowed on this VLAN trunk.

> **NOTE:** Allowing all VLANs is easy for a lab environment, but it is typically not a best practice for a production environment. Only the required VLANs will be allowed on the VLAN trunk ports in a production environment.

8. Review and compare the VLAN port configurations of ports 1/1/1 and 1/1/27. Notice the different modes (trunk and access) between the port VLANs.

```
show vlan port 1/1/1
show vlan port 1/1/27
```

```
sw-edge1(config-if)# show vlan port 1/1/1

--------------------------------------------------------------------------------
VLAN  Name                                Mode            Mapping
--------------------------------------------------------------------------------
11    v11-employee                        access          port
```

```
sw-edge1(config-if)# show vlan port 1/1/27

--------------------------------------------------------------------------------
VLAN  Name                                Mode            Mapping
--------------------------------------------------------------------------------
1     DEFAULT_VLAN_1                      native-untagged port
11    v11-employee                        trunk           port
```

Question: What does the 'native-untagged' mode represents for the VLAN 1?

Answer: Traffic in VLAN 1 will traverse this VLAN trunk as untagged ethernet frames, without a VLAN tag, just like an access port with VLAN 1.

Bounce the Lab NIC on PC1

To renew the IP address on PC1, you can bounce the Lab NIC (disable / enable).

9. On PC1, right-click the **network** icon in the status bar.
10. Click **Open Network & Internet Settings**.

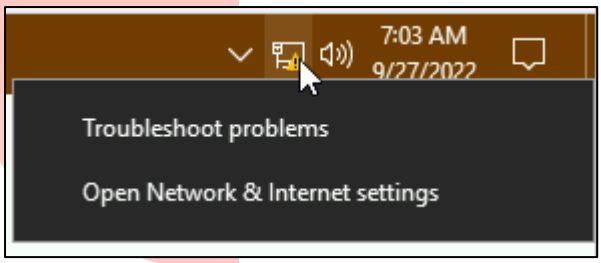

11. Under Change your network settings, click **Change adapter options**.

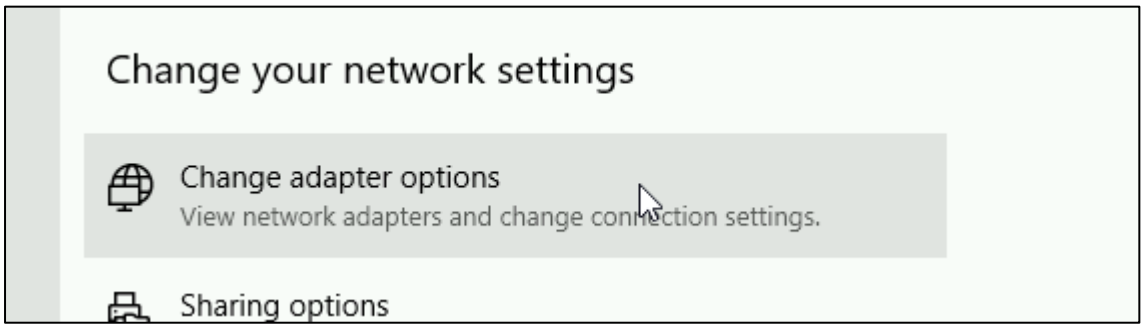

12. Right-click the **Lab NIC** (this is the NIC connected to port 1 on sw-edge1).
13. Click **Disable**.
14. Right-click the **Lab NIC**, click **Enable**.

Verify the MAC Address Table on sw-edge1

15. On sw-edge1, review the MAC-address table for VLAN 11.

```
show mac-address-table vlan 11
```

Example, your output may be different.

```
sw-edge1(config-if)# show mac-address-table vlan 11
MAC age-time            : 300 seconds
Number of MAC addresses : 2

MAC Address          VLAN     Type                        Port
-------------------------------------------------------------
00:50:56:b1:cd:26    11       dynamic                     1/1/1
b8:d4:e7:da:80:00    11       dynamic                     1/1/27
```

Question: Do you see MAC addresses on the uplink port? How Many?

Answer: Yes, the sw-agg1 MAC address.

Task 5: Practice VLAN port configuration

In this task you will practice the VLAN port configuration.

Configure these items:

- Prepare a VLAN trunk between sw-edge1 port 1/1/25 and sw-edge2 port 1/1/26 and enable the ports.
- On sw-edge1, set the port 1/1/25 description to ‘sw-edge2’.
- Ensure that only VLANs 1 and 11 are passing on this link.
- On sw-edge2, assign the port that connects to PC4 (1/1/4) to VLAN 11. Verify the port is enabled and up.

Verify the configuration using these steps:

- Sw-edge1 port 1/1/25 is connected to port 1/1/26 on edge2 using LLDP.
- PC4 is getting an IP address in VLAN 11.
- Verify mac-address tables on sw-edge2, sw-edge1, sw-agg1.
- Save the configuration after successful validation.

Objectives

- Practice the VLAN port configuration.

Steps

There are no steps in this task; use the commands you learned in the previous tasks to configure the objectives.
Make sure to use the lab dashboard to open a console connection to sw-edge2 as well to apply the configuration.

Configuration solution

sw-edge1

```
int 1/1/25
 description sw-edge2
 vlan trunk allowed 1,11
 no shutdown
 exit
```

sw-edge2

```
vlan 11
int 1/1/26
 description sw-edge1
 vlan trunk allowed 1,11
 no shutdown
 exit

int 1/1/4
 description pc4
 vlan access 11
 exit
```

Verify the configuration (both switches)

```
show lldp neighbor
show lldp neighbor 1/1/26
```

```
show mac-address-table vlan 11
```

1. Save the configurations on both edge switches.

```
write memory
```

You have completed this Lab!

Lab 02.02 – Link Aggregation

Overview

In the previous lab you have added a factory default switch to the existing network. Your senior colleague has now asked you to learn about the bundling of links between devices.

This lab you will introduce you to the aggregation of links, also known as a LAG. Link-aggregation provides additional bandwidth capacity and redundancy in case of link failure. Your sw-edge1 and sw-edge2 have two links between each other. You will aggregate these two links into a LAG.

Objectives

After completing this lab, you will be able to:

- Configure link-aggregation
- Verify link aggregation

Task 1: Configure Edge1 with Link-Aggregation to Edge2

In this task, you will configure sw-edge1 with a link aggregation configuration to aggregate ports 1/1/25 and 1/1/26 into a logical port *LAG50.*

First a new logical or parent interface must be created, then the physical member ports must be linked to this logical interface. The parent interface is referred to as a LAG interface.

On sw-edge1, you will configure two ports to the sw-edge2 switch as a LAG bundle.

Objectives

- Verify and configure link aggregation.

Steps

1. Using MGMT PC, open an SSH connection to **sw-edge1** and enter the configuration mode.
2. Define a new Link-Aggregation (LAG) interface with LAG ID 50.

```
interface lag 50
```

```
sw-edge1(config)# interface lag 50
```

3. Review the interface list.

```
show interface brief
```

```
sw-edge1(config-lag-if)# show interface brief
--------------------------------------------------------------------------------
-----
Port      Native  Mode   Type           Enabled Status  Reason                 Speed
Description
          VLAN                                                                 (Mb/s)
--------------------------------------------------------------------------------
-----
1/1/1     11      access 1GbT           yes     up                             1000
pc1
1/1/2     1       access 1GbT           yes     up                             1000
ap1
1/1/3     1       access 1GbT           yes     up                             1000
--
...
lag50     1       access --             no      down    --                     auto
--
```

Question: Did the new interface show up in the list?

Answer: Yes, the lag interface is listed.

Question: Is the interface enabled?

Answer: No, by default it is disabled.

4. Enable the interface LAG 50.

```
no shutdown
```

```
sw-edge1(config-lag-if)# no shutdown
```

5. Review the aggregation mode.

```
show lacp aggregates
```

```
sw-edge1(config-lag-if)# show lacp aggregates

Aggregate name    : lag50
Interfaces        :
Heartbeat rate    : N/A
Hash              : l3-src-dst
Aggregate mode    : Off
```

Question: Based on what you learned in this chapter, what are the types of LAG aggregation modes?

Answer: A LAG can be static (no protocol) or dynamic (LACP protocol is used to validate the link connections).

6. Configure **LACP active** as the aggregation mode.

```
lacp mode active
```

```
sw-edge1(config-lag-if)# lacp mode active
```

7. Review the aggregation mode and verify it now shows the LACP mode as active.

```
show lacp aggregates
```

```
sw-edge1(config-lag-if)# show lacp aggregates

Aggregate name    : lag50
Interfaces        :
Heartbeat rate    : Slow
Hash              : l3-src-dst
Aggregate mode    : Active
```

At this point, the LAG does not contain any physical ports. In the next steps, you will configure port 1/1/25 as member port of the LAG 50.

8. Access the port 1/1/25 context and review the current configuration of the port. In the next steps you will link the port to the LAG 50 and check the impact on the current configuration of the port.

```
interface 1/1/25
 show running-config current-context
```

```
sw-edge1(config-lag-if)# interface 1/1/25
sw-edge1(config-if)#  show running-config current-context
interface 1/1/25
    no shutdown
    description sw-edge2
    no routing

    vlan trunk native 1
    vlan trunk allowed 11
    exit
```

9. Now link the port to LAG 50.

```
 lag 50
```

```
sw-edge1(config-if)# lag 50
```

10. Again, review the current port configuration.

```
show running-config current-context
```

```
sw-edge1(config-if)# show running-config current-context
interface 1/1/25
    no shutdown
    description sw-edge2
    lag 50
    exit
```

Question: What configuration changed after assigning the port to a LAG?

Answer: A LAG member port does not have a VLAN port configuration anymore, this is configured at the LAG parent port context.

11. Review the LACP interface status.

```
show lacp interfaces
```

```
sw-edge1(config-if)# show lacp interfaces

State abbreviations :
A - Active        P - Passive      F - Aggregable I - Individual
S - Short-timeout L - Long-timeout N - InSync     O - OutofSync
C - Collecting    D - Distributing
X - State m/c expired              E - Default neighbor state

Actor details of all interfaces:
------------------------------------------------------------------------------
Intf       Aggr       Port  Port  State   System-ID         System Aggr Forwarding
           Name       Id    Pri                             Pri    Key  State
------------------------------------------------------------------------------
1/1/25     lag50      26    1     ALFOE   64:e8:81:dd:81:80 65534  50   lacp-block


Partner details of all interfaces:
```

```
--------------------------------------------------------------------------------
Intf       Aggr       Port  Port  State   System-ID         System Aggr
           Name       Id    Pri                             Pri    Key
--------------------------------------------------------------------------------
1/1/25     lag50      0     0     PLFOEX  00:00:00:00:00:00 0      0
```

Question: The output shows *lacp-block*, why?

Answer: On AOS-CX switches, LACP requires the peer device to send LACP PDUs. In this case, no LACP PDUs arrive on the port and therefore the port remains in LACP-block state. In this lab, the peer switch sw-edge2 has not been configured yet, so this state is expected at this point in the lab.

12. Now add port 1/1/26 as member port of the LAG 50 and enable the port.

```
interface 1/1/26
 lag 50
 no shutdown
 exit
```

```
sw-edge1(config-if)# interface 1/1/26
sw-edge1(config-if)#  lag 50
sw-edge1(config-if)#  no shutdown
sw-edge1(config-if)#  exit
```

13. Review the LACP interface status.

```
show lacp interfaces
```

```
sw-edge1(config)# show lacp interfaces

State abbreviations :
A - Active        P - Passive      F - Aggregable I - Individual
S - Short-timeout L - Long-timeout N - InSync     O - OutofSync
C - Collecting    D - Distributing
X - State m/c expired              E - Default neighbor state

Actor details of all interfaces:
--------------------------------------------------------------------------------
Intf       Aggr       Port  Port  State   System-ID         System Aggr Forwarding
           Name       Id    Pri                             Pri    Key  State
--------------------------------------------------------------------------------
1/1/25     lag50      26    1     ALFOE   64:e8:81:dd:81:80 65534  50   lacp-block
1/1/26     lag50                                                        down


Partner details of all interfaces:
--------------------------------------------------------------------------------
Intf       Aggr       Port  Port  State   System-ID         System Aggr
           Name       Id    Pri                             Pri    Key
--------------------------------------------------------------------------------
1/1/25     lag50      0     0     PLFOEX  00:00:00:00:00:00 0      0
1/1/26     lag50
```

Question: What is the *Forwarding State*?

Answer: lacp-blocked and down.

14. Review the impact to PC4 network access.
15. On PC4, open a command prompt (cmd) and attempt to ping the default gateway IP address (10.1.11.1).

```
ping 10.1.11.1
```

Question: What is the result for clients connected to Edge2?

Answer: PC4, that is connected to sw-edge2 port 1/1/4, cannot reach the default gateway anymore since the LAG is in blocked state.

You have now completed 1 side (sw-edge1) of the configuration; in the next task you will configure the sw-edge2 side.

Task 2: Practice and verify the LAG configuration on Edge2

Objectives

In this task you will practice the LAG configuration steps on sw-edge2. Configure ports 1/1/25 and 1/1/26 in a LAG with ID 50 with LACP mode active. Make sure all ports are enabled and verify the status. Try to complete this task by yourself, only use the solution below if required.

Solution

On Sw-edge2:

```
interface lag 50
 no shutdown
 lacp mode active
 exit

interface 1/1/25,1/1/26
 lag 50
 no shutdown
 exit
```

Verification Steps

1. On sw-edge2, verify 1/1/25, 1/1/26 and LAG 50 are enabled.

```
show interface brief
```

```
sw-edge2(config)# show interface brief
--------------------------------------------------------------------------------
-----
Port      Native  Mode   Type          Enabled Status  Reason                Speed
Description
          VLAN                                                                 (Mb/s)
--------------------------------------------------------------------------------
-----
...
1/1/25    1       access SFP+DAC1      yes     up                            10000
sw-edge1
1/1/26    1       access SFP+DAC1      yes     up                            10000
sw-edge1
1/1/27    1       trunk  SFP+DAC1      no      down    Administratively down  --
sw-agg1
1/1/28    1       access SFP+DAC1      no      down    Administratively down  --
sw-agg2
vlan1     --              --           yes     up                            --
--
lag50     1       access --            yes     up      --                    20000
--
```

Question: What is the reported speed for the LAG interface?

Answer: The speed should be 20000Mbps, this is the combined speed of the 2 member ports, 10Gbps each.

2. Review the configuration of port 1/1/25.

```
show running-config interface 1/1/25
```

```
sw-edge2(config)# show running-config interface 1/1/25
interface 1/1/25
    no shutdown
    description sw-edge1
    lag 50
    exit
```

3. Review the configuration of port 1/1/26.

```
show running-config interface 1/1/26
```

```
sw-edge2(config)# show running-config interface 1/1/26
interface 1/1/26
    no shutdown
    description sw-edge1
    lag 50
    exit
```

4. Review the configuration of port LAG50.

```
show running-config interface lag 50
```

```
sw-edge2(config)# show running-config interface lag 50
interface lag 50
    no shutdown
    no routing
    vlan access 1
    lacp mode active
    exit
```

5. Review the aggregate interfaces.

```
show lacp aggregates
```

```
sw-edge2(config)# show lacp aggregates

Aggregate name    : lag50
Interfaces        : 1/1/25 1/1/26
Heartbeat rate    : Slow
Hash              : l3-src-dst
Aggregate mode    : Active
```

6. Review the LACP status of the interfaces, verify they are all UP.

```
show lacp interfaces
```

```
sw-edge2(config)# show lacp interfaces

State abbreviations :
A - Active        P - Passive      F - Aggregable I - Individual
S - Short-timeout L - Long-timeout N - InSync     O - OutofSync
C - Collecting    D - Distributing
X - State m/c expired              E - Default neighbor state

Actor details of all interfaces:
------------------------------------------------------------------------------
Intf       Aggr       Port  Port  State   System-ID         System Aggr Forwarding
           Name       Id    Pri                             Pri    Key  State
------------------------------------------------------------------------------
1/1/25     lag50      26    1     ALFNCD  8c:85:c1:49:20:c0 65534  50   up
1/1/26     lag50      27    1     ALFNCD  8c:85:c1:49:20:c0 65534  50   up


Partner details of all interfaces:
------------------------------------------------------------------------------
Intf       Aggr       Port  Port  State   System-ID         System Aggr
           Name       Id    Pri                             Pri    Key
------------------------------------------------------------------------------
1/1/25     lag50      27    1     ALFNCD  64:e8:81:dd:81:80 65534  50
1/1/26     lag50      26    1     ALFNCD  64:e8:81:dd:81:80 65534  50
```

Task 3: Configuring the LAG as a VLAN trunk

In this task you will configure the LAG between sw-edge1 and sw-edge2 and configure it as a VLAN trunk. This ensures that traffic of VLAN 11, for example the PC4 on sw-edge2, can reach its default gateway via this new LAG.

Objectives

- Configure the LAG as a VLAN trunk.

Steps

1. On **sw-edge1**, review the list of interfaces.

```
show interface brief
```

```
sw-edge1(config)# show interface  brief
--------------------------------------------------------------------------------
-----
Port      Native  Mode   Type          Enabled Status  Reason                Speed
Description
          VLAN                                                                 (Mb/s)
--------------------------------------------------------------------------------
-----
...
1/1/27    1       trunk  SFP+DAC1      yes     up                            10000
sw-agg1
1/1/28    1       access SFP+DAC1      no      down    Administratively down --
--
vlan1     --              --            yes     up                            --
--
lag50     1       access --            yes     up      --                    20000
--
```

Question: What is the current VLAN port type of the LAG?

Answer: The LAG is a new interface on the system. New interfaces are access ports in VLAN 1 by default.

2. Review the VLAN port status for LAG 50.

```
show vlan port lag50
```

```
sw-edge1(config)# show vlan port lag50

---------------------------------------------------------------------------
VLAN  Name                            Mode           Mapping
---------------------------------------------------------------------------
1     DEFAULT_VLAN_1                  access         port
```

3. Configure LAG 50 as VLAN trunk and allow all VLANs.

```
interface lag 50
```

```
 vlan trunk allow all
 exit
```

```
sw-edge1(config)# interface lag 50
sw-edge1(config-lag-if)#  vlan trunk allow all
sw-edge1(config-lag-if)#  exit
```

4. Review the VLAN port status for LAG 50.

```
show vlan port lag50
```

```
sw-edge1(config)# show vlan port lag50

--------------------------------------------------------------------------------
VLAN  Name                                Mode            Mapping
--------------------------------------------------------------------------------
1     DEFAULT_VLAN_1                      native-untagged port
11    v11-employee                        trunk           port
```

Question: Is VLAN 11 now allowed on the LAG port?

Answer: Yes.

5. On **sw-edge2**, apply the same configuration.

```
interface lag 50
 vlan trunk allow all
 exit
```

```
sw-edge2(config)# interface lag 50
sw-edge2(config-lag-if)#  vlan trunk allow all
sw-edge2(config-lag-if)#  exit
```

6. Verify that PC4 now has access to the network again.
7. Open a command prompt (cmd) and ping the default gateway IP address (10.1.11.1). This should be successful.

```
ping 10.1.11.1
```

8. Save the configuration on both edge switches.

```
write memory
```

Optional Task 4: Link-Aggregation Troubleshooting

The customer has had issues with link aggregations in the past due to incorrect cabling of the link aggregation member ports. You have been asked to explore how issues with cabling can be found on the switches. In this task you will review some of the common link aggregation issues.

Objectives

- Understand the LAG status
- Troubleshoot a LAG

Steps

LAG ports connected to 2 different LAGs on peer device.

This scenario will show how to detect that the ports of a LAG are connected to 2 different LAGs on the peer side. This could occur due to a patching error or a configuration error on the peer side. In this example, a configuration error will be introduced on sw-edge2.

1. On sw-edge2, create a new LAG with ID 51, set LACP active mode.

```
interface lag 51
 no shutdown
 lacp mode active
```

```
sw-edge2(config)# interface lag 51
sw-edge2(config-lag-if)#  no shutdown
sw-edge2(config-lag-if)#  lacp mode active
```

2. Configure the LAG 51 as VLAN trunk and allow all VLANs.

```
vlan trunk allow all
 exit
```

```
sw-edge2(config-lag-if)# vlan trunk allow all
sw-edge2(config-lag-if)#  exit
```

3. Configure port 1/1/26 as member of LAG 51. You must un-assign the port from LAG 50 before you can assign it to the LAG 51.

```
interface 1/1/26
 no lag 50
 lag 51
 exit
```

```
sw-edge2(config)# interface 1/1/26
sw-edge2(config-if)#  no lag 50
sw-edge2(config-if)#  lag 51
sw-edge2(config-if)#  exit
```

Troubleshoot

4. Review the LACP status on both sides. Pay attention to the aggregation key ID.

```
show lacp interfaces
```

```
sw-edge1(config)# show lacp interfaces

State abbreviations :
A - Active        P - Passive      F - Aggregable I - Individual
S - Short-timeout L - Long-timeout N - InSync     O - OutofSync
C - Collecting    D - Distributing
X - State m/c expired              E - Default neighbor state

Actor details of all interfaces:
------------------------------------------------------------------------------
Intf       Aggr       Port  Port  State   System-ID         System Aggr Forwarding
           Name       Id    Pri                             Pri    Key  State
------------------------------------------------------------------------------
1/1/25     lag50      26    1     ALFO    64:e8:81:dd:81:80 65534  50   lacp-block
1/1/26     lag50      27    1     ALFNCD  64:e8:81:dd:81:80 65534  50   up


Partner details of all interfaces:
------------------------------------------------------------------------------
Intf       Aggr       Port  Port  State   System-ID         System Aggr
           Name       Id    Pri                             Pri    Key
------------------------------------------------------------------------------
1/1/25     lag50      27    1     ALFN    8c:85:c1:49:20:c0 65534  51
1/1/26     lag50      26    1     ALFNCD  8c:85:c1:49:20:c0 65534  50
```

```
sw-edge2(config)# show lacp interfaces

State abbreviations :
A - Active        P - Passive      F - Aggregable I - Individual
S - Short-timeout L - Long-timeout N - InSync     O - OutofSync
C - Collecting    D - Distributing
X - State m/c expired              E - Default neighbor state

Actor details of all interfaces:
------------------------------------------------------------------------------
Intf       Aggr       Port  Port  State   System-ID         System Aggr Forwarding
           Name       Id    Pri                             Pri    Key  State
------------------------------------------------------------------------------
1/1/25     lag50      26    1     ALFNCD  8c:85:c1:49:20:c0 65534  50   up
1/1/26     lag51      27    1     ALFN    8c:85:c1:49:20:c0 65534  51   lacp-block


Partner details of all interfaces:
------------------------------------------------------------------------------
Intf       Aggr       Port  Port  State   System-ID         System Aggr
           Name       Id    Pri                             Pri    Key
------------------------------------------------------------------------------
1/1/25     lag50      27    1     ALFNCD  64:e8:81:dd:81:80 65534  50
1/1/26     lag51      26    1     ALFO    64:e8:81:dd:81:80 65534  50
```

Question: What do you observe?

Answer: Different LAG Aggregation Key IDs, where the result is an LACP block.

5. On sw-edge2, revert the configuration.

```
interface 1/1/26
 no lag 51
 lag 50
 exit
no interface lag 51
```

```
sw-edge2(config)# interface 1/1/26
sw-edge2(config-if)#  no lag 51
sw-edge2(config-if)#  lag 50
sw-edge2(config-if)#  exit
sw-edge2(config)# no interface lag 51
```

Optional Task 5: Verify Link-Aggregation Failover

Your customer has heard that the failover of a link aggregation is very fast in case of a link failure. You want to be able to demonstrate this to the customer.

In this task you will prepare a failover test, verify link aggregation failover operation, and examine the impact on existing traffic.

Objectives

- Verify the Link-Aggregation failover

Steps

1. On PC4, open a command prompt (cmd) and start a continuous ping to the default gateway.

```
ping 10.1.11.1 -t
```

2. On **sw-edge1**, shutdown the first port of the LAG.

```
interface 1/1/25
 shutdown
```

```
sw-edge1(config)# interface 1/1/25
sw-edge1(config-if)# shutdown
```

3. Verify the ping on PC4 continues.

> **NOTE:** A ping packet loss could be observed.

4. On sw-edge1, review the LACP interface state.

```
show lacp interfaces
```

```
sw-edge1(config-if)# show lacp interfaces

State abbreviations :
A - Active        P - Passive      F - Aggregable I - Individual
S - Short-timeout L - Long-timeout N - InSync     O - OutofSync
C - Collecting    D - Distributing
X - State m/c expired              E - Default neighbor state

Actor details of all interfaces:
------------------------------------------------------------------------------
Intf       Aggr       Port  Port  State   System-ID         System Aggr Forwarding
           Name       Id    Pri                             Pri    Key  State
------------------------------------------------------------------------------
1/1/25     lag50                                                      down
1/1/26     lag50      27    1     ALFNCD  64:e8:81:dd:81:80 65534  50   up


Partner details of all interfaces:
------------------------------------------------------------------------------
```

```
Intf       Aggr       Port  Port  State   System-ID         System Aggr
           Name       Id    Pri                             Pri    Key
-------------------------------------------------------------------------------
1/1/25     lag50
1/1/26     lag50      26    1     ALFNCD  8c:85:c1:49:20:c0 65534  50
```

5. Review the port speed in the interface lag brief output.

```
show interface lag brief
```

```
sw-edge1(config-if)# show interface lag brief
-------------------------------------------------------------------------------
-----
Port      Native  Mode   Type           Enabled Status  Reason                 Speed
Description
          VLAN                                                                 (Mb/s)
-------------------------------------------------------------------------------
-----
lag50     1       trunk  --             yes     up      --                     10000
--
```

Question: What happened with the reported interface LAG speed?

Answer: Previously, with both links active, the reported link speed was 20Gbps. With one link down, the reported link speed is 10Gbps. The LAG speed is adjusted to the actual number of active interfaces.

6. Enable port 1/1/25 again.

```
interface 1/1/25
 no shutdown
```

```
sw-edge1(config-if)# interface 1/1/25
sw-edge1(config-if)#  no shutdown
```

7. Now simulate a failure of the second link by shutting down port 1/1/26.

```
interface 1/1/26
 shutdown
```

```
sw-edge1(config-if)# interface 1/1/26
sw-edge1(config-if)#  shutdown
```

8. On PC4, verify that the ping continues; this demonstrates the link aggregation failover test.
9. On sw-edge1, enable port 1/1/26 again.

```
interface 1/1/26
 no shutdown
 exit
```

```
sw-edge1(config-if)# interface 1/1/26
sw-edge1(config-if)#  no shutdown
sw-edge1(config-if)#  exit
```

10. Verify that both interfaces are in LACP UP state.

```
show lacp interfaces
```

```
sw-edge1(config)# show lacp interfaces

State abbreviations :
A - Active        P - Passive      F - Aggregable I - Individual
S - Short-timeout L - Long-timeout N - InSync     O - OutofSync
C - Collecting    D - Distributing
X - State m/c expired              E - Default neighbor state

Actor details of all interfaces:
------------------------------------------------------------------------------
Intf       Aggr       Port  Port  State   System-ID         System Aggr Forwarding
           Name       Id    Pri                             Pri    Key  State
------------------------------------------------------------------------------
1/1/25     lag50      26    1     ALFNCD  64:e8:81:dd:81:80 65534  50   up
1/1/26     lag50      27    1     ALFNCD  64:e8:81:dd:81:80 65534  50   up


Partner details of all interfaces:
------------------------------------------------------------------------------
Intf       Aggr       Port  Port  State   System-ID         System Aggr
           Name       Id    Pri                             Pri    Key
------------------------------------------------------------------------------
1/1/25     lag50      27    1     ALFNCD  8c:85:c1:49:20:c0 65534  50
1/1/26     lag50      26    1     ALFNCD  8c:85:c1:49:20:c0 65534  50
```

Optional Task 6: Static LAG Configuration

The customer is aware that an LACP LAG is the recommended best practice for link aggregation. However, in some rare cases, they expect a device or server that may not support LACP.

After consulting with your colleague, you learn that a static LAG can be configured for those devices. In this task you will practice the configuration of a static LAG. This is a LAG without protocol, so LACP is not used.

To practice this configuration, the existing LACP LAG between the switches will be converted to a static LAG on sw-edge1.

Objectives

- Configure a static LAG

Steps

1. On **sw-edge1**, remove the LACP mode active from the LAG 50.

```
interface lag 50
 no lacp mode active
 exit
```

```
sw-edge1(config)# interface lag 50
sw-edge1(config-lag-if)#  no lacp mode active
sw-edge1(config-lag-if)#  exit
```

2. Review the LACP interfaces.

```
show lacp interfaces
```

```
sw-edge1(config)# show lacp interfaces

State abbreviations :
A - Active        P - Passive      F - Aggregable I - Individual
S - Short-timeout L - Long-timeout N - InSync     O - OutofSync
C - Collecting    D - Distributing
X - State m/c expired              E - Default neighbor state

Actor details of all interfaces:
------------------------------------------------------------------------------
Intf       Aggr       Port  Port  State   System-ID         System Aggr Forwarding
           Name       Id    Pri                             Pri    Key  State
------------------------------------------------------------------------------
1/1/25     lag50                                                      up
1/1/26     lag50                                                      up


Partner details of all interfaces:
------------------------------------------------------------------------------
Intf       Aggr       Port  Port  State   System-ID         System Aggr
           Name       Id    Pri                             Pri    Key
------------------------------------------------------------------------------
```

```
1/1/25      lag50
1/1/26      lag50
```

NOTE: Even though a static LAG has now been configured, the command used to verify the LAG is **show lacp interfaces**.

Question: What is the peer ID reported for a static LAG?

Answer: For a static LAG, no LACP peer id will be reported.

Question: What would happen when a static LAG is connected to an LACP LAG?

_____

Answer: This depends on the platforms that are used. For the AOS-CX switch platform, the LACP side will set all local interfaces to 'blocked'. This means no traffic would traverse the LAG. Other platforms may enable one interface of the LAG when no LACP partner is observed, or even enable all interfaces. You should become familiar with this behavior of the various platforms in the network.

Question: What would happen when a non-configured device is connected to an LACP LAG?

_____

Answer: The LACP side does not receive any LACP frames, so the result will be the same as the previous question (depends on the platform). During zero-touch-provisioning, a newly deployed access switch may not have the correct LAG configuration yet. This would result in the aggregation switch LACP to block all traffic. To handle with this scenario, AOS-CX aggregation switches with VSX support a feature known as *LACP fallback*, this enables the ports when no LACP frames are received, as opposed to the 'blocking' default behavior and preventing a device from becoming provisioned (getting its configuration).

3. Revert the LAG mode to LACP active.

```
interface lag 50
 lacp mode active
 exit
```

```
sw-edge1(config)# interface lag 50
sw-edge1(config-lag-if)#  lacp mode active
sw-edge1(config-lag-if)#  exit
```

4. Verify that the LACP status is *up*.

```
show lacp interfaces
```

```
sw-edge1(config)# show lacp interfaces

State abbreviations :
A - Active        P - Passive      F - Aggregable I - Individual
S - Short-timeout L - Long-timeout N - InSync     O - OutofSync
C - Collecting    D - Distributing
X - State m/c expired              E - Default neighbor state
```

```
Actor details of all interfaces:
------------------------------------------------------------------------------
Intf       Aggr       Port  Port  State   System-ID         System Aggr Forwarding
           Name       Id    Pri                             Pri    Key  State
------------------------------------------------------------------------------
1/1/25     lag50      26    1     ALFNCD  64:e8:81:dd:81:80 65534  50   up
1/1/26     lag50      27    1     ALFNCD  64:e8:81:dd:81:80 65534  50   up

Partner details of all interfaces:
------------------------------------------------------------------------------
Intf       Aggr       Port  Port  State   System-ID         System Aggr
           Name       Id    Pri                             Pri    Key
------------------------------------------------------------------------------
1/1/25     lag50      27    1     ALFNCD  8c:85:c1:49:20:c0 65534  50
1/1/26     lag50      26    1     ALFNCD  8c:85:c1:49:20:c0 65534  50
```

You have completed this Lab!

Lab 03.01 – Basic IP Configuration
Overview

In this lab, the senior network administrator has requested that you add a new VLAN to the network and prepare the IP routing configuration for the new VLAN on the aggregation switch.

Objectives

After completing this lab, you will be able to:

- Add a new VLAN IP interface, also known as SVI, to the aggregation switch.
- Configure DHCP Relay, also known as ip helper, on the aggregation switch.
- Verify routing between the new VLAN and an existing VLAN on the network.

Task 1: Prepare the Edge Switch Uplink to the Aggregation Switch

To ensure that each edge switch has a direct, single uplink to the aggregation switch, you will disable the LAG between the sw-edge1 and sw-edge2. Sw-edge1 has port 1/1/27 already connected to sw-agg1. This was done in the previous lab. Sw-edge2 needs to be configured.

Objectives

- Disable the LAG between the edge switches.

Steps

1. On sw-edge1, disable the LAG 50 (that connects to sw-edge2).

```
interface lag 50
 shutdown
 exit
```

```
sw-edge1(config)# interface lag 50
sw-edge1(config-lag-if)#  shutdown
sw-edge1(config-lag-if)# exit
```

2. Verify that the interface LAG 50 is disabled.

```
show interface lag brief
```

```
sw-edge1(config)# show interface lag brief
--------------------------------------------------------------------------------
-------------------------
Port      Native  Mode   Type           Enabled Status  Reason                 Speed
Description
          VLAN                                                                  (Mb/s)
--------------------------------------------------------------------------------
-------------------------
lag50     1       trunk  --             no      down    --                     auto
--
```

3. Use the LLDP neighbor list to verify that port 1/1/27 is connected to the sw-agg1.

```
show lldp neighbor-info
```

```
sw-edge1(config)# show lldp neighbor-info

LLDP Neighbor Information
=========================

Total Neighbor Entries          : 8
Total Neighbor Entries Deleted  : 0
Total Neighbor Entries Dropped  : 0
Total Neighbor Entries Aged-Out : 0

LOCAL-PORT  CHASSIS-ID         PORT-ID                    PORT-DESC
TTL      SYS-NAME
```

```
--------------------------------------------------------------------------------
-----
...
1/1/25      8c:85:c1:49:20:c0  1/1/26                      sw-edge1
120      sw-edge2
1/1/26      8c:85:c1:49:20:c0  1/1/25                      sw-edge1
120      sw-edge2
1/1/27      b8:d4:e7:da:80:00  1/1/1                       1/1/1
120      sw-agg1
...
```

Question: Do you still observe sw-edge2 in the LLDP neighbor list?

Answer: Yes, even though the connection to sw-edge2 is down, the LLDP neighbor entries have their own aging timers (time to live), therefore, you may still notice the neighbor for some time. This is typically useful in troubleshooting, since you can still detect the connected device for a short period of time, even though the link is down.

Question: What is the TTL shown for the sw-edge2?

Answer: The TTL is 120 (seconds), this means that the neighbor will age out 120 seconds after the last received LLDP PDU.

4. On sw-edge2, enable port 1/1/27. This port connects to sw-agg1.

```
interface 1/1/27
 no shutdown
 exit
```

```
sw-edge2(config)# interface 1/1/27
sw-edge2(config-if)#  no shutdown
sw-edge2(config-if)#  exit
```

5. Verify sw-agg1 is now listed as an LLDP neighbor.

```
show lldp neighbor-info
```

```
sw-edge2(config)# show lldp neighbor-info

LLDP Neighbor Information
=========================

Total Neighbor Entries          : 5
Total Neighbor Entries Deleted  : 2
Total Neighbor Entries Dropped  : 0
Total Neighbor Entries Aged-Out : 2

LOCAL-PORT  CHASSIS-ID         PORT-ID                    PORT-DESC
TTL      SYS-NAME
--------------------------------------------------------------------------------
-----...
1/1/27      b8:d4:e7:da:80:00  1/1/2                      1/1/2
120      sw-agg1
...
```

Task 2: Configure a new Layer2 VLAN in the Edge Infrastructure

In this task you will prepare a new VLAN (VLAN 12) in the environment. Since the lab topology expects a loop-free setup for now, you will disable the LAG between sw-edge1 and sw-edge2. On sw-edge2 you will then enable the uplink to sw-agg1.

Objectives

- Practice the VLAN configuration

Practice

Define VLAN 12 on sw-edge1 and sw-edge2.
Allow all VLANs on the uplink port 1/1/27 on both sw-edge1 and sw-edge2.
On sw-edge2, configure port 1/1/4 as access port in VLAN 12.

Solution

1. Sw-edge1's configuration:

```
vlan 12

interface 1/1/27
 vlan trunk allow all
 no shutdown
 exit
```

2. Sw-edge2's configuration:

```
vlan 12

interface 1/1/27
 vlan trunk allow all
 no shutdown
 exit

interface 1/1/4
 vlan access 12
 exit
```

Task 3: Configure the Aggregation Switch with a new VLAN and SVI

In this task you will create the new VLAN on the aggregation switch. In a real deployment, the VLAN would be created on both aggregation switches. However, to keep the lab setup simple at this stage, you will only configure this on the first aggregation switch: sw-agg1.

In this task, you will configure a new IP interface on VLAN 12. A layer 3 interface configured on a VLAN is known as a switch virtual interface (SVI). This SVI will be the default gateway for clients in VLAN 12.
You will also configure a DHCP Relay on the SVI, to ensure that clients in VLAN 12 can obatin a DHCP-based IP address from the centralized DHCP server.

Objectives

- Configure an SVI on the aggregation switch

Steps

Layer2 VLAN 12

1. On **sw-agg1**, create VLAN 12.

```
vlan 12
 exit
```

```
sw-agg1(config)# vlan 12
sw-agg1(config-vlan-12)#  exit
```

2. Review the VLAN 12 port membership.

```
show vlan 12
```

```
sw-agg1(config)# show vlan 12

-------------------------------------------------------------------------------
-----
VLAN  Name                 Status  Reason                 Type        Interfaces
-------------------------------------------------------------------------------
-----
12    VLAN12               up      ok                     static      1/1/1-
1/1/2,lag256
```

Question: Is port 1/1/1 member of VLAN12? Which ports are members of VLAN12?

_____

Answer: These ports are members of VLAN12: 1/1/1,1/1/2, LAG256.

3. Review the port configuration of port 1/1/1 and 1/1/2.

```
show running-config interface 1/1/1
show running-config interface 1/1/2
```

```
sw-agg1(config)# show running-config interface 1/1/1
interface 1/1/1
    no shutdown
```

```
    mtu 9198
    no routing
    vlan trunk native 1
    vlan trunk allowed all
    exit
```

```
sw-agg1(config)# show running-config interface 1/1/2
interface 1/1/2
    no shutdown
    mtu 9198
    no routing
    vlan trunk native 1
    vlan trunk allowed all
    exit
```

Question: What do you observe?

_____

Answer: Ports 1/1/1 and 1/1/2 have been configured as VLAN trunk and allow all VLANs. This automatically includes the newly created VLAN12.

Now you will investigate what system connects to LAG256.

4. Review the LACP interface status to determine the member ports.

```
show lacp interfaces
```

```
sw-agg1(config)# show lacp interfaces

State abbreviations :
A - Active        P - Passive      F - Aggregable I - Individual
S - Short-timeout L - Long-timeout N - InSync     O - OutofSync
C - Collecting    D - Distributing
X - State m/c expired              E - Default neighbor state

Actor details of all interfaces:
------------------------------------------------------------------------------
Intf       Aggr       Port  Port  State   System-ID         System Aggr Forwarding
           Name       Id    Pri                             Pri    Key  State
------------------------------------------------------------------------------
1/1/46     lag256     47    1     ALFNCD  b8:d4:e7:da:80:00 65534  256  up
1/1/47     lag256     48    1     ALFNCD  b8:d4:e7:da:80:00 65534  256  up


Partner details of all interfaces:
------------------------------------------------------------------------------
Intf       Aggr       Port  Port  State   System-ID         System Aggr
           Name       Id    Pri                             Pri    Key
------------------------------------------------------------------------------
1/1/46     lag256     47    1     ALFNCD  b8:d4:e7:d9:e5:00 65534  256
1/1/47     lag256     48    1     ALFNCD  b8:d4:e7:d9:e5:00 65534  256
```

Question: What are the member ports of the LAG 256?

Answer: Ports 1/1/46 and 1/1/47 are member ports of the LAG256.

5. Review the LLDP neighbor output to see what device connects to these ports.

```
show lldp neighbor-info
```

```
sw-agg1(config)# show lldp neighbor-info

LLDP Neighbor Information
=========================

Total Neighbor Entries          : 6
Total Neighbor Entries Deleted  : 2
Total Neighbor Entries Dropped  : 0
Total Neighbor Entries Aged-Out : 2

LOCAL-PORT  CHASSIS-ID         PORT-ID                    PORT-DESC
TTL      SYS-NAME
-------------------------------------------------------------------------------
-----...
1/1/46      b8:d4:e7:d9:e5:00  1/1/46                     1/1/46
120      sw-agg2
1/1/47      b8:d4:e7:d9:e5:00  1/1/47                     1/1/47
120      sw-agg2
...
```

Question: To what device is LAG 256 connected?

Answer: LAG 256 consists of ports 1/1/46 and 1/1/47 and connects to the sw-agg2.

Layer3 SVI 12

You will now create a SVI for VLAN 12. An SVI is a layer 3 (IP) interface for a VLAN. When a switch doesn't have an SVI for a VLAN, it doesn't have an IP interface in that VLAN, so it cannot directly interact with IP devices in that VLAN.

An access switch typically doesn't need an SVI in the endpoint VLANs. Aggregation switches typically do need an SVI, since they will provide IP default gateway services for the VLAN.

6. On sw-agg1, review the current list of L3 interfaces. Make sure to check the list to the end.

```
show ip interface brief
```

Question: Is there an SVI for VLAN 11 and for VLAN 12?

Answer: In the existing network, only an SVI for VLAN 11 exists. Even though the layer 2 VLAN 12 was created on the switch, it doesn't have an IP interface in this VLAN yet.

```
sw-agg1# show ip interface brief
Interface         IP Address              Interface Status
                                            link/admin
...
vlan1             10.1.1.1/24               up/up
vlan2             10.1.2.2/24               up/up
```

```
vlan3              10.1.3.1/24                up/up
vlan11             10.1.11.1/24               up/up
```

7. Create a new SVI for VLAN 12.

```
interface vlan12
```

```
sw-agg1(config)# interface vlan 12
```

8. Review the list of interfaces. This will now include SVI12, shown as ‘vlan12’

```
show ip interface brief
```

```
sw-agg1# show ip interface brief
Interface         IP Address             Interface Status
                                           link/admin
...
vlan1             10.1.1.1/24              up/up
vlan2             10.1.2.2/24              up/up
vlan3             10.1.3.1/24              up/up
vlan11            10.1.11.1/24             up/up
vlan12            No Address               up/up
```

9. Configure IP address 10.1.12.1/24 on this SVI.

```
 ip address 10.1.12.1/24
```

```
sw-agg1(config-if-vlan)#  ip address 10.1.12.1/24
```

10. Configure the DHCP Relay (also known as ‘ip helper’) on this SVI. The IP address used points to the centralized DHCP server.

```
 ip helper-address 10.254.1.21
 exit
```

```
sw-agg1(config-if-vlan)#  ip helper-address 10.254.1.21
sw-agg1(config-if-vlan)#  exit
```

11. Review the SVI 12 interface status. Take note of the interface MAC address.

```
show ip interface vlan12
```

Example output, your output may be different.

```
sw-agg1(config)# show ip interface vlan12

Interface vlan12 is up
 Admin state is up
 Hardware: Ethernet, MAC Address: b8:d4:e7:da:80:00
 IP MTU 1500
 IP Directed Broadcast is Disabled
 IP Neighbor flood is Disabled
 IPv4 address 10.1.12.1/24
 active-gateway L3 source mac b8:d4:e7:da:80:00
```

SVI 12 MAC: _____

12. On sw-edge1, review the MAC-address table of VLAN 12.

```
show mac-address-table vlan 12
```

```
sw-edge1(config)# show mac-address-table vlan 12
MAC age-time            : 300 seconds
Number of MAC addresses : 2

MAC Address          VLAN     Type                      Port
-------------------------------------------------------------
00:50:56:b1:e3:6f    12       dynamic                   1/1/27
b8:d4:e7:da:80:00    12       dynamic                   1/1/27
```

Question: Do you see the MAC address of the aggregation switch in the MAC table?

_____

Answer: Yes, the aggregation switch SVI is sending ARP broadcast packets. These contain the aggregation switch SVI MAC address as the source MAC, and this is how the edge switches learn the aggregation SVI MAC address.

SVI MAC on different VLANs

13. On sw-agg1, lookup the MAC address of the SVI11.

```
show ip interface vlan11
```

```
sw-agg1(config)# show ip interface vlan11

Interface vlan11 is up
 Admin state is up
 Hardware: Ethernet, MAC Address: b8:d4:e7:da:80:00
 IP MTU 1500
 IP Directed Broadcast is Disabled
 IP Neighbor flood is Disabled
 IPv4 address 10.1.11.1/24
 active-gateway L3 source mac b8:d4:e7:da:80:00
```

Question: Is this MAC address different from the MAC address that is used on SVI12?

Answer: No, the same MAC address is used on different SVIs.

Question: Is it a problem to have the same MAC address on multiple places in the network?

Answer: No, MAC addresses must only be unique within a VLAN. The fact that a MAC address is typically globally unique also ensures that it will be unique within a VLAN. It is typically no problem to re-use the same MAC address in different VLANs.

14. On sw-edge1, review the MAC address table for VLAN 11 and VLAN 12. You should observe the aggregation switch MAC address in both VLANs, with the same MAC address.

```
show mac-address-table vlan 11
show mac-address-table vlan 12
```

```
sw-edge1(config)# show mac-address-table vlan 11
MAC age-time             : 300 seconds
Number of MAC addresses : 2

MAC Address          VLAN     Type                      Port
-------------------------------------------------------------
00:50:56:b1:cd:26    11       dynamic                   1/1/1
b8:d4:e7:da:80:00    11       dynamic                   1/1/27
```

```
sw-edge1(config)# show mac-address-table vlan 12
MAC age-time             : 300 seconds
Number of MAC addresses : 2

MAC Address          VLAN     Type                      Port
-------------------------------------------------------------
00:50:56:b1:e3:6f    12       dynamic                   1/1/27
b8:d4:e7:da:80:00    12       dynamic                   1/1/27
```

Verify VLAN 12 Operation

With VLAN 12 configured on the edge and aggregation switch, and the SVI 12 configured with an IP address and IP helper active on the sw-agg1, your VLAN 12 should now be ready for use.
Test the operation with your PC4, it should get an IP address in VLAN 12.

15. Open a connection to PC4 to check the current IP address.
16. Right-click the **network** icon in the status bar.
17. Click **Open Network & Internet Settings**.

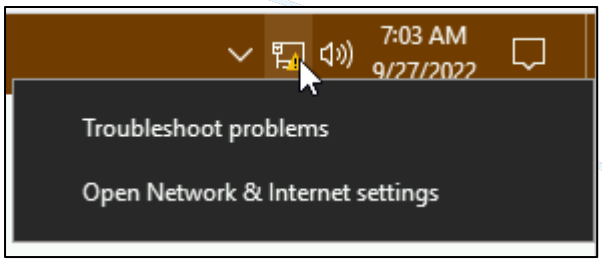

18. Under Change your network settings, click **Change adapter options**.

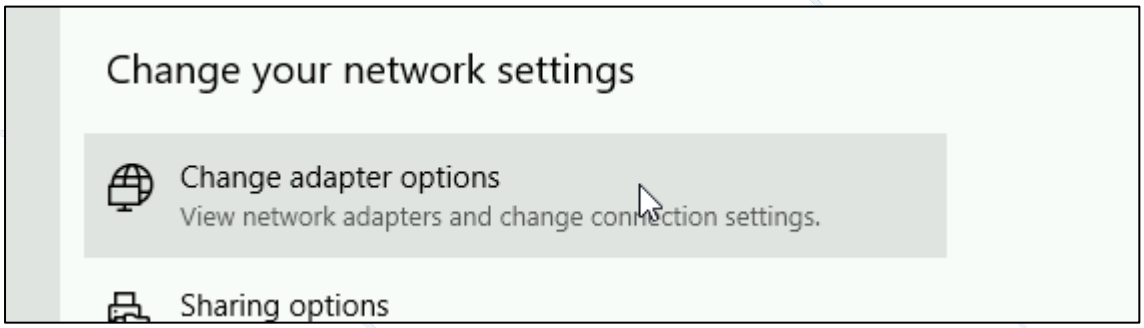

19. Right-click the **Lab NIC** (this is the NIC connected to port 4 on sw-edge2). Click on **Status**.

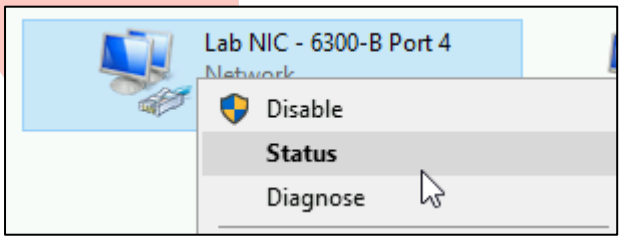

20. Click **Details**.

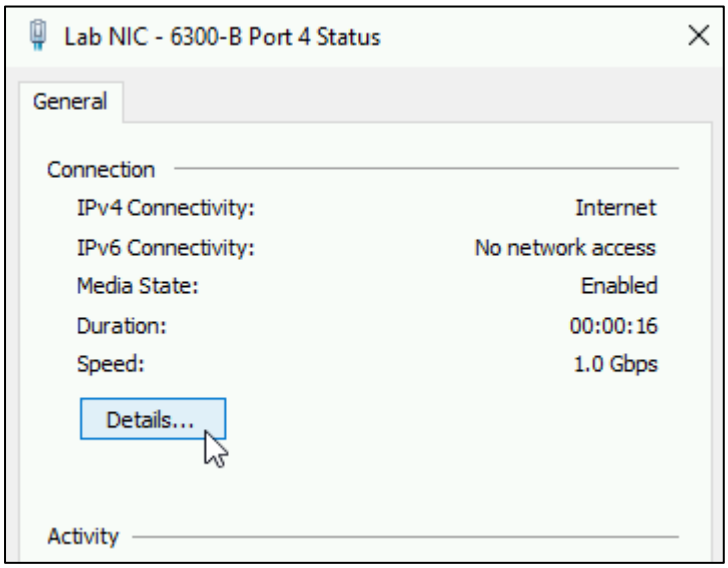

21. Take note of the MAC address and the IP address of the NIC.

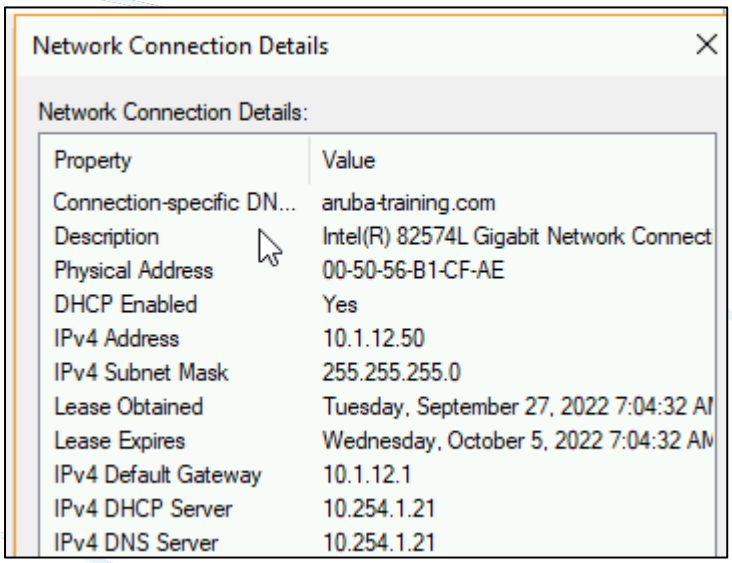

NOTE: You may need to *disable* and *enable* the **Lab NIC** interface to force a DHCP renewal. Make sure that you **don't disable the 'Do Not Touch'** interface.

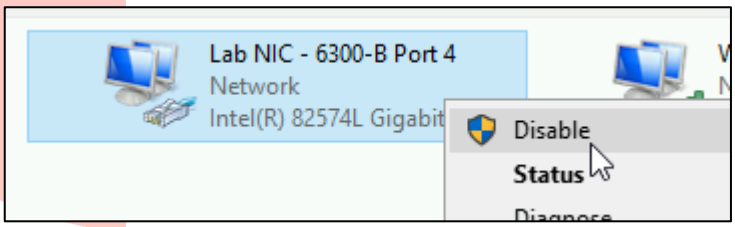

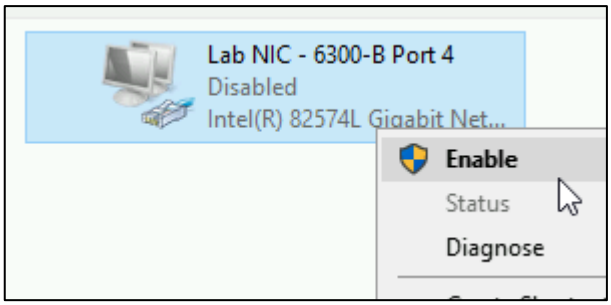

22. On sw-edge2, review the MAC address table of VLAN 12. You should see the PC4 MAC address on port 1/1/4 and the sw-agg1 SVI12 MAC on port 1/1/27.

```
show mac-address-table vlan 12
```

```
sw-edge2(config)# show mac-address-table vlan 12
MAC age-time            : 300 seconds
Number of MAC addresses : 2

MAC Address          VLAN     Type                        Port
-------------------------------------------------------------
00:50:56:b1:e3:6f    12       dynamic                     1/1/4
b8:d4:e7:da:80:00    12       dynamic                     1/1/27
```

23. On sw-agg1, review the MAC address table. You should see the PC4 MAC address on the port 1/1/2.

```
show mac-address-table vlan 12
```

```
sw-agg1(config)# show mac-address-table vlan 12
MAC age-time            : 300 seconds
Number of MAC addresses : 1

MAC Address          VLAN     Type                        Port
-------------------------------------------------------------
00:50:56:b1:e3:6f    12       dynamic                     1/1/2
```

24. On sw-agg1, review the ARP table. You should see the PC4 IP address to MAC entry has been learned on port 1/1/2.

```
show arp
```

```
sw-agg1(config)# show arp

IPv4 Address     MAC                Port         Physical Port               State
--------------------------------------------------------------------------------
-
10.1.11.50       00:50:56:b1:cd:26  vlan11       1/1/1                       reachable
10.1.12.50       00:50:56:b1:e3:6f  vlan12       1/1/2                       reachable
10.254.101.254   08:00:09:e8:fc:ea  1/1/8        1/1/8                       reachable

Total Number Of ARP Entries Listed: 3.
--------------------------------------------------------------------------------
-
```

25. On sw-agg1, review the DHCP Relay statistics, you should see a number higher than 0 for Valid Requests and Responses.

```
show dhcp-relay
```

```
sw-agg1(config)# show dhcp-relay

 DHCP Relay Agent                 : enabled
 DHCP Request Hop Count Increment : enabled
 L2VPN Clients                    : enabled
 Option 82                        : disabled
 Source-Interface                 : disabled
 Response Validation              : disabled
 Option 82 Handle Policy          : replace
 Remote ID                        : mac

 DHCP Relay Statistics:

  Valid Requests Dropped Requests Valid Responses Dropped Responses
  -------------- ---------------- --------------- -----------------
  3              0                3               0

 DHCP Relay Option 82 Statistics:

  Valid Requests Dropped Requests Valid Responses Dropped Responses
  -------------- ---------------- --------------- -----------------
  0              0                0               0
```

26. **Optional step**: on the PC4, you can disable/enable the Lab NIC to trigger a new DHCP request. This will result in an increase in the DHCP Relay statistics.

```
sw-agg1(config)# show dhcp-relay

 DHCP Relay Agent                 : enabled
 DHCP Request Hop Count Increment : enabled
 L2VPN Clients                    : enabled
 Option 82                        : disabled
 Source-Interface                 : disabled
 Response Validation              : disabled
```

```
Option 82 Handle Policy          : replace
Remote ID                        : mac

DHCP Relay Statistics:

 Valid Requests Dropped Requests Valid Responses Dropped Responses
 -------------- ---------------- --------------- -----------------
 5              0                5               0

DHCP Relay Option 82 Statistics:

 Valid Requests Dropped Requests Valid Responses Dropped Responses
 -------------- ---------------- --------------- -----------------
 0              0                0               0
```

27. Save the configuration on all switches.

```
write memory
```

You have completed this Lab!

Lab 03.02 – Edge Switch Management IP

Overview

At this stage in the lab, the PC1 and PC4 are connected to the network on VLAN11 and VLAN12. However, the access switches themselves are not reachable on the network yet. The access switches have an Out of Band Management (OOBM) IP address. This is convenient in the training lab environment. However, this is not considered a standard configuration for access switches in a campus deployment.

In a typical campus deployment, the access switches will be configured through an in-band management VLAN. In this lab the access switches will be configured with an in-band management IP and you will ensure that they can access the Internet. This will be required in later labs to provide cloud management functions using Aruba Central.

Objectives

After completing this lab, you will be able to:

- Configure management VLAN 3 and appropriate IP addresses on edge switches.
- Verify management connectivity on edge switches.

Task 1: Configure In-band Management IP VLAN 3

In this task the access switches will be configured with an SVI on VLAN 3, a default gateway, and DNS server.

Objectives

- Configure an access switch management IP address.
- Configure DNS name lookup on the access switches.
- Configure time services using NTP.

Steps

Review the current management IP address

1. On sw-edge1, attempt to ping 8.8.8.8.

```
ping 8.8.8.8
```

```
sw-edge1(config)# ping 8.8.8.8
ping4: connect: Network is unreachable
```

2. Review the current IP interfaces.

```
show ip interface brief
```

```
sw-edge1(config)# show ip interface brief
Interface         IP Address             Interface Status
                                           link/admin
vlan1           No Address                 up/up
```

Question: Did the ping work? Why?

Answer: No, sw-edge1 currently does not have an in-band IP address.

3. Make sure you can distinguish in-band (**show ip interface brief**) from the OOBM management interface (**show interface mgmt**).

```
show interface mgmt
```

```
sw-edge1(config)# show interface mgmt
  Address Mode                  : static
  Admin State                   : up
  Link State                    : up
  Mac Address                   : 64:e8:81:dd:81:81
  IPv4 address/subnet-mask      : 10.251.1.4/24
  Default gateway IPv4          :
  IPv6 address/prefix           :
  IPv6 link local address/prefix: fe80::66e8:81ff:fedd:8181/64
  Default gateway IPv6          :
  Primary Nameserver            :
  Secondary Nameserver          :
  Tertiary Nameserver           :
```

Configure VLAN 3 In-band IP

4. On sw-edge1, create VLAN3.

```
vlan 3
 exit
```

```
sw-edge1(config)# vlan 3
sw-edge1(config-vlan-3)# exit
```

5. Create SVI3.

```
interface vlan 3
```

```
sw-edge1(config)# interface vlan 3
```

6. Assign IP address 10.1.3.4/24 and return to the global configuration context.

```
 ip address 10.1.3.4/24
 exit
```

```
sw-edge1(config-if-vlan)#  ip address 10.1.3.4/24
sw-edge1(config-if-vlan)#  exit
```

7. Review the IP interface list.

```
show ip interface brief
```

```
sw-edge1(config)# show ip interface brief
Interface         IP Address             Interface Status
                                           link/admin
vlan1             No Address               up/up

vlan3             10.1.3.4/24              up/up
```

VLAN1 - SVI1 is enabled for DHCP by default, this is useful for ZTP deployments, but in case you are using a different VLAN ID for management, it is safer to disable the DHCP client and shutdown SVI1.

8. Disable the DHCP client on SVI 1.

```
interface vlan 1
 no ip dhcp
 shutdown
 exit
```

```
sw-edge1(config)# interface vlan 1
sw-edge1(config-if-vlan)#  no ip dhcp
sw-edge1(config-if-vlan)#  shutdown
sw-edge1(config-if-vlan)#  exit
```

9. Attempt to ping sw-agg1 (IP 10.1.3.1), this should now be successful.

```
ping 10.1.3.1
```

```
sw-edge1(config)# ping 10.1.3.1
PING 10.1.3.1 (10.1.3.1) 100(128) bytes of data.
108 bytes from 10.1.3.1: icmp_seq=1 ttl=64 time=20.3 ms
108 bytes from 10.1.3.1: icmp_seq=2 ttl=64 time=0.173 ms
108 bytes from 10.1.3.1: icmp_seq=3 ttl=64 time=0.173 ms
108 bytes from 10.1.3.1: icmp_seq=4 ttl=64 time=0.224 ms
108 bytes from 10.1.3.1: icmp_seq=5 ttl=64 time=0.174 ms

--- 10.1.3.1 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4079ms
rtt min/avg/max/mdev = 0.173/4.202/20.266/8.032 ms
```

10. Attempt to ping a host on the internet (IP 8.8.8.8).

```
ping 8.8.8.8
```

```
sw-edge1(config)# ping 8.8.8.8
ping4: connect: Network is unreachable
```

Question: Did this work? Why?

Answer: The ping did not work. The edge switch does not have a default route configured.

11. Review the IP routing table.

```
show ip route
```

```
sw-edge1(config)# show ip route

Displaying ipv4 routes selected for forwarding

Origin Codes: C - connected, S - static, L - local
              R - RIP, B - BGP, O - OSPF
Type Codes:   E - External BGP, I - Internal BGP, V - VPN, EV - EVPN
              IA - OSPF internal area, E1 - OSPF external type 1
              E2 - OSPF external type 2

VRF: default

Prefix              Nexthop       Interface     VRF(egress)       Origin/   Distance/
Age
                                                                     Type      Metric
-------------------------------------------------------------------------------------
-----
10.1.3.0/24         -             vlan3         -                 C         [0/0]
-
10.1.3.4/32         -             vlan3         -                 L         [0/0]
-

Total Route Count : 2
```

12. Add a default route to 10.1.3.1.

```
ip route 0.0.0.0/0 10.1.3.1
```

```
sw-edge1(config)# ip route 0.0.0.0/0 10.1.3.1
```

13. Verify the route is active in the IP routing table.

```
show ip route
```

```
sw-edge1(config)# show ip route

Displaying ipv4 routes selected for forwarding

Origin Codes: C - connected, S - static, L - local
              R - RIP, B - BGP, O - OSPF
Type Codes:   E - External BGP, I - Internal BGP, V - VPN, EV - EVPN
              IA - OSPF internal area, E1 - OSPF external type 1
              E2 - OSPF external type 2

VRF: default

Prefix        Nexthop      Interface     VRF(egress)       Origin/    Distance/    Age
                                                           Type       Metric
-------------------------------------------------------------------------------------
-----
0.0.0.0/0      10.1.3.1    vlan3         -                 S          [1/0]
00h:00m:14s
10.1.3.0/24    -           vlan3         -                 C          [0/0]        -
10.1.3.4/32    -           vlan3         -                 L          [0/0]        -

Total Route Count : 3
```

14. Attempt to ping a host on the internet (IP 8.8.8.8). This should now be successful.

```
ping 8.8.8.8
```

```
sw-edge1(config)# ping 8.8.8.8
PING 8.8.8.8 (8.8.8.8) 100(128) bytes of data.
76 bytes from 8.8.8.8: icmp_seq=1 ttl=114 (truncated)
76 bytes from 8.8.8.8: icmp_seq=2 ttl=114 (truncated)
76 bytes from 8.8.8.8: icmp_seq=3 ttl=114 (truncated)
76 bytes from 8.8.8.8: icmp_seq=4 ttl=114 (truncated)
76 bytes from 8.8.8.8: icmp_seq=5 ttl=114 (truncated)

--- 8.8.8.8 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4005ms
rtt min/avg/max/mdev = 8.656/9.189/9.803/0.411 ms
```

Name resolution

You will now test and configure the DNS name lookup function in the network.

15. Attempt to ping the Aruba Path Quality Monitoring (PQM) host.

```
ping pqm.arubanetworks.com
```

```
sw-edge1(config)# ping pqm.arubanetworks.com
ping4: pqm.arubanetworks.com: Temporary failure in name resolution
```

NOTE: The host **pqm.arubanetworks.com** is cloud-hosted by Aruba Networks for branch site connectivity and path quality checks.

Question: Did the ping succeed? Why?

Answer: No, the edge switch does not have any name resolution server (Domain Name Service (DNS)) configured. While it can reach the internet, it currently cannot perform name lookups.

16. Configure a DNS host 10.254.1.21.

```
ip dns server-address 10.254.1.21
```

```
sw-edge1(config)# ip dns server-address 10.254.1.21
```

17. Attempt to ping **pqm.arubanetworks.com** again, which should now be successful.

```
ping pqm.arubanetworks.com
```

```
sw-edge1(config)# ping pqm.arubanetworks.com
PING pqm.arubanetworks.com (54.81.169.15) 100(128) bytes of data.
108 bytes from 54.81.169.15: icmp_seq=1 ttl=34 time=30.9 ms
108 bytes from 54.81.169.15: icmp_seq=2 ttl=34 time=30.9 ms
108 bytes from 54.81.169.15: icmp_seq=3 ttl=34 time=30.7 ms
108 bytes from 54.81.169.15: icmp_seq=4 ttl=34 time=30.7 ms
108 bytes from 54.81.169.15: icmp_seq=5 ttl=34 time=43.7 ms

--- pqm.arubanetworks.com ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4006ms
rtt min/avg/max/mdev = 30.659/33.379/43.708/5.165 ms
```

NOTE: Since it is cloud-hosted, *pqm.arubanetworks.com* runs on many different systems and therefore the IP address may be different for each lookup.

Time Services

The current release of the AOS-CX NTP client configuration includes a default NTP server. This can be overruled in the configuration in case a specific server is required or access to the default server is not allowed.

18. Review the default NTP host in the running-configuration.

```
show ntp servers
```

```
sw-edge1(config)# show ntp servers
-------------------------------------------------------
          NTP SERVER KEYID MINPOLL MAXPOLL OPTION VER
-------------------------------------------------------
pool.ntp.org(default)    --       4       4 iburst   4
-------------------------------------------------------
```

19. Configure the server **10.254.1.21** as NTP server.

```
ntp server 10.254.1.21 iburst
ntp enable
```

```
sw-edge1(config)# ntp server 10.254.1.21 iburst
sw-edge1(config)# ntp enable
```

20. Review NTP time updates. Your system time may have been updates using the pool.ntp.org server already.

```
show event -r | i ntp
```

```
sw-edge1(config)# show event -r | i ntp
2022-10-15T19:18:31.175905+00:00 sw-edge1 ntp-mgrd[4382]: Event|1101|LOG_INFO|||NTP
Association Add : 10.254.1.21 iburst version 4
2022-10-15T19:18:03.342892+00:00 sw-edge1 ntpd[4669]: Event|System date/time changed
from 2022-10-15 19:20:48 to 2022-10-15 19:20:48 using 193.187.181.6
2022-10-15T12:17:48.336546+00:00 6300 ntp-mgrd[4382]: Event|1107|LOG_INFO|||NTP
enable state change : NTP is disabled -> NTP is enabled
2022-10-15T12:17:43.173512+00:00 6300 ntp-mgrd[4382]: Event|1101|LOG_INFO|||NTP
Association Add : pool.ntp.org iburst version 4 minpoll 4 maxpoll 4
```

21. Review the current time.

```
show clock
```

```
sw-edge1(config)# show clock
Sat Oct 15 19:20:14 UTC 2022
System is configured for timezone : UTC
```

22. Configure the correct time zone for the local system.

```
clock timezone us/michigan
```

```
sw-edge1(config)# clock timezone us/michigan
```

23. Verify the current time now reflects the configured time zone.

```
show clock
```

```
sw-edge1(config)# show clock
Sat Oct 15 15:20:31 EDT 2022
System is configured for timezone : US/Michigan
```

The time on the switch should now be the same as the time on the MGMT PC. This system has also been configured in this time zone.

Task 2: Practice on Edge Switch 2

Objectives

In this task you will practice applying the same settings on sw-edge2. You can use the solution if you are unsure about the steps or if you want to verify your steps.

On sw-edge2:

- Configure SVI3 (VLAN 3) IP address 10.1.3.5/24 with default gateway to 10.1.3.1.
- Configure 10.254.1.21 as DNS server.
- Configure 10.254.1.21 as the NTP server and define the appropriate time zone.
- Verify connectivity.

Solution

Here is the configuration for sw-edge2 if you need help or to verify your configuration:

```
vlan 3

interface vlan 3
 ip address 10.1.3.5/24
 exit

interface vlan 1
 no ip dhcp
 shutdown
 exit

ip route 0.0.0.0/0 10.1.3.1
ip dns server-address 10.254.1.21

ntp server 10.254.1.21 iburst
ntp enable
clock timezone us/michigan
```

1. Save the configuration on both edge switches.

```
write memory
```

You have completed this Lab!

Lab 03.03 – Static Routes

Overview

This lab will introduce you to static IP routes and static route summarization.

Objectives

After completing this lab, you will be able to:

- Configure static routes in the routing table.
- Understand the operation of static routes.
- Review the IP routing table.

Task 1: Configure a Static Route between Aggregation and Core

Your senior colleague has asked you to explore the IP routing table and make a static route configuration to reach the datacenter subnet 10.254.1.0/24. In this task the default route on the aggregation switch sw-agg1 will be removed and a static route for the Datacenter 1 subnet will be added. Datacenter 2 will be covered in the next task.

You will also explore the source IP address used when performing a ping on a routing switch.

The last step is to ensure there is bi-directional communication by configuring the core router with a static route for the return traffic.

Objectives

- Configure static routes
- Verify static routes

Steps

Remove default route on sw-agg1

On sw-agg1, the default route must be removed to test the static route function in this lab. The default route was part of the initial configuration of sw-agg1.

1. On sw-agg1, remove the default route.

```
no ip route 0.0.0.0/0 10.254.101.254
```

```
sw-agg1(config)# no ip route 0.0.0.0/0 10.254.101.254
```

2. Verify the default route has been removed from the IP routing table.

```
show ip route | include 0.0.0.0
show ip route 0.0.0.0
```

```
sw-agg1(config)# show ip route | include 0.0.0.0
sw-agg1(config)#
```

```
sw-agg1(config)# show ip route 0.0.0.0

No ipv4 routes configured

sw-agg1(config)#
```

Review the routed uplink to core router

3. On sw-agg1, review the IP interface list, you should have port 1/1/8 configured with an IP address.

```
show ip interface brief
```

```
sw-agg1(config)# show ip interface brief
Interface         IP Address             Interface Status
                                           link/admin
...
1/1/8           10.254.101.2/24           up/up
...
```

4. Review the configuration of port 1/1/8 (uplink to core router). The core router has IP address 10.254.101.254.

```
show running-config interface 1/1/8
```

```
sw-agg1(config)# show running-config interface 1/1/8
interface 1/1/8
    no shutdown
    mtu 9198
    description rtr-core1-1/1/1
    ip address 10.254.101.2/24
    exit
```

5. Review the IP routing table, there should be a directly connected route for the 10.254.101.0/24 subnet.

```
show ip route
```

```
sw-agg1(config)# show ip route

Displaying ipv4 routes selected for forwarding

Origin Codes: C - connected, S - static, L - local
              R - RIP, B - BGP, O - OSPF
Type Codes:   E - External BGP, I - Internal BGP, V - VPN, EV - EVPN
              IA - OSPF internal area, E1 - OSPF external type 1
              E2 - OSPF external type 2

VRF: default

Prefix         Nexthop          Interface     VRF(egress)       Origin/   Distance/
Age
                                                                 Type      Metric
-------------------------------------------------------------------------------------
-----
10.1.1.0/24      -                vlan1         -                 C         [0/0]
-
10.1.1.1/32      -                vlan1         -                 L         [0/0]
-
10.1.2.0/24      -                vlan2         -                 C         [0/0]
-
10.1.2.2/32      -                vlan2         -                 L         [0/0]
-
10.1.3.0/24      -                vlan3         -                 C         [0/0]
-
10.1.3.1/32      -                vlan3         -                 L         [0/0]
-
```

```
10.1.11.0/24      -                 vlan11          -                 C          [0/0]
-
10.1.11.1/32      -                 vlan11          -                 L          [0/0]
-
10.1.12.0/24      -                 vlan12          -                 C          [0/0]
-
10.1.12.1/32      -                 vlan12          -                 L          [0/0]
-
10.254.101.0/24   -                 1/1/8           -                 C          [0/0]
-
10.254.101.2/32   -                 1/1/8           -                 L          [0/0]
-
```

6. For troubleshooting (with a larger routing table), it can be useful to check if the IP routing table has a route available to reach a given destination IP address. Check if there is a route available for IP 10.254.101.254. This should reveal the same directly connected route from the previous step.

```
show ip route 10.254.101.254
```

```
sw-agg1(config)# show ip route 10.254.101.254

VRF: default

   Prefix            : 10.254.101.0/24
VRF(egress)        : -
   Nexthop           : -
Interface          : 1/1/8
   Origin            : connected
Type               : -
   Distance          : 0
Metric             : 0
   Age               : -
Tag                : 0
   Encap Type        : -                                        Encap
Details      : -
```

7. You may also use the pipe function to filter the command output. In this example, you will only receive the output lines that **include** the text **10.254.101**.

```
show ip route | include 10.254.101
```

```
sw-agg1(config)# show ip route | include 10.254.101
10.254.101.0/24   -                 1/1/8           -                 C          [0/0]
-
10.254.101.2/32   -                 1/1/8           -                 L          [0/0]
-
```

8. Verify connectivity to the core router using the directly connected IP, this should be successful.

```
ping 10.254.101.254
```

9. Attempt to reach an IP in the server subnet (10.254.1.21).

```
ping 10.254.1.21
```

```
sw-agg1(config)# ping 10.254.1.21
ping4: connect: Network is unreachable
```

Question: Why did this ping fail?

_____

Answer: There is currently no IP route in the routing table for the requested destination IP.

10. Lookup the requested destination IP in the local routing table to see if there are any matching routes.

```
show ip route 10.254.1.21
```

```
sw-agg1(config)# show ip route 10.254.1.21

No ipv4 routes configured
```

Question: What do you observe?

Answer: There is currently no matching route for the requested address.

Configure static route to the datacenter subnet

11. On sw-agg1, add a static route for the 10.254.1.0/24 subnet with the core router as the next hop IP address.

```
ip route 10.254.1.0/24 10.254.101.254
```

```
sw-agg1(config)# ip route 10.254.1.0/24 10.254.101.254
```

12. Lookup the requested destination IP in the local routing table to see if there are any matching routes.

```
show ip route 10.254.1.21
```

```
sw-agg1(config)# show ip route 10.254.1.21

VRF: default

   Prefix           : 10.254.1.0/24
VRF(egress)        : -
   Nexthop          : 10.254.101.254
Interface          : 1/1/8
   Origin           : static
Type               : -
   Distance         : 1
Metric             : 0
   Age              : 00h:01m:09s
Tag                : 0
   Encap Type       : -                                          Encap
Details      : -
```

13. Attempt to ping the host in the server subnet (10.254.1.21). This should now be successful.

```
ping 10.254.1.21
```

Question: What is the source IP address used for this ping?

_____

Answer: By default, a system with multiple IP interfaces (such as the aggregation switch) will use the outbound interface IP address as the source IP. In this example the IP address that is configured on port 1/1/8.

14. Attempt to ping 10.254.1.21 using source IP address of SVI1

```
ping 10.254.1.21 source vlan1
```

```
sw-agg1(config)# ping 10.254.1.21 source vlan1
PING 10.254.1.21 (10.254.1.21) from 10.1.1.1 : 100(128) bytes of data.

--- 10.254.1.21 ping statistics ---
5 packets transmitted, 0 received, 100% packet loss, time 4081ms
```

Question: Did this ping work?

Answer: No.

Question: In the first line of the ping command output, the system will report *PING … from* …. What is the from (source) IP address used for this ping?

Answer: The source IP address is now 10.1.1.1. This is the SVI 1 IP address, as requested using the source command option.

15. Now try to identify what routing hop (router) is having trouble with this traffic flow. Use **traceroute** to setup an IP trace to destination IP 10.254.1.21 using source interface SVI1 (IP 10.1.1.1).

```
traceroute 10.254.1.21 source vlan1
```

Question: Does any hop show a reply in the output? Why?

_____

Answer: When a different source IP is used, it is possible that there is no return route for the specified source subnet. In this example, it means that the ping packet did arrive at the destination host, but the return route may be missing on the return path. In this example, the fact that not even 1 hop responds, forms an indication that the issue may be on the next-hop (the core router).

16. Stop the traceroute using CTRL-C.

```
sw-agg1(config)# traceroute 10.254.1.21 source vlan1
traceroute to 10.254.1.21 (10.254.1.21), 1 hops min, 30 hops max, 3 sec. timeout, 3
probes
^C
```

17. Use MGMT PC to open an SSH connection to the core router (hostname **rtr-core1**).

18. On rtr-core1, review the IP routing table.

```
show ip route
```

19. First check if any routes are available for the 10.254.101.2 IP address. This is the IP address of sw-agg1 port 1/1/8.

```
show ip route 10.254.101.2
```

```
rtr-core1# show ip route 10.254.101.2

VRF: default

   Prefix            : 10.254.101.0/24                          VRF(egress)
: -
   Nexthop           : -                                        Interface
: 1/1/1
   Origin            : connected                                Type
: -
   Distance          : 0                                        Metric
: 0
   Age               : -                                        Tag
: 0
   Encap Type        : -                                        Encap Details
: -
```

Question: Are there any matching routes for this IP? What is the type of route?

_____

Answer: Yes, this is a directly connected subnet on the core router.

20. Next check if any routes are available for the 10.1.1.1 IP address, this is the sw-agg1 SV1 address that is used in the test ping and traceroute.

```
show ip route 10.1.1.1
```

```
rtr-core1# show ip route 10.1.1.1

VRF: default

   Prefix            : 10.0.0.0/8                               VRF(egress)
: -
   Nexthop           : -                                        Interface
: blackhole
   Origin            : static                                   Type
: -
   Distance          : 1                                        Metric
: 0
   Age               : 00h:03m:23s                              Tag
: 0
   Encap Type        : -                                        Encap Details
: -
```

Question: Are there any matching routes for this IP address? If so, why did the ping not succeed?

_____

Answer: There is a matching route on the core router, but this is a 10.0.0.0/8 summary route that points to the blackhole interface. At this moment, the core router simply drops (blackholes) traffic destined to 10.1.1.1, instead of sending it back to the sw-agg1.

> **NOTE:** This route has been configured on the core router to ensure that traffic to any address in the 10.0.0.0/8 range that is not handled by a more specific subnet is dropped by the core router. This can help prevent routing loops when route summarization occurs at multiple places in the network.

21. Add a specific route for the 10.1.1.0/24 subnet to next-hop IP of sw-agg1.

```
ip route 10.1.1.0/24 10.254.101.2
```

```
rtr-core1(config)# ip route 10.1.1.0/24 10.254.101.2
```

22. Review all the configured static routes in the routing table. Your new static route should be listed in the output.

```
show ip route static
```

```
rtr-core1(config)# show ip route static

Displaying ipv4 routes selected for forwarding

Origin Codes: C - connected, S - static, L - local
              R - RIP, B - BGP, O - OSPF
Type Codes:   E - External BGP, I - Internal BGP, V - VPN, EV - EVPN
              IA - OSPF internal area, E1 - OSPF external type 1
              E2 - OSPF external type 2

VRF: default

Prefix        Nexthop       Interface  VRF(egress)   Origin/  Distance/    Age
                                                     Type     Metric
-------------------------------------------------------------------------------
---------
0.0.0.0/0     10.254.1.253   1/1/9          -           S         [1/0]
00h:04m:07s
10.0.0.0/8    -              blackhole  -               S         [1/0]
00h:04m:07s
10.1.1.0/24   10.254.101.2   1/1/1          -           S         [1/0]
00h:00m:14s
10.1.3.0/24   10.254.101.2   1/1/1          -           S         [255/0]
00h:04m:07s
10.1.11.0/24  10.254.101.2   1/1/1          -           S         [255/0]
00h:04m:07s
10.1.12.0/24  10.254.101.2   1/1/1          -           S         [255/0]
00h:04m:07s

Total Route Count : 6
```

23. Verify in the routing table that the new route will be used to reach destination IP 10.1.1.2. Notice the difference with the previous command output 3 steps earlier.

```
show ip route 10.1.1.1
```

```
rtr-core1(config)# show ip route 10.1.1.1

VRF: default

   Prefix            : 10.1.1.0/24                                VRF(egress)
: -
   Nexthop           : 10.254.101.2                               Interface
: 1/1/1
   Origin            : static                                     Type
: -
   Distance          : 1                                          Metric
: 0
   Age               : 00h:02m:05s                                Tag
: 0
   Encap Type        : -                                          Encap Details
: -
```

24. On sw-agg1, verify that the ping to 10.254.1.21 works using SVI1 as the source IP.

```
ping 10.254.1.21 source vlan1
```

```
sw-agg1(config)# ping 10.254.1.21 source vlan1
PING 10.254.1.21 (10.254.1.21) from 10.1.1.1 : 100(128) bytes of data.
108 bytes from 10.254.1.21: icmp_seq=1 ttl=127 time=2.35 ms
108 bytes from 10.254.1.21: icmp_seq=2 ttl=127 time=2.22 ms
108 bytes from 10.254.1.21: icmp_seq=3 ttl=127 time=1.98 ms
108 bytes from 10.254.1.21: icmp_seq=4 ttl=127 time=1.98 ms
108 bytes from 10.254.1.21: icmp_seq=5 ttl=127 time=1.93 ms

--- 10.254.1.21 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4004ms
rtt min/avg/max/mdev = 1.930/2.091/2.347/0.162 ms
```

Question: Was the ping successful?

Answer: Yes, routes are now available for both transmit and return traffic.

Task 2: Route Aggregation with Static Routes

In this task you will explore route aggregation with static routes. The static route of the previous task only worked for SVI 1 communication. All other subnets connected to the aggregation layer would require their own static routes on the core router.

Your customer wants to know if, instead of configuring a static route for each IP subnet connected to the aggregation switches, a simplified static route configuration can be applied. You have heard that route aggregation can be used to aggregate or summarize subnets into an aggregate route.

In this task you will configure an aggregate route to handle all current and future subnets connected to sw-agg1. You will also ensure that both the current Datacenter 1 and the future Datacenter 2 can be reached using a single route.

Datacenter subnet list:

Network	Subnet
Datacenter 1	10.254.1.0/24
Datacenter 2 (future)	10.254.2.0/24

Objectives

- Configure static route aggregation.

Steps

Prepare core router for this task

The core router has been pre-configured with some static routes in this lab environment to support the initial lab exercises with VLAN 11 and 12. To complete the current lab exercise, you must remove one of these static routes.

1. On the core router, remove the static route for the VLAN 12 subnet (10.1.12.0/24).

```
no ip route 10.1.12.0/24 10.254.101.2
```

```
rtr-core1(config)# no ip route 10.1.12.0/24 10.254.101.2
```

Verify that SVI12 cannot access the datacenter subnet.

2. On sw-agg1, attempt to ping 10.254.1.21 using source IP SVI12, this should fail.

```
ping 10.254.1.21 source vlan12
```

```
sw-agg1(config)# ping 10.254.1.21 source vlan12
PING 10.254.1.21 (10.254.1.21) from 10.1.12.1 : 100(128) bytes of data.

--- 10.254.1.21 ping statistics ---
5 packets transmitted, 0 received, 100% packet loss, time 4088ms
```

Configure aggregate route

3. On the core router, attempt to find a route in the routing table for the sw-agg1 SVI12 IP address (10.1.12.1).

```
show ip route 10.1.12.1
```

```
rtr-core1(config)# show ip route 10.1.12.1

VRF: default

   Prefix            : 10.0.0.0/8                          VRF(egress)       : -
   Nexthop           : -                                   Interface         :
blackhole
   Origin            : static                              Type              : -
   Distance          : 1                                   Metric            : 0
   Age               : 00h:09m:42s                         Tag               : 0
   Encap Type        : -                                   Encap Details     : -
```

Question: Is there any matching route?

Answer: Yes, the requested destination IP matches the 10.0.0.0/8 blackhole route, so traffic to 10.1.12.1 will be dropped by the core router. A destination of *blackhole* in a static route will cause the routing switch to *drop* all traffic that matches this route entry.

In the lab IP plan, all current and future campus IP subnets that connect to the aggregation switch belong to the 10.1.0.0/16 range.

4. On the core router, add an aggregate route 10.1.0.0/16 for the campus subnets using sw-agg1 as the next-hop IP.

```
ip route 10.1.0.0/16 10.254.101.2
```

```
rtr-core1(config)# ip route 10.1.0.0/16 10.254.101.2
```

5. Again, check what route will be used to reach IP 10.1.12.1. The output should now show the 10.1.0.0/16 route.

```
show ip route 10.1.12.1
```

```
rtr-core1(config)# show ip route 10.1.12.1

VRF: default

   Prefix            : 10.1.0.0/16                         VRF(egress)       : -
   Nexthop           : 10.254.101.2                        Interface         :
1/1/1
   Origin            : static                              Type              : -
   Distance          : 1                                   Metric            : 0
   Age               : 00h:00m:10s                         Tag               : 0
   Encap Type        : -                                   Encap Details     : -
```

6. On sw-agg1, verify that the previous ping test is now successful.

```
ping 10.254.1.21 source vlan12
```

```
sw-agg1(config)# ping 10.254.1.21 source vlan12
PING 10.254.1.21 (10.254.1.21) from 10.1.12.1 : 100(128) bytes of data.
108 bytes from 10.254.1.21: icmp_seq=1 ttl=127 time=1.93 ms
108 bytes from 10.254.1.21: icmp_seq=2 ttl=127 time=2.18 ms
108 bytes from 10.254.1.21: icmp_seq=3 ttl=127 time=1.69 ms
108 bytes from 10.254.1.21: icmp_seq=4 ttl=127 time=1.25 ms
108 bytes from 10.254.1.21: icmp_seq=5 ttl=127 time=2.07 ms

--- 10.254.1.21 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4004ms
rtt min/avg/max/mdev = 1.245/1.823/2.178/0.332 ms
```

Datacenter aggregation

Now you have connectivity with the first datacenter (10.254.1.0/24). The customer informs you that there will be a second datacenter in the future. This second datacenter will use IP subnet 10.254.2.0/24.
The customer wants you to update sw-agg1, so a single static route is used to reach any IP address that starts with ***10.254*** in the first 2 octets.

The customer has already activated a test IP address (10.254.2.1) on the core router. This IP address is in the IP range of the second datacenter. This address 10.254.2.1 can be used to test the connectivity for this second datacenter from the client PCs.

7. On sw-agg1, attempt to ping 10.254.**2.**1 using source IP SVI12, which should fail at this point.

```
ping 10.254.2.1 source vlan12
```

```
sw-agg1(config)# ping 10.254.2.1 source vlan12
ping4: sendmsg: Network is unreachable
PING 10.254.2.1 (10.254.2.1) from 10.1.12.1 : 100(128) bytes of data.
ping4: sendmsg: Network is unreachable
ping4: sendmsg: Network is unreachable
ping4: sendmsg: Network is unreachable
ping4: sendmsg: Network is unreachable

--- 10.254.2.1 ping statistics ---
5 packets transmitted, 0 received, 100% packet loss, time 4116ms
```

8. On sw-agg1, add an aggregate route 10.254.0.0/16 for the datacenter subnets using the core router as the next hop IP.

```
ip route 10.254.0.0/16 10.254.101.254
```

```
sw-agg1(config)# ip route 10.254.0.0/16 10.254.101.254
```

9. On sw-agg1, verify 10.254.2.1 is now reachable.

```
ping 10.254.2.1 source vlan12
```

```
sw-agg1(config)# ping 10.254.2.1 source vlan12
PING 10.254.2.1 (10.254.2.1) from 10.1.12.1 : 100(128) bytes of data.
108 bytes from 10.254.2.1: icmp_seq=1 ttl=64 time=1.59 ms
108 bytes from 10.254.2.1: icmp_seq=2 ttl=64 time=1.55 ms
108 bytes from 10.254.2.1: icmp_seq=3 ttl=64 time=1.85 ms
108 bytes from 10.254.2.1: icmp_seq=4 ttl=64 time=1.76 ms
108 bytes from 10.254.2.1: icmp_seq=5 ttl=64 time=1.54 ms

--- 10.254.2.1 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4005ms
rtt min/avg/max/mdev = 1.540/1.657/1.852/0.126 ms
```

Most specific route

It is possible that multiple routes could be used for a given destination, and so routers must decide which route to choose. The first decision is to pick the *most matching* route, this means the route that has the longest mask. In these steps you will explore this behavior.

10. On the sw-agg1, review the route that would be used to reach 10.254.1.21.

```
show ip route 10.254.1.21
```

```
sw-agg1(config)# show ip route 10.254.1.21

VRF: default

   Prefix            : 10.254.1.0/24                       VRF(egress)       : -
   Nexthop           : 10.254.101.254                      Interface         :
1/1/8
   Origin            : static                              Type              : -
   Distance          : 1                                   Metric            : 0
   Age               : 00h:32m:22s                         Tag               : 0
   Encap Type        : -                                   Encap Details     : -
```

11. Now review the route that would be used to reach 10.254.2.1 (the second datacenter).

```
show ip route 10.254.2.1
```

```
sw-agg1(config)# show ip route 10.254.2.1

VRF: default

   Prefix            : 10.254.0.0/16                       VRF(egress)       : -
   Nexthop           : 10.254.101.254                      Interface         :
1/1/8
   Origin            : static                              Type              : -
   Distance          : 1                                   Metric            : 0
   Age               : 00h:01m:02s                         Tag               : 0
   Encap Type        : -                                   Encap Details     : -
```

Question: Do you notice any difference? Why?

_____

Answer: For the 10.254.1.21 IP, there is a more matching route /24 (24 bits of the subnet mask match), for the 10.254.2.1 IP, the only available route is the /16 aggregate route.

Since the aggregate routes have now been added to the routing tables, the more specific routes can be removed.

12. On sw-agg1, remove the route for the 10.254.1.0/24 subnet.

```
no ip route 10.254.1.0/24 10.254.101.254
```

```
sw-agg1(config)# no ip route 10.254.1.0/24 10.254.101.254
```

13. Again, review the route used to reach the 10.254.1.21 host, this should now show the /16 aggregate route instead of the more specific /24 route from the previous command output.

```
show ip route 10.254.1.21
```

```
sw-agg1(config)# show ip route 10.254.1.21

VRF: default

   Prefix            : 10.254.0.0/16                    VRF(egress)       : -
   Nexthop           : 10.254.101.254                   Interface         :
1/1/8
   Origin            : static                           Type              : -
   Distance          : 1                                Metric            : 0
   Age               : 00h:01m:59s                      Tag               : 0
   Encap Type        : -                                Encap Details     : -
```

14. On the **core router**, remove the specific route for the 10.1.1.0/24 subnet.

```
no ip route 10.1.1.0/24 10.254.101.2
```

```
rtr-core1(config)# no ip route 10.1.1.0/24 10.254.101.2
```

15. Verify that the core router is now using the aggregate route 10.1.0.0/16 for the 10.1.1.2 destination.

```
show ip route 10.1.1.1
```

```
rtr-core1(config)# show ip route 10.1.1.1

VRF: default

   Prefix            : 10.1.0.0/16                                VRF(egress)
: -
   Nexthop           : 10.254.101.2                               Interface
: 1/1/1
   Origin            : static                                     Type
: -
   Distance          : 1                                          Metric
: 0
   Age               : 00h:04m:19s                                Tag
: 0
   Encap Type        : -                                          Encap Details
: -
```

16. Save the configuration on all switches and the core router.

```
write memory
```

You have completed this Lab!

Lab 03.04 – Basic OSPF Configuration

Overview

Your customer is pleased to see that the static routes and route aggregation have simplified the routing table. Now they have explained they may need to add several other sites with multiple links, and they are concerned that it may become hard to maintain the routing tables on all these devices.

After consulting with your senior colleague, you hear that Open Shortest Path First (OSPF) is a routing protocol that can help with the distribution of routes and calculate the best paths in the network. In this lab you will introduce OSPF in the network and explore its features.

Objectives

After completing this lab, you will be able to:

- Start OSPF process on AOS-CX devices
- Configure a single area OSPF
- Enable interface for OSPF
- Verify OSPF neighbor adjacency
- Verify OSPF exchange of routes
- Advertise a default route using OSPF

Task 1: Configure Loopback Interfaces on Aggregation Switches

In this task you will configure an IP Loopback interface on the Aggregation switches.

Objectives

- Configure a Loopback Interface

Steps

In the next steps you will configure a Loopback interface on the sw-agg1 aggregation switch. A loopback interface is an internal, software interface. Since there is no physical subnet, there is only one host (the device itself). Therefore, typically a subnet mask of /32 is configured for the IP address of a loopback interface.

A loopback interface is always up. Its state is independent of the state of any physical interface on the system. Loopback interfaces can be used as the source IP address for outbound connections. This ensures that, for example, all outbound SNMP, RADIUS or packets to Aruba Central have the same, stable source IP address.

1. On **sw-agg1**, configure a new loopback interface with id 0.

```
interface loopback 0
```

```
sw-agg1(config)# interface loopback 0
```

2. Configure IP address 10.1.0.2/32.

```
 ip address 10.1.0.2/32
 exit
```

```
sw-agg1(config-loopback-if)#  ip address 10.1.0.2/32
sw-agg1(config-loopback-if)#  exit
```

3. Review the IP interface list shows the correct new Loopback 0 and IP address 10.1.0.2/32.

```
show ip interface brief
```

```
sw-agg1(config)# show ip interface brief
Interface         IP Address              Interface Status
                                            link/admin
...
loopback0        10.1.0.2/32                up/up
...
```

4. Review the IP routing table contains a /32 route for the loopback IP address.

```
show ip route
```

```
sw-agg1(config)# show ip route

Displaying ipv4 routes selected for forwarding
```

```
Origin Codes: C - connected, S - static, L - local
              R - RIP, B - BGP, O - OSPF
Type Codes:   E - External BGP, I - Internal BGP, V - VPN, EV - EVPN
              IA - OSPF internal area, E1 - OSPF external type 1
              E2 - OSPF external type 2

VRF: default

Prefix              Nexthop     Interface     VRF(egress)       Origin/    Distance/
Age
                                                                Type       Metric
-------------------------------------------------------------------------------------
-----
10.1.0.2/32         -           loopback0     -                 L          [0/0]
-
...
```

5. Verify you can locally ping the loopback IP address, this should be successful.

```
ping 10.1.0.2
```

6. On **sw-agg2**, configure Loopback 0 with IP address 10.1.0.3/32.

```
interface loopback 0
 ip address 10.1.0.3/32
 exit
```

```
sw-agg2(config)# interface loopback 0
sw-agg2(config-loopback-if)#  ip address 10.1.0.3/32
sw-agg2(config-loopback-if)#  exit
```

7. Verify it is listed in the IP interface list.

```
show ip interface brief
```

```
sw-agg1(config)# show ip interface brief
Interface         IP Address              Interface Status
                                            link/admin
...
loopback0         10.1.0.3/32               up/up
...
```

Task 2: Configure OSPF on Aggregation Switches

OSPF uses Link State Advertisements (LSA) to advertises the topology and the routes. These LSAs are added to the Link State Database (LSDB), but they require a unique identifier. The unique identifier for each OSPF router is configured using an OSPF router ID. You must create an OSPF area to store the LSAs and enable the layer3 interfaces for OSPF.

Objectives

In this task you will:

- Enable OSPF on the aggregation switches
- Configure OSPF to use the loopback IP as the router ID
- Enable OSPF on the SVI interface between the aggregation switches
- Review the status of the OSPF process

Steps

Configure OSPF process on sw-agg1

1. On **sw-agg1**, configure OSPF process 1.

```
router ospf 1
```

```
sw-agg1(config)# router ospf 1
```

2. Set the router-id to 10.1.0.2.

```
 router-id 10.1.0.2
```

```
sw-agg1(config-ospf-1)# router-id 10.1.0.2
```

3. Configure the area 0 (backbone area) and exit the OSPF context.

```
 area 0
 exit
```

```
sw-agg1(config-ospf-1)# area 0
sw-agg1(config-ospf-1)# exit
```

4. Verify that OSPF is enabled.

```
show ip ospf
```

```
sw-agg1(config)# show ip ospf
VRF : default                          Process  : 1
-----------------------------------------------------

RouterID               : 10.1.0.2          OSPFv2                 : Enabled
BFD                    : Disabled          SPF Start Interval     : 200   ms
SPF Hold Interval      : 1000  ms          SPF Max Wait Interval  : 5000  ms
LSA Start Time         : 5000  ms          LSA Hold Time          : 0     ms
```

```
LSA Max Wait Time      : 0     ms            LSA Arrival             : 1000  ms
External LSAs          : 0                   Checksum Sum            : 0
ECMP                   : 4                   Reference Bandwidth     : 100000 Mbps
Area Border            : false               AS Border               : false
GR Status              : Enabled             GR Interval             : 120 sec
GR State               : inactive            GR Exit Status          : none
GR Helper              : Enabled             GR Strict LSA Check     : Enabled
GR Ignore Lost I/F     : Disabled
Summary address:

Area      Total     Active
-----------------------------
Normal     1         0
Stub       0         0
NSSA       0         0

Area  : 0.0.0.0
----------------
Area Type              : Normal          Status                : Inactive
Total Interfaces       : 0               Active Interfaces     : 0
Passive Interfaces     : 0               Loopback Interfaces   : 0
SPF Calculation Count  : 1
Area ranges     :
Number of LSAs         : 0               Checksum Sum          : 0
```

Question: What is the router-id used by OSPF?

Answer: The router id should be 10.1.0.2

5. Check the OSPF Link State Database (LSDB).

```
show ip ospf lsdb
```

```
sw-agg1(config)# show ip ospf lsdb
No OSPF LSAs found on VRF default.
```

Question: Why is the LSDB empty?

Answer: Since there are no interfaces enabled for OSPF, no Router LSA is generated.

6. Activate the loopback interface for OSPF process 1 area 0.

```
interface loopback 0
 ip ospf 1 area 0
 exit
```

```
sw-agg1(config)# interface loopback 0
sw-agg1(config-loopback-if)#  ip ospf 1 area 0
sw-agg1(config-loopback-if)#  exit
```

7. Verify the LSDB again.

```
show ip ospf lsdb
```

```
sw-agg1(config)# show ip ospf lsdb
OSPF Router with ID (10.1.0.2) (Process ID 1 VRF default)
=========================================================

Router Link State Advertisements (Area 0.0.0.0)
-----------------------------------------------

LSID            ADV Router      Age       Seq#       Checksum       Link Count
-------------------------------------------------------------------------------
10.1.0.2        10.1.0.2        13        0x80000001 0x00001c10     1
```

Question: How can you recognize that the Router LSA is generated by your own system?

_____

Answer: The LSA ID is based on the local OSPF router ID, in this example 10.1.0.2.

Question: What is the link count for this Router LSA?

_____

Answer: Only one link has been enabled for OSPF (the loopback). This is the only link that shows up in the Router LSA. This also demonstrates that a router may have many IP interfaces or SVIs, but only the OSPF-enabled interfaces are included in the Router LSA advertisement.

Adding additional interfaces

8. On sw-agg1, enable interfaces SVI 11 and SVI 12 for OSPF. This will ensure that these subnets are included in the sw-agg1 Router LSA. Routes that are included in the Router LSA are ‘advertised’ to the network. This means they will be reachable for other OSPF routers.

```
interface vlan11
 ip ospf 1 area 0
 exit
interface vlan12
 ip ospf 1 area 0
 exit
```

```
sw-agg1(config)# interface vlan11
sw-agg1(config-if-vlan)#  ip ospf 1 area 0
sw-agg1(config-if-vlan)#  exit
sw-agg1(config)# interface vlan12
sw-agg1(config-if-vlan)#  ip ospf 1 area 0
sw-agg1(config-if-vlan)#  exit
```

9. Review the LSDB and the link count.

```
show ip ospf lsdb
```

```
sw-agg1(config)# show ip ospf lsdb
OSPF Router with ID (10.1.0.2) (Process ID 1 VRF default)
=========================================================

Router Link State Advertisements (Area 0.0.0.0)
```

```
-----------------------------------------------

LSID            ADV Router      Age       Seq#       Checksum       Link Count
-------------------------------------------------------------------------------
10.1.0.2        10.1.0.2        16        0x80000001 0x0000ac69     3
```

Question: How many links are included in the Router LSA?

_____

Answer: There are now three links in this Router LSA: the loopback subnet, the SVI 11, and SVI 12 subnets.

Task 3: Configure and Tune OSPF Links to Core Router

In this task you will configure the connection between the sw-agg1 and core router. You will observe the OSPF interface statistics and adjust the OSPF interface timers.

A link with only two routers is commonly configured as an OSPF point-to-point (P2P) link. P2P links speed up the OSPF neighbor adjacency setup and simplify the LSDB topology.

Objectives

- Configure an OSPF Point to Point link

Steps

1. On sw-agg1, enable port 1/1/8 (to the core router) for OSPF process 1 area 0.

```
interface 1/1/8
 ip ospf 1 area 0
 exit
```

```
sw-agg1(config)# interface 1/1/8
sw-agg1(config-if)#  ip ospf 1 area 0
sw-agg1(config-if)#  exit
```

2. Review the LSDB link count for your Router LSA.

```
show ip ospf lsdb
```

```
sw-agg1(config)# show ip ospf lsdb
OSPF Router with ID (10.1.0.2) (Process ID 1 VRF default)
==========================================================

Router Link State Advertisements (Area 0.0.0.0)
-----------------------------------------------

LSID            ADV Router      Age       Seq#       Checksum       Link Count
-------------------------------------------------------------------------------
10.1.0.2        10.1.0.2        12        0x80000001 0x0000d8b4     4
```

NOTE: Between OSPF routers on the same link, the OSPF Hello and Dead timers must match. A mismatch will prevent OSPF routers from establishing an adjacency. It is common practice to lower the Hello and Dead values from the default 10 and 40 seconds to 1 and 4 seconds. This change was already made on the core router side.

If you want to review the existing OSPF configuration (including the timer configuration) on the core router, access the core router and run this command:

```
show running-configuration ospf
```

3. Review the OSPF interface statistics for port 1/1/8.

```
show ip ospf statistics interface 1/1/8
```

```
sw-agg1(config)# show ip ospf statistics interface 1/1/8
OSPF Process ID 1 VRF default, interface 1/1/8 statistics  (cleared 0h0m38s ago)
================================================================================

Tx Hello Packets      : 4            Rx Hello Packets      : 0
Tx Hello Bytes        : 256          Rx Hello Bytes        : 0
Tx DD Packets         : 0            Rx DD Packets         : 0
Tx DD Bytes           : 0            Rx DD Bytes           : 0
Tx LS Request Packets : 0            Rx LS Request Packets : 0
Tx LS Request Bytes   : 0            Rx LS Request Bytes   : 0
Tx LS Update Packets  : 0            Rx LS Update Packets  : 0
Tx LS Update Bytes    : 0            Rx LS Update Bytes    : 0
Tx LS Ack Packets     : 0            Rx LS Ack Packets     : 0
Tx LS Ack Bytes       : 0            Rx LS Ack Bytes       : 0

Total Number of State Changes : 39
Number of LSAs                : 0
LSA Checksum Sum              : 0
Total Transmit Failures       : 0
Total OSPF Packets Discarded  : 76

Reason                        Packets Dropped
--------------------------------------------
Invalid type                  0
Invalid length                0
Invalid checksum              0
Invalid version               0
Bad or unknown source         0
Area mismatch                 0
Self-originated               0
Duplicate router ID           0
Interface standby             0
Total Hello packets dropped   38
  Network Mask mismatch       0
  Hello interval mismatch     38
  Dead interval mismatch      0
  Options mismatch            0
  MTU mismatch                0
  Neighbor ignored            0
Authentication errors         0
  Type mismatch               0
  Authentication failures     0
Wrong protocol                0
Resource failures             0
Bad LSA length                0
Bad DD packets                0
Others                        38

Total LSAs Ignored : 0
Bad Type           : 0
Bad Length         : 0
Invalid Data       : 0
Invalid Checksum   : 0
```

Question: Are there any non-zero statistics for Packets dropped?

_____

Answer: Yes, the 'Hello interval mismatch' shows several dropped packets.

4. Adjust the local timers of port 1/1/8 to match the core router side. Set the Hello timer to 1 second and the Dead timer to 4 seconds.

```
interface 1/1/8
 ip ospf hello-interval 1
 ip ospf dead-interval 4
```

```
sw-agg1(config)# interface 1/1/8
sw-agg1(config-if)#  ip ospf hello-interval 1
sw-agg1(config-if)#  ip ospf dead-interval 4
```

5. Review the OSPF interface statistics again for port 1/1/8. The packets dropped should no longer increase, and Rx Hello packets should be a non-zero value now.

```
show ip ospf 1 statistics interface 1/1/8
```

```
sw-agg1(config-if)# show ip ospf 1 statistics interface 1/1/8
OSPF Process ID 1 VRF default, interface 1/1/8 statistics  (cleared 0h1m42s ago)
================================================================================

Tx Hello Packets      : 25          Rx Hello Packets      : 17
Tx Hello Bytes        : 1664        Rx Hello Bytes        : 1152
Tx DD Packets         : 2           Rx DD Packets         : 2
Tx DD Bytes           : 124         Rx DD Bytes           : 144
Tx LS Request Packets : 1           Rx LS Request Packets : 1
Tx LS Request Bytes   : 68          Rx LS Request Bytes   : 56
Tx LS Update Packets  : 3           Rx LS Update Packets  : 2
Tx LS Update Bytes    : 320         Rx LS Update Bytes    : 288
Tx LS Ack Packets     : 2           Rx LS Ack Packets     : 3
Tx LS Ack Bytes       : 148         Rx LS Ack Bytes       : 192

Total Number of State Changes : 87
Number of LSAs                : 0
LSA Checksum Sum              : 0
Total Transmit Failures       : 0
Total OSPF Packets Discarded  : 170

Reason                        Packets Dropped
--------------------------------------------
Invalid type                  0
Invalid length                0
Invalid checksum              0
Invalid version               0
Bad or unknown source         0
Area mismatch                 0
Self-originated               0
Duplicate router ID           0
Interface standby             0
Total Hello packets dropped   85
  Network Mask mismatch       0
```

```
  Hello interval mismatch     85
  Dead interval mismatch      0
...
```

TIP: The change in statistics can be easier to read when you reset them to 0. If you want to reset the OSPF statistics, you can use the command:

```
clear ip ospf statistics.
```

OSPF Link-Type Point to Point

6. Review if you have any active OSPF adjacencies.

```
show ip ospf neighbors
```

```
sw-agg1(config-if)#    show ip ospf neighbors
VRF : default                          Process : 1
================================================

Total Number of Neighbors : 1

Neighbor ID      Priority  State             Nbr Address       Interface
-------------------------------------------------------------------------
10.254.0.1       1         FULL/BDR          10.254.101.254     1/1/8
```

7. Review the OSPF states in the event log for the **hpe-routing** daemon (background process). This daemon is responsible for routing protocols such as OSPF and BGP in AOS-CX.

```
show event -r -d hpe-routing -n 6
```

```
sw-agg1(config-if)# show event -r -d hpe-routing -n 6
-------------------------------------------------
Event logs from current boot
-------------------------------------------------
2022-10-15T20:32:59.378180+00:00 sw-agg1 hpe-routing[6468]:
Event|2401|LOG_INFO|AMM|1/1|AdjChg: Nbr  rtr ID 10.254.0.1 on IP addr
10.254.101.2( area ID 0.0.0.0): Loading -> Full
2022-10-15T20:32:59.375870+00:00 sw-agg1 hpe-routing[6468]:
Event|2401|LOG_INFO|AMM|1/1|AdjChg: Nbr  rtr ID 10.254.0.1 on IP addr
10.254.101.2( area ID 0.0.0.0): Exchange -> Loading
2022-10-15T20:32:59.372921+00:00 sw-agg1 hpe-routing[6468]:
Event|2401|LOG_INFO|AMM|1/1|AdjChg: Nbr  rtr ID 10.254.0.1 on IP addr
10.254.101.2( area ID 0.0.0.0): Exstart -> Exchange
2022-10-15T20:32:59.372375+00:00 sw-agg1 hpe-routing[6468]:
Event|2401|LOG_INFO|AMM|1/1|AdjChg: Nbr  rtr ID 10.254.0.1 on IP addr
10.254.101.2( area ID 0.0.0.0): Two-way -> Exstart
2022-10-15T20:32:59.372276+00:00 sw-agg1 hpe-routing[6468]:
Event|2401|LOG_INFO|AMM|1/1|AdjChg: Nbr  rtr ID 10.254.0.1 on IP addr
10.254.101.2( area ID 0.0.0.0): Init -> Two-way
2022-10-15T20:32:58.690624+00:00 sw-agg1 hpe-routing[6468]:
Event|2401|LOG_INFO|AMM|1/1|AdjChg: Nbr  rtr ID 10.254.0.1 on IP addr
10.254.101.2( area ID 0.0.0.0): Down -> Init
```

8. Review the OSPF LSDB

```
show ip ospf lsdb
```

```
sw-agg1(config-if)# show ip ospf lsdb
OSPF Router with ID (10.1.0.2) (Process ID 1 VRF default)
============================================================

Router Link State Advertisements (Area 0.0.0.0)
-----------------------------------------------

LSID            ADV Router      Age       Seq#       Checksum       Link Count
-------------------------------------------------------------------------------
10.1.0.2        10.1.0.2        138       0x80000002 0x000074a6     4
10.254.0.1      10.254.0.1      139       0x80000003 0x000058a6     5

Network Link State Advertisements (Area 0.0.0.0)
------------------------------------------------

LSID            ADV Router      Age       Seq#       Checksum
-------------------------------------------------------------
10.254.101.2    10.1.0.2        143       0x80000001 0x00000cb9

AS External Link State Advertisements
-------------------------------------

LSID            ADV Router      Age       Seq#       Checksum
-------------------------------------------------------------
0.0.0.0         10.254.0.1      907       0x80000002 0x00009234
```

Question: Are there any new entries in the LSDB?

Answer: Yes, there is one new router LSA, one new Network LSA and one AS External LSA. This indicates that routing information could now be calculated by OSPF.

9. Review the learned OSPF routes (from OSPF point of view).

```
show ip ospf routes
```

```
sw-agg1(config-if)# show ip ospf routes
Codes: i - Intra-area route, I - Inter-area route
       E1 - External type-1, E2 - External type-2

OSPF Process ID 1 VRF default, Routing Table
--------------------------------------------

Total Number of Routes : 3

10.1.11.0/24       (i) area: 0.0.0.0
     directly attached to interface vlan11, cost 100 distance 110
10.1.12.0/24       (i) area: 0.0.0.0
     directly attached to interface vlan12, cost 100 distance 110
10.254.101.0/24    (i) area: 0.0.0.0
```

```
     directly attached to interface 1/1/8, cost 10 distance 110
```

NOTE: Make sure to enter the correct command, there is also a command **show ip route ospf**. (**route ospf** versus **ospf route**). This command would show you the OSPF routes from the IP routing table.

Question: Although there is an active OSPF adjacency, it seems no OSPF routes are learned from the OSPF core router. What could be causing this issue?

Answer: When only two OSPF routers exist on an Ethernet link, it is common practice to make the OSPF link a point-to-point (P2P) connection. By default, OSPF links over Ethernet are Broadcast (BCAST). Your sw-agg1 believes it is connected to a BCAST network (the network LSA), while the core router has been configured with a point-to-point link. It is important to configure both sides of the OSPF link with the same network type, otherwise OSPF cannot use this link when calculating the best path.

10. Configure the OSPF link on port 1/1/8 as OSPF point-to-point.

```
interface 1/1/8
 ip ospf network point-to-point
 exit
```

```
sw-agg1(config-if)# interface 1/1/8
sw-agg1(config-if)#  ip ospf network point-to-point
sw-agg1(config-if)#  exit
```

11. Verify that OSP routes can now be calculated.

```
show ip ospf routes
```

```
sw-agg1(config)# show ip ospf routes
Codes: i - Intra-area route, I - Inter-area route
       E1 - External type-1, E2 - External type-2

OSPF Process ID 1 VRF default, Routing Table
--------------------------------------------

Total Number of Routes : 7

0.0.0.0/0          (E2)
     via 10.254.101.254 interface 1/1/8, cost 1 distance 110
10.1.11.0/24       (i) area: 0.0.0.0
     directly attached to interface vlan11, cost 100 distance 110
10.1.12.0/24       (i) area: 0.0.0.0
     directly attached to interface vlan12, cost 100 distance 110
10.254.0.1/32      (i) area: 0.0.0.0
     via 10.254.101.254 interface 1/1/8, cost 10 distance 110
10.254.1.0/24      (i) area: 0.0.0.0
     via 10.254.101.254 interface 1/1/8, cost 110 distance 110
10.254.101.0/24    (i) area: 0.0.0.0
     directly attached to interface 1/1/8, cost 10 distance 110
```

```
10.254.102.0/24    (i) area: 0.0.0.0
     via 10.254.101.254 interface 1/1/8, cost 110 distance 110
```

Review the OSPF routes

Based on the new LSAs, sw-agg1 has completed a topology calculation and discovered new routes. The OSPF process will propose these routes to the IP routing table.

12. Review the full IP routing table.

```
show ip route
```

```
sw-agg1(config)# show ip route

Displaying ipv4 routes selected for forwarding

Origin Codes: C - connected, S - static, L - local
              R - RIP, B - BGP, O - OSPF
Type Codes:   E - External BGP, I - Internal BGP, V - VPN, EV - EVPN
              IA - OSPF internal area, E1 - OSPF external type 1
              E2 - OSPF external type 2

VRF: default

Prefix              Nexthop       Interface VRF(egress) Origin/   Distance/    Age
                                                        Type      Metric
-------------------------------------------------------------------------------
-----
0.0.0.0/0           10.254.101.254   1/1/8       -       O/E2      [110/1]
00h:00m:42s
10.1.0.2/32         -                loopback0   -       L         [0/0]        -
10.1.1.0/24         -                vlan1       -       C         [0/0]        -
10.1.1.1/32         -                vlan1       -       L         [0/0]        -
10.1.2.0/24         -                vlan2       -       C         [0/0]        -
10.1.2.2/32         -                vlan2       -       L         [0/0]        -
10.1.3.0/24         -                vlan3       -       C         [0/0]        -
10.1.3.1/32         -                vlan3       -       L         [0/0]        -
10.1.11.0/24        -                vlan11      -       C         [0/0]        -
10.1.11.1/32        -                vlan11      -       L         [0/0]        -
10.1.12.0/24        -                vlan12      -       C         [0/0]        -
10.1.12.1/32        -                vlan12      -       L         [0/0]        -
10.254.0.0/16       10.254.101.254   1/1/8       -       S         [1/0]
00h:35m:07s
10.254.0.1/32       10.254.101.254   1/1/8       -       O         [110/10]
00h:00m:42s
10.254.1.0/24       10.254.101.254   1/1/8       -       O         [110/110]
00h:00m:42s
10.254.101.0/24     -                1/1/8       -       C         [0/0]        -
10.254.101.2/32     -                1/1/8       -       L         [0/0]        -
10.254.102.0/24     10.254.101.254   1/1/8       -       O         [110/110]
00h:00m:42s

Total Route Count : 18
```

Question: How can you identify the OSPF routes in the IP routing table?

Answer: The OSPF routes have an origin code *O*.

Question: What is the administrative distance assigned to the OSPF routes?

Answer: OSPF routes have a distance of 110 by default.

Question: What is the administrative distance assigned to directly connected routes?

Answer: Directly connected routes have a distance of 1.

13. Now lookup the list of OSPF routes in the IP routing table. Then check the output of the IP routes from the OSPF point of view.

```
show ip route ospf
show ip ospf routes
```

```
sw-agg1(config)# show ip ospf routes
Codes: i - Intra-area route, I - Inter-area route
       E1 - External type-1, E2 - External type-2

OSPF Process ID 1 VRF default, Routing Table
--------------------------------------------

Total Number of Routes : 7

0.0.0.0/0          (E2)
     via 10.254.101.254 interface 1/1/8, cost 1 distance 110
10.1.11.0/24       (i) area: 0.0.0.0
     directly attached to interface vlan11, cost 100 distance 110
10.1.12.0/24       (i) area: 0.0.0.0
     directly attached to interface vlan12, cost 100 distance 110
10.254.0.1/32      (i) area: 0.0.0.0
     via 10.254.101.254 interface 1/1/8, cost 10 distance 110
10.254.1.0/24      (i) area: 0.0.0.0
     via 10.254.101.254 interface 1/1/8, cost 110 distance 110
10.254.101.0/24    (i) area: 0.0.0.0
     directly attached to interface 1/1/8, cost 10 distance 110
10.254.102.0/24    (i) area: 0.0.0.0
     via 10.254.101.254 interface 1/1/8, cost 110 distance 110
```

Question: Is there any difference in the list of routes that were calculated by OSPF and the list of routes in the IP routing table?

Answer: Some routes in the IP routing table, such as the directly connected routes for example, have a better administrative distance (1) than the OSPF routes (110). In this case, the directly connected route will make it to the IP routing table instead of the proposed OSPF route.

Most matching route

During the previous lab, you have configured a static aggregate route for the 10.254.0.0/16 datacenter subnets.

14. On sw-agg1, now verify what route will be used when trying to reach 10.254.1.0.

```
show ip route 10.254.1.0
```

```
sw-agg1(config)# show ip route 10.254.1.0

VRF: default

   Prefix            : 10.254.1.0/24                 VRF(egress)       : -
   Nexthop           : 10.254.101.254                Interface         : 1/1/8
   Origin            : ospf                          Type              :
ospf_intra_area
   Distance          : 110                           Metric            : 110
   Age               : 00h:17m:29s                   Tag               : 0
   Encap Type        : -                             Encap Details     : -
```

Question: What route will be used for the 10.254.1.0 destination?

Answer: The most matching route for the 10.254.1.0 subnet is now an OSPF route.

Question: Do you still need the static aggregation route that was made in the previous lab activity?

Answer: No, in this lab setup, OSPF advertises the exact routes from the datacenter. These OSPF routes are now used to reach the datacenter subnets.

15. On **sw-agg1**, review the configured static routes in the running configuration.

```
show running-config | include "ip route"
```

```
sw-agg1(config)# show running-config | include "ip route"
ip route 10.254.0.0/16 10.254.101.254
```

> **NOTE:** Make sure to use double quotes to enclose filter strings that contain a space, like “ip route” in the previous example.

Question: How many static routes are currently defined on sw-agg1?

Answer: There is one static route for 10.254.0.0/16.

Question: Is there a default static route in the configuration?

Answer: No, it was removed during the “static routes” lab activity.

16. Remove the aggregate route for the datacenter subnets.

```
no ip route 10.254.0.0/16 10.254.101.254
```

```
sw-agg1(config)# no ip route 10.254.0.0/16 10.254.101.254
```

17. Verify you can still reach 10.254.1.21.

```
ping 10.254.1.21
```

Default route

The core router uses OSPF to advertise the default route (0.0.0.0/0) to the OSPF neighbors.

18. On **sw-agg1**, check the IP routing table for the default route (0.0.0.0).

```
show ip route
show ip route 0.0.0.0
```

```
sw-agg1(config)# show ip route

Displaying ipv4 routes selected for forwarding

Origin Codes: C - connected, S - static, L - local
              R - RIP, B - BGP, O - OSPF
Type Codes:   E - External BGP, I - Internal BGP, V - VPN, EV - EVPN
              IA - OSPF internal area, E1 - OSPF external type 1
              E2 - OSPF external type 2

VRF: default

Prefix      Nexthop          Interface VRF(egress)   Origin/    Distance/    Age
                                                     Type       Metric
-------------------------------------------------------------------------------
-----
0.0.0.0/0   10.254.101.254   1/1/8       -           O/E2       [110/1]
00h:19m:35s
...
```

```
sw-agg1(config)# show ip route 0.0.0.0

VRF: default

   Prefix           : 0.0.0.0/0
VRF(egress)        : -
   Nexthop          : 10.254.101.254
Interface          : 1/1/8
   Origin           : ospf
Type               : ospf_type2_ext
   Distance         : 110
Metric             : 1
   Age              : 00h:20m:35s
Tag                : 0
   Encap Type       : -                                          Encap
Details      : -
```

Question: Is there a default route present in the IP routing table?

_____

Answer: Yes, the default route is present, even when no local static default route is defined in the sw-agg1 configuration.

Question: What is the advantage of a dynamic static route versus a static default route?

Answer: A dynamic default route can still be advertised using alternate paths in case of a link failure. This can be more challenging to configure using static routes.

19. Verify sw-agg1 can reach the internet. Perform a ping to 8.8.8.8, this should be successful.

```
ping 8.8.8.8
```

```
sw-agg1(config)# ping 8.8.8.8
PING 8.8.8.8 (8.8.8.8) 100(128) bytes of data.
76 bytes from 8.8.8.8: icmp_seq=1 ttl=115 (truncated)
76 bytes from 8.8.8.8: icmp_seq=2 ttl=115 (truncated)
76 bytes from 8.8.8.8: icmp_seq=3 ttl=115 (truncated)
76 bytes from 8.8.8.8: icmp_seq=4 ttl=115 (truncated)
76 bytes from 8.8.8.8: icmp_seq=5 ttl=115 (truncated)

--- 8.8.8.8 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4004ms
rtt min/avg/max/mdev = 8.593/8.849/9.168/0.188 ms
```

Task 4: Configure Passive Interface for Campus Access Subnets

After the successful setup of OSPF routing in the network, your customer contacts you regarding a security concern. The IP interfaces (typically SVI) that connect clients to the network must be advertised using OSPF; this has been done by enabling the interface for OSPF. The side-effect is that OSPF will start sending Hello messages and attempt to find OSPF neighbors on that interface. This means that any 'rogue' OSPF routers on the client subnets would be able to establish an adjacency with the customer OSPF router.

Your senior colleague informs you that it is possible to disable OSPF Hellos on an interface, while keeping the interfaces enabled in the router LSA advertisement. This is known as an OSPF *passive* interface. In this task you will configure OSPF passive interfaces.

Objectives

- Configure an IP interface as passive.
- Configure the OSPF process to use passive interface by default.
- Enable uplink IP interface when Passive interface default is active.

Steps

Passive single interface

1. On PC1, start a Wireshark capture of the lab-facing NIC, wait about 10-15 seconds, and stop the capture.

Question: Are there any OSPF packets seen by the client?

_____

Answer: Yes, by default, OSPF interfaces are active. Any client or rogue router in the client network could attempt to establish an OSPF adjacency with the sw-agg1.

Question: What is the interval of the hello packets?

Answer: By default, OSPF hello packets are sent every 10 seconds on a broadcast network type, such as Ethernet.

> **NOTE:** An example network trace file of OSPF Hello Packets can also be found in the ACAF Student Files on the MGMT PC desktop:
>
> lab03-04-ospf-hello-pc4.pcapng

2. On sw-agg1, make the SVI11 interface an OSPF passive interface.

```
interface vlan 11
 ip ospf passive
 exit
```

```
sw-agg1(config)# interface vlan 11
sw-agg1(config-if-vlan)#  ip ospf passive
sw-agg1(config-if-vlan)#  exit
```

3. Review the brief list of OSPF interfaces, check the Flags column for Active and Passive.

```
show ip ospf interface brief
```

```
sw-agg1(config)# show ip ospf interface brief
OSPF Process ID 1 VRF default
==============================

Total Number of Interfaces: 4

Flags: P - Passive  A - Active

Interface     Area            IP Address/Mask     Cost    State           Status
Flags
-------------------------------------------------------------------------------
--------
1/1/8         0.0.0.0         10.254.101.2/24     10      Point-to-point  Up
A
loopback0     0.0.0.0         10.1.0.2/32         0       Loopback        Up
A
vlan11        0.0.0.0         10.1.11.1/24        100     DR-other        Up
P
vlan12        0.0.0.0         10.1.12.1/24        100     DR              Up
A
```

4. On PC1, start the Wireshark trace again, wait about 10-15 second and stop the trace.

Question: Are there any OSPF packets in the trace?

_____

Answer: No, the sw-agg1 is no longer sending or processing OSPF packets on the interface SVI11.

Passive interface default

When the aggregation switch has many client-facing IP subnets (SVIs), the configuration can be simplified by making all OSPF interfaces passive by default. The network administrator must only enable the uplinks during the initial setup. Any SVIs that are added in the future will automatically be passive.

5. On sw-agg1, remove the SVI 11 OSPF passive configuration.

```
interface vlan 11
 no ip ospf passive
 exit
```

```
sw-agg1(config)# interface vlan 11
sw-agg1(config-if-vlan)#  no ip ospf passive
sw-agg1(config-if-vlan)#  exit
```

6. On sw-agg1, configure passive interface by default.

```
router ospf 1
 passive-interface default
 exit
```

```
sw-agg1(config)# router ospf 1
sw-agg1(config-ospf-1)#  passive-interface default
sw-agg1(config-ospf-1)# exit
```

7. On sw-agg1, verify with the IP OSPF interface details that all SVI interfaces are now passive.

```
show ip ospf interface brief
```

```
sw-agg1(config)# show ip ospf interface brief
OSPF Process ID 1 VRF default
==============================

Total Number of Interfaces: 4

Flags: P - Passive  A - Active

Interface     Area            IP Address/Mask     Cost    State           Status
Flags
-------------------------------------------------------------------------------
--------
1/1/8         0.0.0.0         10.254.101.2/24     10      Point-to-point  Up
P
loopback0     0.0.0.0         10.1.0.2/32         0       Loopback        Up
P
vlan11        0.0.0.0         10.1.11.1/24        100     DR-other        Up
P
vlan12        0.0.0.0         10.1.12.1/24        100     DR-other        Up
P
```

NOTE: You may also use Wireshark on the PC again to verify that no more OSPF packets are advertised on the SVI11.

However, since all interfaces are passive now, the adjacency with the core router has been lost.

8. Verify the OSPF neighbor list.

```
show ip ospf neighbors
```

```
sw-agg1(config)# show ip ospf neighbors
No OSPF neighbor found on VRF default.
```

9. Adjust the configuration by making port 1/1/8 OSPF ‘no passive’.

```
interface 1/1/8
 no ip ospf passive
 exit
```

```
sw-agg1(config)# interface 1/1/8
sw-agg1(config-if)#  no ip ospf passive
sw-agg1(config-if)#  exit
```

10. Verify with the IP OSPF interface details that all interface 1/1/8 is now an active OSPF interface.

```
show ip ospf interface brief
```

```
sw-agg1(config)# show ip ospf interface brief
OSPF Process ID 1 VRF default
==============================

Total Number of Interfaces: 4

Flags: P - Passive  A - Active

Interface     Area            IP Address/Mask     Cost    State           Status
Flags
--------------------------------------------------------------------------------
--------
1/1/8         0.0.0.0         10.254.101.2/24     10      Point-to-point  Up
A
loopback0     0.0.0.0         10.1.0.2/32         0       Loopback        Up
P
vlan11        0.0.0.0         10.1.11.1/24        100     DR-other        Up
P
vlan12        0.0.0.0         10.1.12.1/24        100     DR-other        Up
P
```

11. Verify the OSPF neighbor list contains a FULL adjacency with the core router again.

```
show ip ospf neighbors
```

```
sw-agg1(config)# show ip ospf neighbors
VRF : default                          Process : 1
================================================

Total Number of Neighbors : 1

Neighbor ID      Priority  State             Nbr Address       Interface
------------------------------------------------------------------------
10.254.0.1       n/a       FULL              10.254.101.254     1/1/8
```

Task 5: Practice OSPF Configuration on sw-agg2

In this task you will add sw-agg2 to the OSPF topology. You will practice the OSPF configuration by applying it on agg2. Note that passive interfaces ensure no OSPF peering is possible on endpoint subnets. This also means that peering is not possible between agg1 and agg2 on any of the endpoint subnets. The uplink should be the only interface that is active for OSPF peering.

If the sw-agg1 or sw-agg2 uplink to the core router fails, they would not have an alternative path to reach the core router. To solve this, a direct IP (routed) link is required between the two aggregation switches.
In the lab topology, SVI2 (VLAN 2) has been configured on the LAG between the aggregation switches for this purpose.

It will be used as the routed link between the aggregation switches and, if an uplink fails, it can be used as an alternate path via the peer aggregation switch to reach the core router.

Complete these configuration steps:

On sw-agg2:
- Verify loopback 0 with ip 10.1.0.3/32
- Define OSPF process 1 with rid ip of lo0
- Activate uplink 1/1/8 with the correct OSPF interface timers and OSPF network type
- Activate link SVI2 to agg1 with the adjusted OSPF timers and OSPF network type.
- Ensure all future interfaces are passive by default
- Verify uplink adjacency with core router and sw-agg1 is FULL

Adjust the configuration on sw-agg1 as needed to support the above requirements.

Objectives

- Practice the OSPF configuration.

Solution

Sw-agg2

```
# sw-agg2 solution

router ospf 1
 router-id 10.1.0.3
 passive-interface default
 area 0.0.0.0

interface 1/1/8
 ip ospf 1 area 0.0.0.0
 no ip ospf passive
 ip ospf hello-interval 1
 ip ospf dead-interval 4
 ip ospf network point-to-point

interface loopback 0
 ip ospf 1 area 0.0.0.0

interface vlan2
 ip ospf 1 area 0.0.0.0
 ip ospf network point-to-point
 no ip ospf passive
 exit
```

Sw-agg1

```
# sw-agg1 solution

interface vlan2
 ip ospf 1 area 0.0.0.0
 ip ospf network point-to-point
 no ip ospf passive
 exit
```

Validation steps

```
show ip ospf neighbors
show ip ospf interface brief
show ip ospf route
```

1. Verify sw-agg1 has two OSPF peers (sw-agg2 and rtr-core1).
2. Verify sw-agg2 has two OSPF peers (sw-agg1 and rtr-core1).
3. Verify sw-agg2 has OSPF routes with next-hop to both sw-agg1 and rtr-core1.

Troubleshooting steps

In case you would have issues, check your OSPF configuration and the OSPF interface configuration and compare these settings with the peer device (sw-agg1 or rtr-core1).

```
show run ospf
show run interface vlan2
```

IMPORTANT: Once you have completed and validated the OSPF configuration , save the configuration on all the switches!

```
write memory
```

Optional Task 6: Verify Routing Failover with OSPF

Now that you have completed the OSPF setup, the customer would like to see a demonstration of how OSPF handles a link failure and re-routes traffic in the network. In this task you will explore how OSPF responds to a link failure in the network.

Objectives

- Understand how OSPF handles link failures

Steps

1. On PC1, open a command prompt (**cmd**) and start a continuous ping to 10.254.1.21

```
ping 10.254.1.21 -t
```

2. Start a second command prompt and perform a traceroute to 10.254.1.21. Take note of the IP addresses in the trace output.

> **TIP:** On the Windows client, use **tracert -d 10.254.1.21** (-d = do not resolve host names) to speed up the traceroute output.

Trace route IP list:

_____

_____

_____

3. On sw-agg1, check the current number SPF calculations in the **show ip ospf** command output.

```
show ip ospf
```

```
sw-agg1(config)# show ip ospf
VRF : default                          Process  : 1
-----------------------------------------------------

RouterID               : 10.1.0.2          OSPFv2                 : Enabled
BFD                    : Disabled          SPF Start Interval     : 200   ms
SPF Hold Interval      : 1000  ms          SPF Max Wait Interval  : 5000  ms
LSA Start Time         : 5000  ms          LSA Hold Time          : 0     ms
LSA Max Wait Time      : 0     ms          LSA Arrival            : 1000  ms
External LSAs          : 1                 Checksum Sum           : 36917
ECMP                   : 4                 Reference Bandwidth    : 100000 Mbps
Area Border            : false             AS Border              : false
GR Status              : Enabled           GR Interval            : 120 sec
GR State               : inactive          GR Exit Status         : none
GR Helper              : Enabled           GR Strict LSA Check    : Enabled
GR Ignore Lost I/F     : Disabled
Summary address:
```

```
Area      Total     Active
-----------------------------
Normal     1         1
Stub       0         0
NSSA       0         0

Area  : 0.0.0.0
----------------
Area Type              : Normal         Status               : Active
Total Interfaces       : 5              Active Interfaces    : 2
Passive Interfaces     : 3              Loopback Interfaces  : 1
SPF Calculation Count  : 20
Area ranges     :
Number of LSAs         : 3              Checksum Sum         : 126962
```

4. Shutdown the uplink port to the core router (1/1/8).

```
interface 1/1/8
 shutdown
```

```
sw-agg1(config)# interface 1/1/8
sw-agg1(config-if)#  shutdown
```

5. On PC1, verify that the ping continued.

Question: Do you notice any difference in the ping output?

_____

Answer: Yes, the TTL has changed because the path now requires one additional hop (via sw-agg2).

6. In the second command prompt, run the traceroute again. Compare this with the previous output.

Question: What do you observe?

_____

Answer: The path takes sw-agg1, then the routed connection to sw-agg2 (via SVI 2), and finally the connection to the core router.

7. On sw-agg1, review the SPF count. This should have increased compared to the previous check.

```
show ip ospf
```

```
sw-agg1(config-if)# show ip ospf
VRF : default                          Process  : 1
-------------------------------------------------

RouterID               : 10.1.0.2           OSPFv2                 : Enabled
...
Area  : 0.0.0.0
----------------
Area Type              : Normal         Status               : Active
Total Interfaces       : 5              Active Interfaces    : 2
```

```
Passive Interfaces     : 3              Loopback Interfaces   : 1
SPF Calculation Count  : 22
Area ranges      :
Number of LSAs         : 3              Checksum Sum          : 45623
```

8. On sw-agg1, enable the port to the core router again (1/1/8).

```
interface 1/1/8
 no shutdown
 exit
```

```
sw-agg1(config-if)# interface 1/1/8
sw-agg1(config-if)#  no shutdown
sw-agg1(config-if)#  exit
```

9. Cleanup: On the PC close the command prompt windows.

Optional Task 7: Verify Equal Cost Multiple Paths Routing with OSPF

In the network setup, there is one link from the core router to each aggregation switch. For traffic that comes from the datacenter, two paths seem to be available to reach the client subnets that are connected to the aggregation switches. Your customer would like to understand if both links will be used and how this process works.

You learn from your colleague that both links can indeed be used thanks to a system that is known as Equal Cost Multiple Paths (ECMP). In this task, you will observe the operation of ECMP with regards to OSPF. ECMP is by default enabled. When multiple paths with the same cost exist for a destination network, all the paths will be placed in the routing table. The routing switch will use a load-balancing algorithm to distribute the load over these multiple paths.

Objectives

- Understand the ECMP operation
- Configure ECMP

Steps

1. Open an SSH connection to the core router.
2. In the IP routing table, lookup the entries for the 10.1.2.0 network.

```
show ip route
show ip route 10.1.2.0
```

```
rtr-core1(config)# show ip route

Displaying ipv4 routes selected for forwarding

Origin Codes: C - connected, S - static, L - local
              R - RIP, B - BGP, O - OSPF
Type Codes:   E - External BGP, I - Internal BGP, V - VPN, EV - EVPN
              IA - OSPF internal area, E1 - OSPF external type 1
              E2 - OSPF external type 2

VRF: default

Prefix               Nexthop           Interface     VRF(egress)   Origin/  Distance/
Age
                                                                     Type     Metric
-------------------------------------------------------------------------------------
-----
0.0.0.0/0      10.254.1.253     1/1/9         -         S         [1/0]
01h:14m:34s
10.0.0.0/8     -                blackhole     -         S         [1/0]
01h:14m:34s
10.1.0.0/16    10.254.101.2     1/1/1         -         S         [1/0]
01h:04m:04s
10.1.0.2/32    10.254.101.2     1/1/1         -         O         [110/100]
00h:01m:38s
```

```
10.1.0.3/32     10.254.102.3     1/1/2          -          O          [110/100]
00h:03m:21s
10.1.2.0/24     10.254.101.2     1/1/1          -          O          [110/200]
00h:01m:38s
                10.254.102.3     1/1/2          -                     [110/200]
00h:01m:38s
10.1.3.0/24     10.254.101.2     1/1/1          -          S          [255/0]
01h:14m:34s
10.1.11.0/24    10.254.101.2     1/1/1          -          O          [110/200]
00h:01m:38s
...
```

```
rtr-core1(config)# show ip route 10.1.2.0

VRF: default

   Prefix            : 10.1.2.0/24                 VRF(egress)       : -
   Nexthop           : 10.254.101.2                Interface         : 1/1/1
   Origin            : ospf                        Type              :
ospf_intra_area
   Distance          : 110                         Metric            : 200
   Age               : 00h:04m:50s                 Tag               : 0
   Encap Type        : -                           Encap Details     : -

   Prefix            : 10.1.2.0/24                 VRF(egress)       : -
   Nexthop           : 10.254.102.3                Interface         : 1/1/2
   Origin            : ospf                        Type              :
ospf_intra_area
   Distance          : 110                         Metric            : 200
   Age               : 00h:04m:50s                 Tag               : 0
   Encap Type        : -                           Encap Details     : -
```

Question: How many routes are available for the 10.1.2.0/24 network?

_____

Answer: The core router has 2 ECMP paths for the 10.1.2.0/24 network. The first path via sw-agg1, the second path via sw-agg2.

3. Review the OSPF default maximum paths for ECMP.

```
show ip ospf
```

```
rtr-core1(config)# show ip ospf
VRF : default                          Process  : 1
-----------------------------------------------------

RouterID              : 10.254.0.1          OSPFv2                 : Enabled
BFD                   : Disabled            SPF Start Interval     : 200   ms
SPF Hold Interval     : 1000  ms            SPF Max Wait Interval  : 5000  ms
LSA Start Time        : 5000  ms            LSA Hold Time          : 0     ms
LSA Max Wait Time     : 0     ms            LSA Arrival            : 1000  ms
External LSAs         : 1                   Checksum Sum           : 36917
ECMP                  : 4                   Reference Bandwidth    : 100000 Mbps
```

```
Area Border            : false              AS Border              : true
GR Status              : Enabled            GR Interval            : 120 sec
GR State               : inactive           GR Exit Status         : none
GR Helper              : Enabled            GR Strict LSA Check    : Enabled
GR Ignore Lost I/F     : Disabled
...
```

4. Adjust the OSPF maximum paths to 1.

```
router ospf 1
 maximum-paths 1
 exit
```

```
rtr-core1(config)# router ospf 1
rtr-core1(config-ospf-1)#  maximum-paths 1
rtr-core1(config-ospf-1)#  exit
```

5. In the IP routing table, lookup the route entries for the 10.1.2.0 network.

```
show ip route
```

Question: How many routes are available?

Answer: Only one route is present in the IP routing table.

6. Review the 10.1.2.0 network in the OSPF routing table.

```
show ip ospf route
```

```
rtr-core1(config)# show ip route

Displaying ipv4 routes selected for forwarding

Origin Codes: C - connected, S - static, L - local
              R - RIP, B - BGP, O - OSPF
Type Codes:   E - External BGP, I - Internal BGP, V - VPN, EV - EVPN
              IA - OSPF internal area, E1 - OSPF external type 1
              E2 - OSPF external type 2

VRF: default

Prefix         Nexthop      Interface    VRF(egress)   Origin/  Distance/    Age
                                                       Type     Metric
-------------------------------------------------------------------------------
---------
0.0.0.0/0      10.254.1.253     1/1/9          -       S         [1/0]
01h:20m:01s
10.0.0.0/8     -                blackhole      -       S         [1/0]
01h:20m:01s
10.1.0.0/16    10.254.101.2     1/1/1          -       S         [1/0]
01h:09m:31s
10.1.0.2/32    10.254.101.2     1/1/1          -       O         [110/100]
00h:07m:05s
10.1.0.3/32    10.254.102.3     1/1/2          -       O         [110/100]
00h:08m:48s
```

```
10.1.2.0/24     10.254.102.3      1/1/2          -       O          [110/200]
00h:00m:21s
10.1.3.0/24     10.254.101.2      1/1/1          -       S          [255/0]
01h:20m:01s
```

Question: How many routes are available?

Answer: Only one route has been calculated by OSPF as the best route.

7. Restore the OSPF configuration by removing the maximum paths command.

```
router ospf 1
 no maximum-paths
 exit
```

```
rtr-core1(config)# router ospf 1
rtr-core1(config-ospf-1)#  no maximum-paths
rtr-core1(config-ospf-1)#  exit
```

8. Verify that both routes are available again in the routing table.

```
show ip route
```

NOTE: If you want to see the default values for any setting or command, you can use the command **show running-configuration all**. When you want to show the default OSPF commands on an interface, you can use **show running-config all | begin 50 1/1/8 | include ospf**. This will show the running-configuration beginning at the line **1/1/8** and showing the previous **50 lines**. Out of those lines, only the lines that include the string **ospf** will be shown.

```
rtr-core1(config)# show running-config all | begin 50 1/1/8 | include
ospf
    ip ospf 1 area 0.0.0.0
    ip ospf dead-interval 4
    ip ospf hello-interval 1
    ip ospf retransmit-interval 5
    ip ospf transit-delay 1
    ip ospf priority 1
    ip ospf cost 100
    ip ospf network point-to-point
    ip ospf 1 area 0.0.0.0
    ip ospf dead-interval 4
    ip ospf hello-interval 1
    ip ospf retransmit-interval 5
    ip ospf transit-delay 1
    ip ospf priority 1
    ip ospf cost 100
    ip ospf network point-to-point
```

You have completed this Lab!

Lab 04.01 Spanning-Tree

Overview

Your customer informs you that they have experienced some network loops in the past. They would like to understand what options are available to deal with network loops. After contact with your senior colleague, you learn that the STP and the loop protect protocol can be used for network loop detection.

STP can also be used to select the best path in a redundant layer 2 topology. However, in most networks today, this function will be performed by a VSX cluster and the core or aggregation and MCLAG connections to the edge switches. In this type of setup, the role of the STP is reduced to a loop detection protocol.

Since STP is considered a core network administrator competency, this lab will still provide you a brief introduction to STP layer 2 best path selection and path failover. In this lab you will explore the operation and basic configuration of the STP and loop protection.

Objectives

After completing this lab, you will be able to:

- Review the STP default settings
- Understand STP path cost and failover
- Configure STP access security with BPDU-Guard
- Configure and verify loop protection

Task 1: Explore STP costs and path failover

In this task you will explore the default settings of Spanning-Tree. You will explore the root path cost and Topology Change Notifications (TCNs). To get a redundant layer 2 path in the network, the link between sw-edge1 and sw-agg2 will be enabled. This will result in a loop in the network that will be resolved by STP.

Note that this is only an example of a looped topology. In a later lab activity, you will configure a LAG between the access switch and the VSX aggregation pair of switches to achieve an active-active forwarding topology.

Objectives

- Explore default STP settings
- Review the Root Path Cost
- Review the Root Port

Steps

Verify Spanning-Tree on the Root Bridge

For this lab, STP should be enabled on all switches. This is the default state on the 6xxx series switches.

In the next steps you will review the Spanning-Tree protocol status on your lab switches.

1. On sw-agg1, review the spanning-tree status. STP should be enabled. Take note of the bridge ID priority and MAC-address.

```
show spanning-tree
```

```
sw-agg1(config)# show spanning-tree
Spanning tree status      : Enabled Protocol: MSTP

MST0
  Root ID    Priority    : 0
             MAC-Address: b8:d4:e7:da:80:00
             This bridge is the root
             Hello time(in seconds):2  Max Age(in seconds):20
             Forward Delay(in seconds):15

  Bridge ID  Priority    : 0
             MAC-Address: b8:d4:e7:da:80:00
             Hello time(in seconds):2  Max Age(in seconds):20
             Forward Delay(in seconds):15

Port    Role           State      Cost  Priority   Type  BPDU-Tx  BPDU-Rx  TCN-Tx
TCN-Rx
------------ -------------- ---------- -------------- ---------- ---------------- ---
-----
1/1/1   Designated     Forwarding 2000  128        P2P   75994    14       10      14
1/1/2   Designated     Forwarding 2000  128        P2P   64085    9        6       6
lag256  Designated     Forwarding 800   64         P2P   76128    2        22      2
```

```
Number of topology changes    : 19
Last topology change occurred : 128119 seconds ago
```

Question: What is the priority configured on sw-agg1?

_____

Answer: Sw-agg1 is the target Root Bridge. In the lab configuration it was configured with priority 0, the lowest priority value. The lowest value is the best value for STP.

Verify Spanning-Tree status and Root Port on the Access Switches

2. On sw-edge1, lookup the switch base system MAC Address.

```
show system
```

Example output, our output may be different.

```
sw-edge1# show system
Hostname           : sw-edge1
System Description : FL.10.09.1040
System Contact     :
System Location    :

Vendor             : Aruba
Product Name       : JL666A 6300F 24G CL4 PoE 4SFP56 Sw
Chassis Serial Nbr : SG09KN500B
Base MAC Address   : 64e881-dd8180
ArubaOS-CX Version : FL.10.09.1040

Time Zone          : US/Michigan

Up Time            : 1 day, 18 hours, 34 minutes
CPU Util (%)       : 1
Memory Usage (%)   : 22
```

3. Review the spanning-tree status. STP should be enabled. Take note of the bridge ID priority and MAC-address.

```
show spanning-tree
```

```
sw-edge1# show spanning-tree
Spanning tree status      : Enabled Protocol: MSTP

MST0
  Root ID    Priority    : 0
             MAC-Address: b8:d4:e7:da:80:00
             Hello time(in seconds):2  Max Age(in seconds):20
             Forward Delay(in seconds):15

  Bridge ID  Priority    : 32768
             MAC-Address: 64:e8:81:dd:81:80
             Hello time(in seconds):2  Max Age(in seconds):20
             Forward Delay(in seconds):15
...
```

Question: Is there a relation between the Bridge ID MAC and the system MAC?

Answer: Yes, they are the same.

Question: What is the (default) bridge priority on sw-edge1?

Answer: The default priority for spanning tree is 32768.

Question: Based on the bridge IDs you have seen on sw-agg1 and the edge switches, which switch will be elected by STP as the Root Bridge?

Answer: Sw-agg1 has the lowest bridge ID and will be elected as the Root Bridge.

Verify the Root Path Cost

In the next steps you will explore the path cost used by Spanning Tree to measure the cost to reach the Root Bridge. This is described as the Root Path Cost.

4. Use the **show spanning-tree mst** command to take note of the Root Path Cost.

```
show spanning-tree mst
```

```
sw-edge1# show spanning-tree mst
#### MST0
Vlans mapped:  1-4094
Bridge         Address:64:e8:81:dd:81:80    priority:32768
Regional Root
Operational    Hello time(in seconds): 2  Forward delay(in seconds):15  Max-age(in
seconds):20  txHoldCount(in pps): 6
Configured     Hello time(in seconds): 2  Forward delay(in seconds):15  Max-age(in
seconds):20  Max-Hops:20
Root           Address:b8:d4:e7:da:80:00  Priority:0
               Port:1/1/27                Path cost:2000
Regional Root  Address:64:e8:81:dd:81:80  Priority:32768
               Internal cost:0            Rem Hops:20
...
```

Question: What is the current Root Path Cost?

Answer: The Root Path Cost is 2000.

Question: What type of link will result in a cost of 2000?

Answer: A single 10Gbps link has a default cost value of 2000.

Question: How does the STP cost 2000 link to 10Gbps?

Answer: As part of the STP standard, 2000000 has been set as the reference cost for a 10Mbps connection. Other speeds are based on this reference cost:

Speed	STP Link Cost

10Mbps	2000000
100Mbps	200000
1Gbps	20000
10Gbps	2000
100Gbps	200

Review the port list of the 'show spanning-tree mst' command output.
Question: For port 1/1/27, what is the Role, State and Cost?

Answer: Port 1/1/27 is the Root port, it is Forwarding, and the cost is 2000.

5. Review port 1/1/27 in the brief interface list.

```
show interface 1/1/27 brief
```

```
sw-edge1# show int 1/1/27 brief
--------------------------------------------------------------------------------
-------------------------
Port      Native  Mode   Type       Enabled Status  Reason                Speed
Description
          VLAN                                                             (Mb/s)
--------------------------------------------------------------------------------
-------------------------
1/1/27    1       trunk  SFP+DAC1   yes     up                            10000
sw-agg1
```

Question: What is the speed for the port 1/1/27?

Answer: The speed is 10000 (Mb/s). This represents 10Gbps and results in the STP cost of 2000.

6. On sw-edge2, review the spanning-tree status. STP should be enabled. Take note of the bridge ID priority and MAC-address.

```
show spanning-tree
```

```
sw-edge2# show spanning-tree
Spanning tree status      : Enabled Protocol: MSTP

MST0
  Root ID    Priority    : 0
             MAC-Address: b8:d4:e7:da:80:00
             Hello time(in seconds):2  Max Age(in seconds):20
             Forward Delay(in seconds):15

  Bridge ID  Priority  : 32768
             MAC-Address: 8c:85:c1:49:20:c0
             Hello time(in seconds):2  Max Age(in seconds):20
             Forward Delay(in seconds):15
...
```

Verify the Root Path Cost

7. Use the **show spanning-tree mst** command to take note of the Root Path Cost. This should also be 2000, based on a single 10Gbps link to the sw-agg1.

```
show spanning-tree mst
```

```
sw-edge2# show spanning-tree mst
#### MST0
Vlans mapped:  1-4094
Bridge         Address:8c:85:c1:49:20:c0    priority:32768
Regional Root
Operational    Hello time(in seconds): 2  Forward delay(in seconds):15  Max-age(in
seconds):20  txHoldCount(in pps): 6
Configured     Hello time(in seconds): 2  Forward delay(in seconds):15  Max-age(in
seconds):20  Max-Hops:20
Root           Address:b8:d4:e7:da:80:00  Priority:0
               Port:1/1/27                Path cost:2000
Regional Root  Address:8c:85:c1:49:20:c0  Priority:32768
               Internal cost:0            Rem Hops:20
...
```

Investigate Interface Cost for LAG ports

Each interface link speed has its own assigned cost. In the next steps you will explore what the impact of the LAG speed is with regards to the Spanning Tree cost.

8. On sw-agg2, review the spanning-tree status. STP should be enabled. Take note of the bridge ID priority and MAC-address.

```
show spanning-tree
```

```
sw-agg2# show spanning-tree
Spanning tree status      : Enabled Protocol: MSTP

MST0
  Root ID    Priority    : 0
             MAC-Address: b8:d4:e7:da:80:00
             Hello time(in seconds):2  Max Age(in seconds):20
             Forward Delay(in seconds):15

  Bridge ID  Priority  : 4096
             MAC-Address: b8:d4:e7:d9:e5:00
             Hello time(in seconds):2  Max Age(in seconds):20
             Forward Delay(in seconds):15
```

9. On sw-agg2, use the **show spanning-tree mst** command to see the Root Path Cost.

```
show spanning-tree mst
```

```
sw-agg2# show spanning-tree mst
#### MST0
Vlans mapped:  1-4040
```

```
Bridge          Address:b8:d4:e7:d9:e5:00    priority:4096
Regional Root
Operational     Hello time(in seconds): 2  Forward delay(in seconds):15  Max-age(in
seconds):20  txHoldCount(in pps): 6
Configured      Hello time(in seconds): 2  Forward delay(in seconds):15  Max-age(in
seconds):20  Max-Hops:20
Root            Address:b8:d4:e7:da:80:00  Priority:0
                Port:lag256                Path cost:800
Regional Root   Address:b8:d4:e7:d9:e5:00  Priority:4096
                Internal cost:0            Rem Hops:20
```

Question: What is the current Root Port?

Answer: The current Root Port is lag256.

Question: What is the current Root Path Cost?

Answer: The current Root Path Cost is 800.

Question: What is the cost for port lag256?

Answer: Port lag256 has a cost of 800.

10. Review the interface brief output, filter ‘up’ in the output. Pay attention to the speed.

```
show interface brief | include up
```

```
sw-agg2# show interface brief | include up
1/1/8     --      routed SFP+DAC3       yes     up                          10000
rtr-core1-1/1/2
1/1/46    1       trunk  SFP28DAC0.65   yes     up                          25000
--
1/1/47    1       trunk  SFP28DAC0.65   yes     up                          25000
--
loopback0 --      routed --             yes     up                          --
--
vlan2     --              --            yes     up                          --
--
lag256    1       trunk  --             yes     up      --                  50000
--
```

Question: What is the speed of port 1/1/46 and 1/1/47?

Answer: Port 1/1/46 and 1/1/47 are connected at 25Gbps.

Question: What would be the standard STP cost for a 25Gbps link?

Answer: 25Gbps is ‘4 times slower’ than a 100Gbps link, so 200x4=800.

Question: What was the Root Path Cost on sw-agg2?

Answer: The Root Path Cost is 800. This is the cost of a 25Gbps link.

Question: What is the speed of the LAG256?

Answer: The interface speed is 50Gbps, based on 2 member ports of 25Gbps.

Question: What could you expect the LAG STP cost to be, based on a speed of 50 Gbps?

Answer: 50Gbps is '2 times slower' than a 100Gbps, so 200x2=400. However, the LAG is using an STP cost of 800.

Question: Then why is the LAG cost 800?

Answer: This is a vendor choice. For AOS-CX, the choice was made to keep the link speed of a LAG to the individual member port speed. This ensures that a link down event for a LAG member will not impact STP, since the STP port speed did not change.

Conclusion: On AOS-CX, the STP cost of a LAG is based on the member port speed, not the total reported speed of the LAG to ensure STP stability during LAG member port changes.

Add redundant link to the network

In the next steps you will activate a new link in the network. This will result in a loop in the network topology. You will enable the link between sw-edge1 and sw-agg2.

11. On **sw-agg2**, enable the interface 1/1/1. It connects to sw-edge1.

```
interface 1/1/1
 no shutdown
 exit
```

```
sw-agg2(config)# interface 1/1/1
sw-agg2(config-if)# no shutdown
sw-agg2(config-if)# exit
```

12. On **sw-edge1**, enable the interface 1/1/28. It connects to sw-agg2.

```
interface 1/1/28
 description sw-agg2
 no shutdown
 exit
```

```
sw-edge1(config)# interface 1/1/28
sw-edge1(config-if)#  description sw-agg2
sw-edge1(config-if)#  no shutdown
sw-edge1(config-if)#  exit
```

13. Verify the interface 1/1/28 is up and connects to port 1/1/1 on sw-agg2.

```
show interface 1/1/28 brief
show lldp neighbor
```

```
sw-edge1(config)# show interface 1/1/28 brief
-------------------------------------------------------------------------------
-------------------------
Port      Native  Mode   Type           Enabled Status  Reason                 Speed
Description
          VLAN                                                                 (Mb/s)
-------------------------------------------------------------------------------
-------------------------
1/1/28    1       access SFP+DAC1       yes     up                             10000
sw-agg2
```

```
sw-edge1(config)# show lldp neighbor-info

LLDP Neighbor Information
=========================

Total Neighbor Entries          : 7
Total Neighbor Entries Deleted  : 4
Total Neighbor Entries Dropped  : 0
Total Neighbor Entries Aged-Out : 4

LOCAL-PORT  CHASSIS-ID         PORT-ID                PORT-DESC   TTL      SYS-NAME
-------------------------------------------------------------------------------
-----
...
1/1/27      b8:d4:e7:da:80:00  1/1/1                  1/1/1       120      sw-agg1
1/1/28      b8:d4:e7:d9:e5:00  1/1/1                  1/1/1       120      sw-agg2
...
```

Review the Spanning-Tree for the redundant link

14. On **sw-edge1**, review the spanning-tree mst status. Pay attention to the State column. This is the forwarding mode.

```
show spanning-tree mst
```

```
sw-edge1(config)# show spanning-tree mst
...
Port           Role           State      Cost Priority Type      BPDU-Tx Rx  TCN-Tx
Rx
-------------- -------------- ---------- ---------- ---------- ---------------- -----
-----
...
1/1/24         Disabled       Down       20000  128    P2P          0   0      0
0
1/1/27         Root           Forwarding 2000   128    P2P  Bound   16  76904  14
12
1/1/28         Alternate      Blocking   2000   128    P2P  Bound   2   63     0
2
...
```

Question: Is there any interface in Blocking state?

_____

Answer: Yes, port 1/1/28 is shown with a **Blocking** state.

Task 2: Test STP Path failover

In this task you will test the STP failover on the access switch sw-edge1 in case of an uplink failure. The uplink to sw-agg1 will be disabled. As a result, STP should enable the path to sw-agg2.

Objectives

- Verify STP Path Failover

Steps

1. On **sw-edge1**, check the Spanning-Tree MST output and take note of the current Root Path Cost.

```
show spanning-tree mst
```

Example output, your output may be different.

```
sw-edge1(config)# show spanning-tree mst
#### MST0
Vlans mapped:  1-4094
Bridge         Address:64:e8:81:dd:81:80    priority:32768
Regional Root
Operational    Hello time(in seconds): 2  Forward delay(in seconds):15  Max-age(in
seconds):20  txHoldCount(in pps): 6
Configured     Hello time(in seconds): 2  Forward delay(in seconds):15  Max-age(in
seconds):20  Max-Hops:20
Root           Address:b8:d4:e7:da:80:00  Priority:0
               Port:1/1/27                Path cost:2000
Regional Root  Address:64:e8:81:dd:81:80  Priority:32768
               Internal cost:0            Rem Hops:20
...
Topology change flag          : False
Number of topology changes    : 45
Last topology change occurred : 50 seconds ago
```

Question: What is the current root path cost?

Answer: The current root path cost is 2000. By default, this represents a single 10Gbps connection to reach the Root Bridge.

2. Take note of the current number of Topology Changes.

Question: What is the current Root Port?

Answer: The current Root Port is port 1/1/27, which connects to sw-agg1.

3. On **PC1**, open a command prompt and start a continuous ping to server IP 10.254.1.21.

```
ping 10.254.1.21 -t
```

4. On **sw-edge1**, shutdown the uplink to sw-agg1, this is port 1/1/27.

```
interface 1/1/27
 shutdown
 exit
```

5. On **sw-edge1**, review the spanning-tree mst output.

```
show spanning-tree mst
```

Example output, your output may be different.

```
sw-edge1(config)# show spanning-tree mst
#### MST0
Vlans mapped:  1-4094
Bridge         Address:64:e8:81:dd:81:80    priority:32768
Regional Root
Operational    Hello time(in seconds): 2  Forward delay(in seconds):15  Max-age(in
seconds):20  txHoldCount(in pps): 6
Configured     Hello time(in seconds): 2  Forward delay(in seconds):15  Max-age(in
seconds):20  Max-Hops:20
Root           Address:b8:d4:e7:da:80:00  Priority:0
               Port:1/1/28                Path cost:2800
Regional Root  Address:64:e8:81:dd:81:80  Priority:32768
               Internal cost:0            Rem Hops:20
...
Topology change flag          : False
Number of topology changes    : 47
Last topology change occurred : 2 seconds ago
```

Question: What port is the Root Port now?

Answer: Port 1/1/28.

Question: What is the Root Path Cost now?

Answer: Previously it was 2000 (single 10Gbps connection), now it shows 2800 (10Gbps connection to reach sw-agg2, then the 25Gbps connection to reach sw-agg1).

Question: What is the number of topology changes? Did this change from the previous output?

Answer: The number of topology changes has now increased, due to the uplink port change.

6. On PC1, check your ping status.

Question: Is the ping successful?

Answer: No, although it seems from the previous output that STP successfully calculated a new best path to the Root Bridge, the user traffic has stopped.

Let's investigate. What could be the issue?

7. On **sw-edge1**, review the MAC address table for VLAN 11. You should expect to see a client MAC address and a MAC address for the default gateway (sw-agg1) on the uplink port.

```
show mac-address-table vlan 11
```

Example output, your output may be different.

```
sw-edge1(config)# show mac-address-table vlan 11
MAC age-time            : 300 seconds
Number of MAC addresses : 1

MAC Address          VLAN     Type                      Port
-------------------------------------------------------------
00:50:56:b1:cd:26    11       dynamic                   1/1/1
```

Question: Do you see a MAC address for the client and the default gateway?

Answer: No, only a MAC address on port 1/1/1 for the client is shown. This indicates there is something wrong with the uplink port VLAN.

8. Review the VLAN 11 port membership.

```
show vlan 11
```

```
sw-edge1(config)# show vlan 11

------------------------------------------------------------------------------------------
----------------------------
VLAN  Name                                Status  Reason                  Type
Interfaces
------------------------------------------------------------------------------------------
----------------------------
11    v11-employee                        up      ok                      static
1/1/1,1/1/27,lag50
```

Question: What are the port members of VLAN 11?

Answer: Ports 1/1/1,1/1/27 and LAG50.

Question: What port is missing from this list?

Answer: Port 1/1/28 is now the uplink port, but it is not member of the VLAN 11.

9. Review the port 1/1/28 configuration.

```
show running-config interface 1/1/28
```

```
sw-edge1(config)# show running-config interface 1/1/28
interface 1/1/28
    no shutdown
    description sw-agg2
    no routing
    vlan access 1
```

```
    exit
```

Question: What do you observe?

Answer: The port is an access port in VLAN 1. This is still the default configuration.

10. Adjust the configuration. Make port 1/1/28 a VLAN trunk that allows all VLANs.

```
interface 1/1/28
 vlan trunk allowed all
 exit
```

11. Verify your configuration. Check the VLAN 11 port membership.

```
show vlan 11
```

```
sw-edge1(config)# show vlan 11

------------------------------------------------------------------------------
----------------------------
VLAN  Name                                Status  Reason                  Type
Interfaces
------------------------------------------------------------------------------
----------------------------
11    v11-employee                        up      ok                      static
1/1/1,1/1/27-1/1/28,lag50
```

12. Check the MAC-address table for VLAN 11 again.

```
show mac-address-table vlan 11
```

```
sw-edge1(config)# show mac-address-table vlan 11
MAC age-time            : 300 seconds
Number of MAC addresses : 2

MAC Address          VLAN     Type                      Port
-------------------------------------------------------------
00:50:56:b1:cd:26    11       dynamic                   1/1/1
b8:d4:e7:da:80:00    11       dynamic                   1/1/28
```

Question: Were any MAC addresses learned on the uplink port?

Answer: Yes, a MAC address has been learned on the port 1/1/28 now.

13. On PC1, verify that the ping has resumed.
14. On **sw-edge1**, repeat the uplink failure to confirm the STP failover. Enable the original uplink again.

```
interface 1/1/27
 no shutdown
```

15. Wait about 10 seconds, then shutdown the link again.

```
 shutdown
```

16. On **PC1**, verify that the ping continues.

NOTE: It is expected to see the loss of a ping packet during a topology change.

This demonstrates the STP failover.

17. On **sw-edge1**, restore the uplink connection by enabling port 1/1/27.

```
interface 1/1/27
 no shutdown
 exit
```

18. On **sw-edge1**, review the Spanning-Tree MST output.

```
show spanning-tree mst
```

```
sw-edge1(config-if)# show spanning-tree mst
#### MST0
Vlans mapped:  1-4094
Bridge         Address:64:e8:81:dd:81:80    priority:32768
Regional Root
Operational    Hello time(in seconds): 2  Forward delay(in seconds):15  Max-age(in
seconds):20  txHoldCount(in pps): 6
Configured     Hello time(in seconds): 2  Forward delay(in seconds):15  Max-age(in
seconds):20  Max-Hops:20
Root           Address:b8:d4:e7:da:80:00  Priority:0
               Port:1/1/27                Path cost:2000
Regional Root  Address:64:e8:81:dd:81:80  Priority:32768
               Internal cost:0            Rem Hops:20
...
```

Question: What is the current Root Port?

_____

Answer: The Root Port should have reverted to port 1/1/27.

19. On **PC1**, verify the ping has continued when the original uplink was restored.

Cleanup

20. On **PC1**, close the command prompt.

Task 3: Configure Spanning-Tree Access Security

In this task you will configure access ports that connect to clients with STP BPDU Guard. BPDU Guard will shut down a port when it receives an STP BPDU. This could occur when two wall outlets are connected to each other (network loop) or when a rogue switch is introduced into the network. In such cases, BPDU Guard will shut down the ports to contain the problem.

A BPDU Guard timeout can be configured to automatically restore the link after a set amount of time.
In the lab, the LAG between sw-edge1 and sw-edge2 will be used to test the feature. Since sw-edge2 is advertising BPDUs, sw-edge1 will shut down the link.

Objectives

- Configure STP access security with BPDU guard.
- Configure automatic recovery for BPDU guard disabled ports.

Steps

1. On sw-edge1, configure the BPDU Guard recovery timer. Set it to 60 seconds.

```
spanning-tree bpdu-guard timeout 60
```

```
sw-edge1(config)# spanning-tree bpdu-guard timeout 60
```

2. Configure a port range with ports 1/1/1 to 1/1/24 and enable BPDU Guard.

```
interface 1/1/1-1/1/24
 spanning-tree bpdu-guard
 exit
```

```
sw-edge1(config)# interface 1/1/1-1/1/24
sw-edge1(config-if-<1/1/1-1/1/24>)#  spanning-tree bpdu-guard
sw-edge1(config-if-<1/1/1-1/1/24>)#  exit
```

> **NOTE:** In the training lab, there are no STP switches connected to the access ports, so the previous step is only an example configuration.

To test the BPDU Guard feature, you will enable it on the LAG to sw-edge2, this is LAG50. Since sw-edge2 generates STP BPDUs, sw-edge1 will consider it a violation of the BPDU Guard and shutdown the LAG port.

3. On sw-edge1, enable BPDU Guard on the LAG 50 (link that connects to sw-edge2).

```
interface lag 50
 spanning-tree bpdu-guard
```

```
sw-edge1(config)# interface lag 50
sw-edge1(config-lag-if)#  spanning-tree bpdu-guard
```

4. Enable the LAG 50. Now STP BPDU frames that are generated by sw-edge2 will be received on a protected port.

```
no shutdown
```

```
sw-edge1(config-lag-if)# no shutdown
```

5. Review the Spanning-Tree output.

```
show spanning-tree
```

```
sw-edge1(config-lag-if)# show spanning-tree
...
lag50        Disabled       Bpdu-Error 20000          64         P2P              913
8          8          7
...
```

Question: What is the status for port lag50?

Answer: Port lag50 is reported with Bpdu-Error.

6. To see more details, review the Spanning-Tree inconsistent-ports output.

```
show spanning-tree inconsistent-ports
```

```
sw-edge1(config-lag-if)# show spanning-tree inconsistent-ports
Instance ID  Blocked Port   Reason
------------ -------------- ------------
0            lag50          BPDU Guard
```

Question: Do you see any inconsistent ports?

_____

Answer: The LAG 50 is reported as inconsistent port with reason BPDU Guard.

7. Wait 60 seconds and check the event log in reverse order.

```
show events -r -d hpe-mstpd
```

```
sw-edge1(config-lag-if)# show events -r -d hpe-mstpd
-------------------------------------------------
Event logs from current boot
-------------------------------------------------
2022-10-17T03:14:55.024057-04:00 sw-edge1 hpe-mstpd[3243]:
Event|2007|LOG_WARN|CDTR|1|Port lag50 disabled - BPDU received on protected port
2022-10-17T03:14:55.022192-04:00 sw-edge1 hpe-mstpd[3243]:
Event|2012|LOG_INFO|CDTR|1|CIST - Topology Change generated on port lag50 going in to
forwarding
2022-10-17T03:13:49.594762-04:00 sw-edge1 hpe-mstpd[3243]:
Event|2007|LOG_WARN|CDTR|1|Port lag50 disabled - BPDU received on protected port
2022-10-17T03:13:47.506076-04:00 sw-edge1 hpe-mstpd[3243]:
Event|2012|LOG_INFO|CDTR|1|CIST - Topology Change generated on port lag50 going in to
forwarding
...
```

Question: Did the port LAG 50 come back up?

_____

Answer: Yes, after the BPDU Guard timeout expired (60 seconds), sw-edge1 enabled the port LAG 50 again, but since it received another BPDU, the port was also disabled again.

Cleanup

8. On sw-edge1, remove the Spanning-Tree BPDU Guard command.

```
interface lag 50
 no spanning-tree bpdu-guard
```

```
sw-edge1(config-lag-if)# interface lag 50
sw-edge1(config-lag-if)#  no spanning-tree bpdu-guard
```

9. Shutdown the port LAG 50.

```
 shutdown
 exit
```

```
sw-edge1(config-lag-if)# shutdown
sw-edge1(config-lag-if)# exit
```

Task 4: Configure Loop Protection

After the successful STP setup, your customer informs you that they have had some issues with loops in the past that were not handled by STP. Some of their contractors had connected their own switches to some switch ports, and it turned out that these switches were blocking the STP BPDU frames, while at the same time, they were not sending out STP BPDU frames themselves. Since STP did not receive its own BPDUs back on the looped connection, the loop went unnoticed.

They have asked if you have a solution for this scenario.

After consulting with your senior colleague, you learn that Aruba switches have a feature that is known as Loop Protect. It is an Aruba-specific technology that can also detect loops.

The Aruba best practice is to have Loop Protect enabled on the access ports.

In this task you will enabled Loop Protect and verify its operation.

Objectives

- Configure Loop Protection.
- Verify the Loop Protection operation.

Prepare the setup

Introducing a loop in the network can have significant impact on the network. Spanning-Tree and Loop Protect are designed to handle loops in the network. However, if you make a configuration mistake while configuring Loop Protect or Spanning-Tree, the loop will remain effective in the remote lab environment.
To limit the impact of the loop, you will apply a rate-limiter on the looped port. This will ensure that Broadcasts, Multicast and Unknown Unicasts will only pass at 100 packets per second (PPS).

1. On sw-edge1, configure port lag50 with a rate-limiter.

```
interface lag 50
 rate-limit unknown-unicast 100 pps
 rate-limit broadcast 100 pps
 rate-limit multicast 100 pps
```

```
sw-edge1(config)# interface lag 50
sw-edge1(config-lag-if)#  rate-limit unknown-unicast 100 pps
sw-edge1(config-lag-if)#  rate-limit broadcast 100 pps
sw-edge1(config-lag-if)#  rate-limit multicast 100 pps
```

2. To demonstrate Loop Protect, you will disable STP on this port lag50. If you would leave STP enabled, STP may be faster with the loop detection, and you may not have a chance to see Loop Protect in action. BPDU-filter can be used to disable STP on a port. No STP BPDUs will be transmitted, and any received BPDUs will be dropped and ignored.

```
 spanning-tree bpdu-filter
```

```
 exit
```

```
sw-edge1(config-lag-if)#  spanning-tree bpdu-filter
This filter command allows the port to go into a continuous
forwarding mode and spanning tree will not interfere, even if
the port would cause a loop to form in the network topology.
If you suddenly experience high traffic load, shutdown the
port and remove the applicable filter configuration under
interface context using the following CLI command(s):
no spanning-tree bpdu-filter
no spanning-tree rpvst-filter
sw-edge1(config-lag-if)#  exit
```

Configure Loop Protect

Now that STP has been disabled, you can enable Loop Protection on the port and enable the port. This will introduce the loop in the network.

3. On the LAG 50, enabled Loop Protection.

```
interface lag 50
 loop-protect
```

```
sw-edge1(config)# interface lag 50
sw-edge1(config-lag-if)#  loop-protect
```

4. Enable the port LAG 50 to activate the loop.

```
 no shutdown
```

```
sw-edge1(config-lag-if)# no shutdown
```

Verify Loop Protect Operation

5. Review the event log file. Pay attention to the hpe-lpd daemon events.

```
show events -r
```

```
sw-edge1(config-lag-if)# show events -r
-------------------------------------------------
Event logs from current boot
-------------------------------------------------
2022-10-17T03:18:32.409479-04:00 sw-edge1 lacpd[3150]: Event|1321|LOG_INFO|CDTR|1|LAG
50 State change for interface 1/1/25: Actor state: ALFO, Partner state ALFNCD
2022-10-17T03:18:32.398850-04:00 sw-edge1 lacpd[3150]: Event|1321|LOG_INFO|CDTR|1|LAG
50 State change for interface 1/1/26: Actor state: ALFO, Partner state ALFNCD
2022-10-17T03:18:32.377134-04:00 sw-edge1 intfd[718]: Event|404|LOG_INFO|UKWN|1|Link
status for interface 1/1/26 is down - Network loop detected
2022-10-17T03:18:32.376924-04:00 sw-edge1 intfd[718]: Event|404|LOG_INFO|UKWN|1|Link
status for interface 1/1/25 is down - Network loop detected
2022-10-17T03:18:32.348222-04:00 sw-edge1 hpe-lpd[3131]:
Event|2808|LOG_INFO|CDTR|1|Ports TX lag50 and RX 1/1/27 are involved during TX port
disabling
```

```
2022-10-17T03:18:32.348056-04:00 sw-edge1 hpe-lpd[3131]:
Event|2801|LOG_WARN|CDTR|1|Port lag50 is disabled by Loop-protection after loop
detection on VLAN 1
```

Question: What are the TX and RX ports involved in the Loop Protect port disabling?

Answer: Ports lag50 and 1/1/27 (the uplink port to sw-agg1) are involved.

Question: Did you enable port 1/1/27 for loop protect?

Answer: No. Only the transmitting port (lag50) was enabled. When the looped packet comes back on any interface (In this example, via sw-edge2 > sw-agg1, back to port 1/1/27), the switch can disable the transmitting (TX) interface. In this example, this is the lag50.

6. Check the operational status for the ports.

```
show interface brief
```

```
sw-edge1(config)# show interface brief
--------------------------------------------------------------------------------
-------------------------
Port      Native  Mode   Type           Enabled Status  Reason                 Speed
Description
          VLAN                                                                  (Mb/s)
--------------------------------------------------------------------------------
-------------------------
...
1/1/25    1       access SFP+DAC1       yes     down    Network loop detected  --
--
1/1/26    1       access SFP+DAC1       yes     down    Network loop detected  --
--

...
```

7. Check the operational status for the individual port 1/1/25 (this is the member port of the lag 50).

```
show interface 1/1/25
```

```
sw-edge1(config)# show interface 1/1/25

Interface 1/1/25 is down
 Admin state is up
 State information: Network loop detected
 Link state: down for 3 minutes (since Thu Sep 22 05:25:00 EDT 2022)
 Link transitions: 192
 Description:
```

8. Check the loop protect status.

```
show loop-protect loop-detected
```

```
sw-edge1# show loop-protect loop-detected

Status and Counters - Loop Protection Information
```

```
Transmit Interval             : 5 (sec)
Port Re-enable Timer          : Disabled
Loop Detected Trap            : Disabled

Interface 1/1/25
  Loop-protect enabled        : Yes
  Action on loop detection    : TX disable
  Loop detected count         : 1
  Loop detected               : Yes
    Detected on VLAN          : 1
    Detected at               : 2022-09-22T05:25:00
  Interface status            : down
```

9. Review the interface 1/1/25.

```
show interface 1/1/25
```

```
sw-edge1(config-lag-if)# show interface 1/1/25

Interface 1/1/25 is down
 Admin state is up
 State information: Network loop detected
...
```

Question: What is the state information?

Answer: The state information reports network loop detected.

10. Check the current configuration of the looped port.

```
show running-config interface lag50
```

```
sw-edge1(config-lag-if)# show running-config int lag50
interface lag 50
    no shutdown
    no routing
    vlan trunk native 1
    vlan trunk allowed all
    lacp mode active
    rate-limit unknown-unicast 100 pps
    rate-limit broadcast 100 pps
    rate-limit multicast 100 pps
    loop-protect
    spanning-tree bpdu-filter
    exit
sw-edge1(config-lag-if)#
```

Question: Is the port disabled in the configuration with a shutdown command?

Answer: No, the port is operationally disabled, not through the ‘shutdown’ configuration command.

Cleanup

First you will recover the looped status on the ports, the last step will be disabling the LAG 50.

11. On sw-edge1, remove the STP BPDU filter on port lag 50.

```
interface lag 50
 no spanning-tree bpdu-filter
```

```
sw-edge1(config-lag-if)# interface lag 50
sw-edge1(config-lag-if)#  no spanning-tree bpdu-filter
```

12. Remove loop protect on the port lag 50.

```
 no loop-protect
 exit
```

```
sw-edge1(config-lag-if)#  no loop-protect
sw-edge1(config-lag-if)#  exit
```

13. Bounce the physical ports to clear the loop status, STP is now active again and it will set the port in Discarding state.

```
interface 1/1/25-1/1/26
 shutdown
 no shutdown
 exit
```

```
sw-edge1(config)# interface 1/1/25-1/1/26
sw-edge1(config-if-<1/1/25-1/1/26>)#  shutdown
sw-edge1(config-if-<1/1/25-1/1/26>)#  no shutdown
sw-edge1(config-if-<1/1/25-1/1/26>)#  exit
```

14. Verify with **show interface brief** that the ports 1/1/25, 1/1/26 and the LAG50 are **up** (This is physical up, STP is now blocking traffic on the ports).

```
show interface brief
```

```
sw-edge1(config)# show interface brief
--------------------------------------------------------------------------------
-------------------------
Port      Native  Mode   Type     Enabled Status  Reason               Speed
Description
          VLAN                                                          (Mb/s)
--------------------------------------------------------------------------------
-----...
1/1/27    1       trunk  SFP+DAC1 yes     up                           10000   sw-
agg1
1/1/28    1       trunk  SFP+DAC1 yes     up                           10000   sw-
agg2
vlan1     --              --       no      down    Administratively down  --      --
vlan3     --              --       yes     up                           --      --
lag50     1       trunk  --       yes     up      --                   20000   --
...
```

15. As a last step, shutdown the LAG 50.

```
interface lag 50
 shutdown
```

```
  exit
```

```
sw-edge1(config)# interface lag 50
sw-edge1(config-lag-if)#  shutdown
sw-edge1(config-lag-if)#  exit
```

16. Save the configurations on **all** switches.

```
write memory
```

You have completed the Lab!

Lab 04.02 VRRP

Overview

In your customer deployment, the final setup will use Aruba VSX and Active Gateway as the first hop redundancy protocol. This will offer an active-active default gateway function for the endpoints.

The active gateway feature is specific for the AOS-CX switches with VSX. In some environments, such as an SD-Branch gateway deployment, it may not be available and an alternative default gateway solution can be used, such as Virtual Router Redundancy Protocol (VRRP).Your senior colleague wants you to become familiar with the VRRP technology and has asked you to configure it on one of the VLANs.

The VRRP configuration will be replaced with the Active Gateway configuration in a later lab activity.

Objectives

After completing this lab, you will be able to:

- Configure VRRP between 2 Aruba switches.
- Review the status of VRRP.
- Verify the VRRP failover.

Task 1: Reconfigure sw-agg1 with VRRP on SVI11

In this task you will change the sw-agg1 SVI11 configuration. You will set the local IP address to 10.1.11.2/24 and you will activate VRRP with IP Address 10.1.11.1. This is not a redundant configuration yet, sw-agg2 will be added as the redundant VRRP peer in the next task.

Objectives

- Configure VRRP
- Verify VRRP Operation

Steps

1. On **PC1**, open a command prompt and review the ARP cache entry for 10.1.11.1.

```
arp -a 10.1.11.1
```

Example output.

```
C:\Users\student>arp -a 10.1.11.1

Interface: 10.1.11.50 --- 0xe
  Internet Address      Physical Address      Type
  10.1.11.1             b8-d4-e7-da-80-00     dynamic
```

Take note of the current MAC address of 10.1.11.1.

_____

2. Using MGMT PC, open an SSH connection to sw-agg1.
3. On **sw-agg1**, check the current SVI11 MAC address.

```
show interface vlan11
```

Example output.

```
sw-agg1(config)# show interface vlan11

Interface vlan11 is up
 Admin state is up
 Description:
 Hardware: Ethernet, MAC Address: b8:d4:e7:da:80:00
 IPv4 address 10.1.11.1/24
    active-gateway L3 source mac b8:d4:e7:da:80:00
 L3 Counters: Rx Disabled, Tx Disabled
```

Question: Does the SVI11 MAC address match the ARP cache entry of PC1?

Answer: Yes.

4. Check the global VRRP status.

```
show vrrp
```

```
sw-agg1(config)# show vrrp
```

```
VRRP is not enabled
```

5. Enable VRRP globally and verify the status again.

```
router vrrp enable
show vrrp
```

```
sw-agg1(config)# router vrrp enable
sw-agg1(config)# show vrrp

VRRP is enabled
```

6. Change the SVI11 IP address to 10.1.11.2/24

```
interface vlan 11
 ip address 10.1.11.2/24
```

```
sw-agg1(config)# interface vlan 11
sw-agg1(config-if-vlan)#  ip address 10.1.11.2/24
```

7. Create a VRRP VRID instance on SVI11, this will put you to the VRID context.

```
interface vlan 11
 vrrp 11 address-family ipv4
```

```
sw-agg1(config-if-vlan)# interface vlan 11
sw-agg1(config-if-vlan)#  vrrp 11 address-family ipv4
```

8. Set the VRRP primary address to 10.1.11.1.

```
 address 10.1.11.1 primary
```

```
sw-agg1(config-if-vrrp)#  address 10.1.11.1 primary
```

9. Set the VRRP VRID priority to 110

```
 priority 110
```

```
sw-agg1(config-if-vrrp)# priority 110
```

10. Enable the VRRP VRID 11 and return to the SVI configuration context.

```
 no shutdown
 exit
```

```
sw-agg1(config-if-vrrp)#  no shutdown
sw-agg1(config-if-vrrp)#  exit
```

11. Review the current SVI configuration. Verify you have the updated SVI IP address and the VRRP configuration.

```
show running-config current-context
```

```
sw-agg1(config-if-vlan)# show running-config current-context
```

```
interface vlan 11
    ip address 10.1.11.2/24
    ip helper-address 10.254.1.21
    ip ospf 1 area 0.0.0.0
    vrrp 11 address-family ipv4
        address 10.1.11.1 primary
        priority 110
        no shutdown
        exit
```

12. Review the VRRP status.

```
show vrrp
```

```
sw-agg1(config-if-vlan)# show vrrp

VRRP is enabled

Interface vlan11 - Group 11 - Address-Family IPv4
  State is ACTIVE
  State duration 31.386 secs
  Virtual IP address is 10.1.11.1
  Virtual MAC address is 00:00:5e:00:01:0b
  Advertisement interval is 1000 msec
  Version is 2
  Preemption is enabled
   min delay is 0 sec
  Priority is 110
  Active Router is 10.1.11.2 (local)
  Active Advertisement interval is 1000 msec
  Active Down interval is 3570 msec
```

Question: Do you see interface vlan11 in the VRRP list?

Answer: Yes

Question: What is the current state?

Answer: Active.

Question: What is the Virtual MAC Address for this VRID?

Answer: 00:00:5e:00:01:0b

13. Connect to **PC1** and check the ARP cache entry for 10.1.11.1 again.

```
arp -a 10.1.11.1
```

```
C:\Users\student>arp -a 10.1.11.1

Interface: 10.1.11.50 --- 0xe
  Internet Address      Physical Address      Type
  10.1.11.1             00-00-5e-00-01-0b     dynamic
```

Question: What do you observe?

Answer: The MAC address for the entry has changed to the VRRP Virtual MAC address 00:00:5e:00:01:0b.

Task 2: Configure sw-agg2 with VRRP on SVI11

In this task you will configure sw-agg2 with VRRP on SVI11. Sw-agg2 will act as the standby VRRP host, while sw-agg1 will be the active VRRP host.

On sw-agg2, you will first add SVI11 and assign it the .3 IP address in the subnet. Next you will activate VRRP with the default priority and with VIP .1. Finally, you will review the VRRP status on the active (sw-agg1) and standby (sw-agg2) systems.

Objectives

- Configure the standby VRRP host.
- Verify the standby VRRP status.

Steps

1. Use MGMT PC to open an SSH connection to **sw-agg2.**
2. **Sw-agg2** does not have an SVI for VLAN 11 yet. Review the L3 IP interfaces to verify that vlan11 is **not** listed.

```
show ip interface brief
```

3. Create SVI11 and assign static IP address 10.1.11.3/24.

```
interface vlan 11
 ip address 10.1.11.3/24
```

```
sw-agg2(config)#     interface vlan 11
sw-agg2(config-if-vlan)#  ip address 10.1.11.3/24
```

4. Enable the SVI for OSPF.

```
 ip ospf 1 area 0
```

```
sw-agg2(config-if-vlan)#  ip ospf 1 area 0
```

5. Now configure sw-agg2 SVI11 with the following VRRP properties:

VRID 11
Primary IP 10.1.11.1

```
 vrrp 11 address-family ipv4
  address 10.1.11.1 primary
```

```
sw-agg2(config-if-vlan)#  vrrp 11 address-family ipv4
sw-agg2(config-if-vrrp)#   address 10.1.11.1 primary
```

6. Review the VRRP status

```
show vrrp
```

```
sw-agg2(config-if-vrrp)# show vrrp

VRRP is not enabled

Interface vlan11 - Group 11 - Address-Family IPv4
  State is INIT (Group Disabled)
  State duration
  Virtual IP address is 10.1.11.1
  Virtual MAC address is 00:00:5e:00:01:0b
  Advertisement interval is 1000 msec
  Version is 2
  Preemption is enabled
   min delay is 0 sec
  Priority is 100
  Active Router is 0.0.0.0
  Active Advertisement interval is 1000 msec
  Active Down interval is 3609 msec
```

Question: What is the state of VRRP on Interface vlan11?

Answer: The state is INIT (Group Disabled).

Question: What should you change to enable the VRID group?

Answer: Enable the VRID.

7. Enable the VRID

```
 no shutdown
```

```
sw-agg2(config-if-vrrp)# no shutdown
```

8. Review the VRRP status again.

```
show vrrp
```

```
sw-agg2(config-if-vrrp)# show vrrp

VRRP is not enabled

Interface vlan11 - Group 11 - Address-Family IPv4
  State is INIT
  State duration
  Virtual IP address is 10.1.11.1
  Virtual MAC address is 00:00:5e:00:01:0b
  Advertisement interval is 1000 msec
  Version is 2
  Preemption is enabled
   min delay is 0 sec
  Priority is 100
  Active Router is 0.0.0.0
  Active Advertisement interval is 1000 msec
  Active Down interval is 3609 msec
```

Question: What is the VRRP state?

Answer: Interface VLAN 11 state is INIT, the global state is not enabled.

Question: What should you change to enable VRRP?

Answer: Enable VRRP at the global level.

9. Return to the global configuration level.

```
exit
exit
```

```
sw-agg2(config-if-vrrp)# exit
sw-agg2(config-if-vlan)# exit
```

10. Enable VRRP at the switch global configuration level.

```
router vrrp enable
```

```
sw-agg2(config)# router vrrp enable
```

11. Verify the VRRP state again.

```
show vrrp
```

```
sw-agg2(config)# show vrrp

VRRP is enabled

Interface vlan11 - Group 11 - Address-Family IPv4
  State is STANDBY
  State duration 02.158 secs
  Virtual IP address is 10.1.11.1
  Virtual MAC address is 00:00:5e:00:01:0b
  Advertisement interval is 1000 msec
  Version is 2
  Preemption is enabled
   min delay is 0 sec
  Priority is 100
  Active Router is 10.1.11.2
  Active Advertisement interval is 1000 msec
  Active Down interval is 3609 msec
```

Question: What is the VRRP state for Interface vlan11?

Answer: STANDBY.

Question: What is currently the active router?

Answer: The active router is 10.1.11.2 (sw-agg1).

Question: What is the default priority of VRRP?

Answer: The default priority is 100. Since sw-agg1 has been configured with VRRP priority 110, it will be the active VRRP host.

Task 3: Test the VRRP Failover

In this task you will test the VRRP VIP failover. On the PC1 you will initiate some test traffic with a continuous ping. While this ping is active, you will disable the sw-agg1 SVI11 interface. VRRP on sw-agg2 will detect that the VRRP Keep-Alive packets are missing, and it will resume the active state on the configured VIP.

Objectives

- Verify the VRRP Failover from Active to Standby.
- Verify the VRRP Pre-empt function.

Steps

1. Open a connection to PC1.
2. Open a command prompt and start a continuous ping to 10.254.1.21.

```
ping 10.254.1.21 -t
```

3. Open a second command prompt and run a traceroute to 10.254.1.21

```
tracert -d 10.254.1.21
```

```
C:\Users\student>tracert -d 10.254.1.21

Tracing route to 10.254.1.21 over a maximum of 30 hops

  1    <1 ms    <1 ms    <1 ms  10.1.11.2
  2     2 ms     1 ms     1 ms  10.254.101.254
  3     2 ms     1 ms     1 ms  10.254.1.21

Trace complete.
```

Question: What is the IP address of the first router in the path?

Answer: Currently, sw-agg1 with IP 10.1.11.2 is the first router that responds in the traceroute path.

4. Use MGMT PC to open an SSH connection to **sw-agg1**
5. Enter the SVI11 context and shutdown the SVI interface.

```
interface vlan 11
 shutdown
```

```
sw-agg1(config)# interface vlan 11
sw-agg1(config-if-vlan)# shutdown
```

6. On PC-1, verify that the ping continued after a few seconds.

```
Reply from 10.254.1.21: bytes=32 time=2ms TTL=126
Reply from 10.254.1.21: bytes=32 time=1ms TTL=126
Request timed out.
Reply from 10.254.1.21: bytes=32 time=2ms TTL=126
Reply from 10.254.1.21: bytes=32 time=2ms TTL=126
```

Question: What is the failover time for VRRP?

Answer: VRRP takes about 3-4 seconds to detect and take over the Virtual IP address by default.

7. On PC-1, run the traceroute to 10.254.1.21 again.

```
tracert -d 10.254.1.21
```

```
C:\Users\student>tracert -d 10.254.1.21

Tracing route to 10.254.1.21 over a maximum of 30 hops

  1    <1 ms    <1 ms    <1 ms  10.1.11.3
  2     2 ms     1 ms     1 ms  10.254.102.254
  3     1 ms     1 ms     1 ms  10.254.1.21

Trace complete.
```

Question: What is the first router in the path?

Answer: The first router is now sw-agg2 with IP address 10.1.11.3.

8. On sw-agg1, review the VRRP state.

```
show vrrp
```

```
sw-agg1(config-if-vlan)# show vrrp

VRRP is enabled

Interface vlan11 - Group 11 - Address-Family IPv4
  State is INIT (Interface Down)
  State duration 01 mins 49.683 secs
...
```

9. On sw-agg2, review the VRRP state.

```
show vrrp
```

```
sw-agg2(config)# show vrrp

VRRP is enabled

Interface vlan11 - Group 11 - Address-Family IPv4
  State is ACTIVE
  State duration 02 mins 22.561 secs
...
```

Question: What is the sw-agg2 state?

Answer: Sw-agg2 is now ACTIVE.

This demonstrates the VRRP failover function

Now you will explore what happens when the original sw-agg1 SVI11 comes back again.

10. On sw-agg1, review the VRRP state.

```
show vrrp
```

```
sw-agg1(config-if-vlan)# show vrrp

VRRP is enabled

Interface vlan11 - Group 11 - Address-Family IPv4
  State is INIT (Interface Down)
  State duration 01 mins 49.683 secs
  Virtual IP address is 10.1.11.1
  Virtual MAC address is 00:00:5e:00:01:0b
  Advertisement interval is 1000 msec
  Version is 2
  Preemption is enabled
    min delay is 0 sec
  Priority is 110
  Active Router is 0.0.0.0
  Active Advertisement interval is 1000 msec
  Active Down interval is 3570 msec
```

Question: What is the preemption state for interface vlan11?

Answer: Preemption is enabled by default.

Question: What is the expected behavior when preemption is enabled?

Answer: With preemption enabled, when a VRRP with a better priority returns to service, it will claim control of the Virtual IP address.

11. Verify the preemption behavior. On sw-agg1, enable SVI11 again.

```
interface vlan 11
 no shutdown
```

```
sw-agg1(config-if-vlan)# interface vlan 11
sw-agg1(config-if-vlan)#  no shutdown
```

12. Wait about 3 seconds, then review the VRRP state.

```
show vrrp
```

```
sw-agg1(config-if-vlan)# show vrrp

VRRP is enabled

Interface vlan11 - Group 11 - Address-Family IPv4
  State is ACTIVE
  State duration 10.691 secs
```

13. On sw-agg2, review the VRRP state.

```
show vrrp
```

```
sw-agg2(config)# show vrrp

VRRP is enabled

Interface vlan11 - Group 11 - Address-Family IPv4
  State is STANDBY
  State duration 35.574 secs
```

Question: What is the state on sw-agg2?

Answer: Sw-agg2 has now reverted to STANDBY, since sw-agg1 performed a preempt of the ACTIVE role.

14. Save the configuration on both aggregation switches.

```
write memory
```

Task 4: Optional – MAC address tracing of the aggregation switches

Trace the MAC address of traffic sent to the router and traffic from the router to the endpoint when a failover occurs.

Objectives

- Understand VRRP MAC address use in the network.

Steps

Prepare the setup

1. Use MGMT PC to open an SSH connection to **sw-agg2**.

By default, OSPF ECMP would have the core router distribute traffic destined to 10.1.11.0/24 to both sw-agg1 and sw-agg2. In this lab we need all the traffic to return via sw-agg1. You will achieve this by increasing the cost of SVI11 on sw-agg2.

2. Enter the SVI 11 context and set OSPF cost 200. The default cost is 100.

```
interface vlan 11
 ip ospf cost 200
```

```
sw-agg2(config)# interface vlan 11
sw-agg2(config-if-vlan)#  ip ospf cost 200
```

Start test traffic

3. On PC1, start a continuous ping to 10.254.1.21

```
ping 10.254.1.21 -t
```

4. On PC-1, start Wireshark trace on the LAB NIC.
5. Set a display filter to '**icmp**'.Don't forget to press <ENTER> to activate the filter.

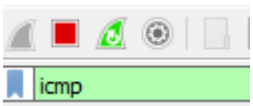

6. Open an SSH connection to sw-agg1.
7. Enter the SVI 11 and shutdown the SVI.

```
interface vlan 11
 shutdown
 exit
```

```
sw-agg1(config)# interface vlan 11
sw-agg1(config-if-vlan)#  shutdown
sw-agg1(config-if-vlan)#  exit
```

8. On PC-1, wait until the ping continues, then stop the Wireshark trace.

Analyze the return traffic

9. Explore **ICMP request** from the beginning of the trace, check the destination MAC. This should be the VRRP MAC.

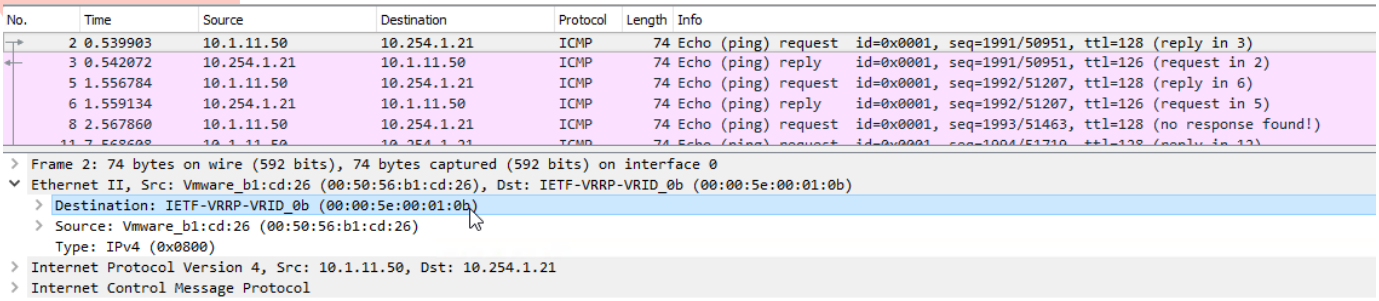

10. Explore **ICMP reply** from the beginning of the trace, check the source MAC. This is the sw-agg1 SVI11 MAC.

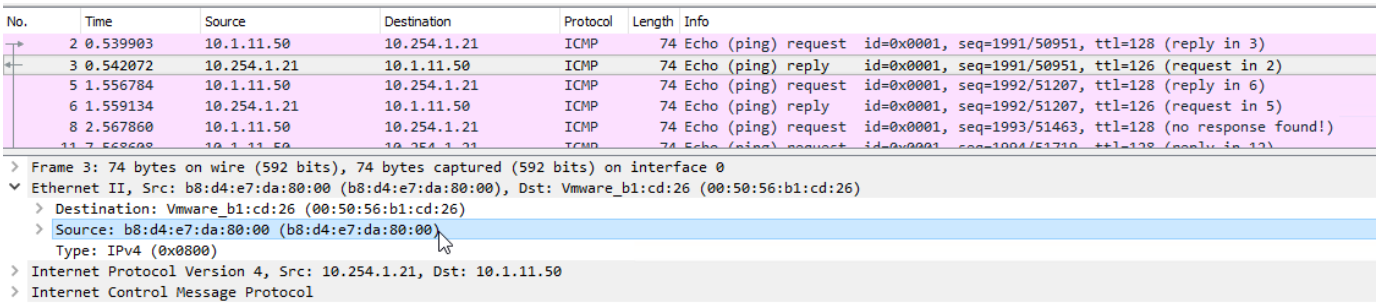

11. Now move to the end of the trace.
12. Explore **ICMP request** from the end of the trace, check the destination MAC. This should be the VRRP MAC. This shows that from the endpoint point of view, the Default MAC address does not change.

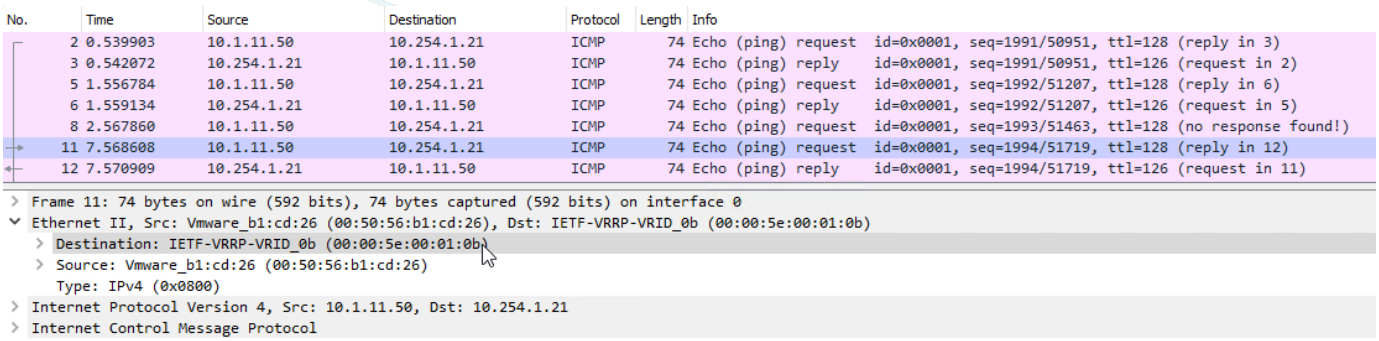

13. Explore **ICMP reply** from the end of the trace - check the source MAC. This is the sw-agg2 SVI11 MAC.

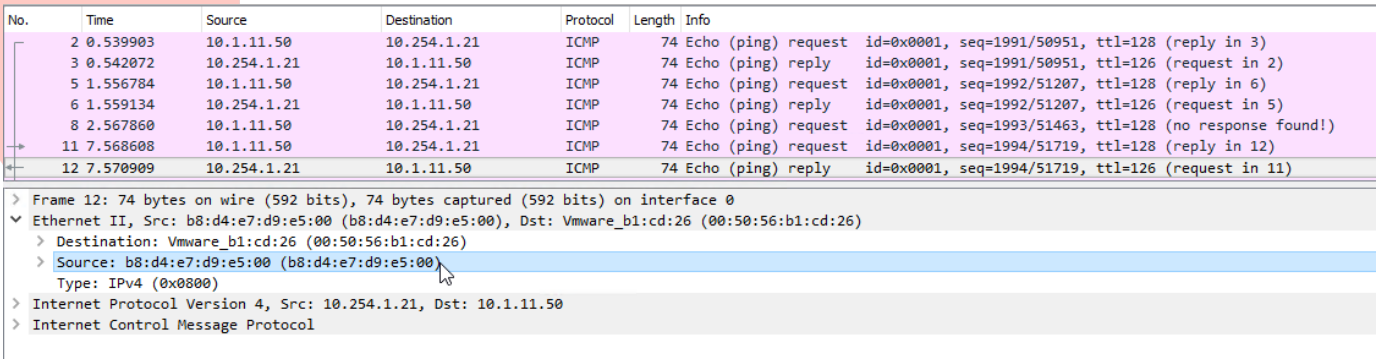

No.	Time	Source	Destination	Protocol	Length	Info
2	0.539903	10.1.11.50	10.254.1.21	ICMP	74	Echo (ping) request id=0x0001, seq=1991/50951, ttl=128 (reply in 3)
3	0.542072	10.254.1.21	10.1.11.50	ICMP	74	Echo (ping) reply id=0x0001, seq=1991/50951, ttl=126 (request in 2)
5	1.556784	10.1.11.50	10.254.1.21	ICMP	74	Echo (ping) request id=0x0001, seq=1992/51207, ttl=128 (reply in 6)
6	1.559134	10.254.1.21	10.1.11.50	ICMP	74	Echo (ping) reply id=0x0001, seq=1992/51207, ttl=126 (request in 5)
8	2.567860	10.1.11.50	10.254.1.21	ICMP	74	Echo (ping) request id=0x0001, seq=1993/51463, ttl=128 (no response found!)
11	7.568608	10.1.11.50	10.254.1.21	ICMP	74	Echo (ping) request id=0x0001, seq=1994/51719, ttl=128 (reply in 12)
12	7.570909	10.254.1.21	10.1.11.50	ICMP	74	Echo (ping) reply id=0x0001, seq=1994/51719, ttl=126 (request in 11)

NOTE: Example traces from the screenshots can be found in the ACAF Student Files on the desktop of MGMT PC:

lab04.02-vrrp-svi-mac-trace.pcapng

Cleanup

14. On sw-agg2, restore the SVI11 OSPF cost to the default by removing the configured cost.

```
interface vlan 11
 no ip ospf cost
 exit
```

```
sw-agg2(config-if-vlan)# interface vlan 11
sw-agg2(config-if-vlan)#  no ip ospf cost
sw-agg2(config-if-vlan)#  exit
```

15. On sw-agg1, enable the SVI11 again.

```
interface vlan 11
 no shutdown
 exit
```

```
sw-agg1(config)# interface vlan 11
sw-agg1(config-if-vlan)#  no shutdown
sw-agg1(config-if-vlan)#  exit
```

16. On PC1, close the command prompt and Wireshark.

Task 5: Optional – VRRP Preempt Delay

In some scenario's, performing immediate VRRP pre-emption may have a negative side-effect. When the router needs to establish upstream routing adjacency, for example with OSPF, the router may take control of the VIP address before it has completed the OSPF peering. In that case, the clients will have a (new) default gateway, but that default gateway does not have routes (yet) to forward the traffic. The result is that traffic will be dropped.

To solve this, you can introduce a delay timer before taking control of the endpoint default gateway function. This gives upstream routing a chance to learn routes before claiming control of the default gateway VIP.

In this lab activity, it applies to upstream OSPF adjacency, but in case of SD-WAN gateways, it may apply to establishing upstream IPsec connections and learning SD-WAN routes.

Objectives

- Configure VRRP Pre-empt delay timer.
- Verify the operation of the Pre-empt delay timer.

Steps

Verify the problem

1. On **PC1**, open a command prompt and start a continuous ping to 10.254.1.21.

```
ping 10.254.1.21 -t
```

2. Use MGMT PC to open an SSH connection to **sw-agg1**.
3. On **sw-agg1**, save the configuration.

```
write memory
```

4. On **sw-agg1**, initiate the reboot and confirm the reboot with **y**.

```
boot system
y
```

```
sw-agg1(config)# boot system
Checking if the configuration needs to be saved...

Checking for updates needed to programmable devices...
Done checking for updates.

1 non-failsafe device(s) also need to be updated.
Please run the 'allow-unsafe-updates' command to enable these updates.


This will reboot the entire switch and render it unavailable
until the process is complete.
Continue (y/n)? y
```

5. On **PC1**, verify the failover takes about 3 seconds (1 lost ping).

```
Reply from 10.254.1.21: bytes=32 time=1ms TTL=126
Reply from 10.254.1.21: bytes=32 time=3ms TTL=126
Request timed out.
Reply from 10.254.1.21: bytes=32 time=4ms TTL=126
Reply from 10.254.1.21: bytes=32 time=1ms TTL=126
```

6. Wait for the sw-agg1 to complete the reboot (about 2 minutes), keep checking the PC1 ping output.

```
Reply from 10.254.1.21: bytes=32 time=3ms TTL=126
Reply from 10.254.1.21: bytes=32 time=2ms TTL=126
Reply from 10.254.1.21: bytes=32 time=2ms TTL=126
Request timed out.
Reply from 10.1.11.2: Destination net unreachable.
Reply from 10.1.11.2: Destination net unreachable.
Reply from 10.1.11.2: Destination net unreachable.
Reply from 10.1.11.2: Destination net unreachable.
Request timed out.
Reply from 10.1.11.2: Destination net unreachable.
Reply from 10.1.11.2: Destination net unreachable.
Reply from 10.1.11.2: Destination net unreachable.
Reply from 10.1.11.2: Destination net unreachable.
Request timed out.
Reply from 10.1.11.2: Destination net unreachable.
Reply from 10.1.11.2: Destination net unreachable.
Request timed out.
Reply from 10.254.1.21: bytes=32 time=1ms TTL=126
Reply from 10.254.1.21: bytes=32 time=2ms TTL=126
Reply from 10.254.1.21: bytes=32 time=2ms TTL=126
```

Question: What do you observe when sw-agg1 comes back online?

Answer: Several ping packets are lost when the sw-agg1 returned to service.

Question: What is the reported ping error and what IP address reports this error?

Answer: The Reply comes from 10.1.11.2, with the message “Destination net unreachable”.

Question: What does this error indicate?

Answer: This indicates that 10.1.11.2 (sw-agg1) is receiving the ICMP request, since it is the active VRRP host. However, it does not have a route to the requested destination (10.254.1.21), since the OSPF adjacency was not established yet. Some seconds later, when the OSPF adjacency is established and the routing tables has been populated, the ping works again.

Question: How could this timing problem be solved?

Answer: The VRRP system should be delayed a little bit to give the upstream routing protocol (OSPF) the opportunity to learn the upstream routes. A typical delay is 60 seconds.

Configure Preempt Delay Timer on sw-agg1

7. On **sw-agg1**, configure a preempt delay timer of 60 seconds.

```
interface vlan 11
```

```
 vrrp 11 address-family ipv4
  preempt delay minimum 60
  exit
 exit
```

```
sw-agg1(config)# interface vlan 11
sw-agg1(config-if-vlan)#  vrrp 11 address-family ipv4
sw-agg1(config-if-vrrp)#   preempt delay minimum 60
sw-agg1(config-if-vrrp)#   exit
sw-agg1(config-if-vlan)#  exit
```

Verify the configured solution

8. On **PC1**, verify the ping to 10.254.1.21 is still running.
9. On **sw-agg1**, save the configuration and reboot the switch.

```
write memory
boot system
```

10. On **PC1**, verify the ping continues when sw-agg1 completes the reboot.

Cleanup steps

11. On PC1, close the command prompt.

There is no need to revert the preempt delay timer on sw-agg1.

You have completed the lab!

Lab 04.03 Introduction to VSX

Overview

In this lab you will be introduced to an Aruba redundant aggregation switch configuration, known as Virtual Switching Extension (VSX). Your senior colleague has completed the aggregation switch configuration for the customer setup and wants you to review the VSX configuration that will be active on the aggregation switches. The working VSX configuration will be loaded to the sw-agg1 and sw-agg2 systems, and you will be able to explore the status and the basic configuration of VSX.

As the associate engineer, you will then connect the sw-edge1 in a redundant way to the VSX aggregation using a VSX LAG. A VSX LAG is also known as Multi-Chassis Link-Aggregation (MCLAG).

In the last task you will explore Aruba VSX active gateway. This is an active-active default gateway solution and replaces the active-standby VRRP protocol in a VSX cluster.

Objectives

After completing this lab, you will be able to:

- Review a working VSX system status.
- Review MCLAG port configuration.
- Verify the VSX MCLAG operation from the access switch point of view.
- Understand the LACP fallback feature.
- Verify Active Gateway configuration.

Task 1: Prepare the topology

In this task you will load the correct start configuration for this lab on

- Sw-agg1
- Sw-agg2

Your senior colleague has prepared the VSX configuration for the customer aggregation switches. You have been asked to apply the configuration to the sw-agg1 and sw-agg2 switches.

Objectives

- Apply a prepared configuration to the Aggregation switches.

Steps

> **IMPORTANT:** Make sure you apply the correct configuration to each aggregation switch in the next steps. Please double-check that you are applying sw-agg1 configuration to the actual sw-agg1 host, and the sw-agg2 configuration to the sw-agg2 host.

sw-agg1

1. Use MGMT PC to open an SSH connection to sw-agg1.
2. On MGMT PC, open the folder **ACAF Student Files** on the desktop, copy the contents of the file named **lab04.03-vsx-to-load-on-sw-agg1.txt.**
3. Paste the commands in the SSH connection of sw-agg1.

sw-agg1

4. Use MGMT PC to open an SSH connection to sw-agg2.
5. On MGMT PC, open the folder **ACAF Student Files** on the desktop, copy the contents of the file named **lab04.03-vsx-to-load-on-sw-agg2**.
6. Paste the commands in the SSH connection of sw-agg2.

With this configuration applied, the aggregation switches have been configured with VSX. In the next task you will explore the status of VSX on the aggregation switches.

Task 2: Review Global VSX configuration and status

In this task you will review the VSX configuration that was prepared by your senior network administrator colleague. This will allow you to understand the basic concept of VSX redundancy at the aggregation layer. You will review the global VSX configuration and status, the roles assigned to the switches in the VSX cluster, and the VSX system MAC.

Objectives

- Understand the VSX components.
- Describe the VSX ISL.
- Describe the VSX device roles.
- Understand the VSX system MAC.

Steps

1. Open an SSH connection to sw-agg1.
2. First review the connections to sw-agg2.

```
show lldp neighbor-info
```

```
sw-agg1(config)# show lldp neighbor-info

LLDP Neighbor Information
=========================

Total Neighbor Entries          : 6
Total Neighbor Entries Deleted  : 0
Total Neighbor Entries Dropped  : 0
Total Neighbor Entries Aged-Out : 0

LOCAL-PORT  CHASSIS-ID         PORT-ID                    PORT-DESC
TTL      SYS-NAME
-------------------------------------------------------------------------------------------
----------------------
1/1/1       64:e8:81:dd:81:80  1/1/27                     sw-agg1
120      sw-edge1
1/1/2       8c:85:c1:49:20:c0  1/1/27                     sw-agg1
120      sw-edge2
...
1/1/46      b8:d4:e7:d9:e5:00  1/1/46                     1/1/46
120      sw-agg2
1/1/47      b8:d4:e7:d9:e5:00  1/1/47                     1/1/47
120      sw-agg2
...
```

Question: What ports are connected to sw-agg2?

Answer: Ports 1/1/46 and 1/1/47.

3. Review the LAG port members.

```
show lacp interfaces
```

```
sw-agg1(config)# show lacp interfaces

State abbreviations :
A - Active        P - Passive      F - Aggregable I - Individual
S - Short-timeout L - Long-timeout N - InSync     O - OutofSync
C - Collecting    D - Distributing
X - State m/c expired              E - Default neighbor state

Actor details of all interfaces:
------------------------------------------------------------------------------
Intf       Aggr       Port  Port  State   System-ID         System Aggr Forwarding
           Name       Id    Pri                             Pri    Key  State
------------------------------------------------------------------------------
1/1/1      lag1(mc)   1     1     IE      02:01:00:00:01:00 65534  1    up
1/1/2      lag2(mc)   2     1     IE      02:01:00:00:01:00 65534  2    up
1/1/46     lag256     47    1     ALFNCD  b8:d4:e7:da:80:00 65534  256  up
1/1/47     lag256     48    1     ALFNCD  b8:d4:e7:da:80:00 65534  256  up


Partner details of all interfaces:
------------------------------------------------------------------------------
Intf       Aggr       Port  Port  State   System-ID         System Aggr
           Name       Id    Pri                             Pri    Key
------------------------------------------------------------------------------
1/1/1      lag1(mc)   0     0     IE      00:00:00:00:00:00 0      0
1/1/2      lag2(mc)   0     0     IE      00:00:00:00:00:00 0      0
1/1/46     lag256     47    1     ALFNCD  b8:d4:e7:d9:e5:00 65534  256
1/1/47     lag256     48    1     ALFNCD  b8:d4:e7:d9:e5:00 65534  256
```

Question: What LAG do ports 1/1/46 and 1/1/47 belong to?

Answer: These ports belong to LAG256.

4. Review the VSX status output.

```
show vsx status
```

```
sw-agg1(config)# show vsx status
VSX Operational State
---------------------
  ISL channel             : In-Sync
  ISL mgmt channel        : operational
  Config Sync Status      : In-Sync
  NAE                     : peer_reachable
  HTTPS Server            : peer_reachable

Attribute           Local               Peer
------------        --------            --------
ISL link            lag256              lag256
ISL version         2                   2
System MAC          02:01:00:00:01:00   02:01:00:00:01:00
Platform            8325                8325
Software Version    GL.10.09.1040       GL.10.09.1040
Device Role         primary             secondary
```

Question: What is the VSX role of sw-agg1?

Answer: Sw-agg1 has been configured as primary.

Question: What is the Inter Switch Link (ISL) used for VSX communication between sw-agg1 and sw-agg2?

Answer: LAG256. Typically, a redundant link, such as a Link Aggregation, is used as the ISL.

Question: What is the system MAC?

Answer: On the lab VSX system, the system MAC has been set to 02:01:00:00:01:00.

Question: What is the difference between the VSX system MAC and the base switch MAC?

Answer: The base switch MAC is used as the 'local' Layer2 identifier of the switch, for example in LLDP communication. The VSX system MAC is used in VSX 'shared' Layer 2 protocols, such as LACP on the MCLAG and Spanning-Tree Bridge ID on a VSX system.

Review the VSX LAG (MCLAG) configuration

In the previous labs, you have configured a LAG between sw-edge1 and sw-edge2. A traditional LAG requires all the member ports to be connected between the same two devices.

With VSX, the adjacent devices can be connected with a LAG, using one link to the primary VSX node and another link to the secondary VSX node (multiple links are also possible).

Examples for adjacent device would be:

- an access switch with a LAG
- a server with NIC teaming
- a firewall with port bonding
- an Aruba Gateway with a port-channel

The senior administrator has prepared the VSX system with VSX LAG 1 for sw-edge1 and VSX LAG 2 connecting to sw-edge2.

5. Review the VSX settings in the running configuration, look for the LAG interfaces.

```
show running-config vsx
```

```
sw-agg1(config)# show running-config vsx
vsx
    system-mac 02:01:00:00:01:00
    inter-switch-link lag 256
    role primary
    vsx-sync dhcp-relay mclag-interfaces ospf vsx-global
interface lag 256
    no shutdown
    no routing
    vlan trunk native 1
    vlan trunk allowed all
    lacp mode active
interface 1/1/47
```

```
    no shutdown
    lag 256
interface 1/1/46
    no shutdown
    lag 256
interface lag 1 multi-chassis
    no shutdown
    no routing
    vlan trunk native 1
    vlan trunk allowed all
    lacp mode active
...
```

Question: Do you see anything special for interface lag 1?

Answer: The interface lag 1 has a new option: multi-chassis. This indicates that the state will be shared with the VSX peer.

6. Review the LACP interfaces.

```
show lacp interfaces
```

```
sw-agg1(config)# show lacp interfaces

State abbreviations :
A - Active        P - Passive      F - Aggregable I - Individual
S - Short-timeout L - Long-timeout N - InSync     O - OutofSync
C - Collecting    D - Distributing
X - State m/c expired              E - Default neighbor state

Actor details of all interfaces:
------------------------------------------------------------------------------
Intf       Aggr       Port  Port  State   System-ID         System Aggr Forwarding
           Name       Id    Pri                             Pri    Key  State
------------------------------------------------------------------------------
1/1/1      lag1(mc)   1     1     IE      02:01:00:00:01:00 65534  1    up
1/1/2      lag2(mc)   2     1     IE      02:01:00:00:01:00 65534  2    up
1/1/46     lag256     47    1     ALFNCD  b8:d4:e7:da:80:00 65534  256  up
1/1/47     lag256     48    1     ALFNCD  b8:d4:e7:da:80:00 65534  256  up


Partner details of all interfaces:
------------------------------------------------------------------------------
Intf       Aggr       Port  Port  State   System-ID         System Aggr
           Name       Id    Pri                             Pri    Key
------------------------------------------------------------------------------
1/1/1      lag1(mc)   0     0     IE      00:00:00:00:00:00 0      0
1/1/2      lag2(mc)   0     0     IE      00:00:00:00:00:00 0      0
1/1/46     lag256     47    1     ALFNCD  b8:d4:e7:d9:e5:00 65534  256
1/1/47     lag256     48    1     ALFNCD  b8:d4:e7:d9:e5:00 65534  256
```

Question: What does **mc** represents in the lag1(mc) output?

Answer: This indicates that LAG1 is a MCLAG (multi-chassis), as opposed to the LAG256, that is a 'local' LAG.

Question: What do the sections ‘Actor’ and ‘Partner’ represent?

Answer:
Actor: Represents the local LACP communication information - what this side of the LAG will advertise to the adjacent device.

Partner: Represents the LACP information that has been received from the partner adjacent device. In case no LACP information has been received from the partner, the System-ID will be shown as 00:00:00:00:00:00. In this lab environment, the sw-edge1 or sw-edge2 will be the LACP partner.

Question: In the ‘Actor’ section, why is the System-ID for the LAG1 different from the LAG256?

Answer: LAG1 is an MCLAG, so both sw-agg1 and sw-agg2 will use this VSX system MAC in their LACP communication with the sw-edge1 access switch. LAG256 is a local LAG between sw-agg1 and sw-agg2, so each aggregation switch will use its own Layer 2 identifier (MAC Address).

7. Since both sw-agg1 and sw-agg2 are part of the MCLAG configuration and status, they need to be aware of each other’s LACP states. This will be synchronized by VSX over the ISL. To see the LACP state for an MCLAG, use the command **show lacp interfaces multi-chassis**.

```
show lacp interfaces multi-chassis
```

```
sw-agg1(config)# show lacp interfaces multi-chassis

State abbreviations :
A - Active        P - Passive      F - Aggregable I - Individual
S - Short-timeout L - Long-timeout N - InSync     O - OutofSync
C - Collecting    D - Distributing
X - State m/c expired              E - Default neighbor state

Actor details of all interfaces:
------------------------------------------------------------------------------
Intf       Aggregate  Port    Port     State   System-ID         System   Aggr
           name       id      Priority                           Priority Key
------------------------------------------------------------------------------
1/1/1      lag1(mc)   1       1        IE      02:01:00:00:01:00 65534    1
1/1/2      lag2(mc)   2       1        IE      02:01:00:00:01:00 65534    2


Partner details of all interfaces:
------------------------------------------------------------------------------
Intf       Aggregate  Partner Port     State   System-ID         System   Aggr
           name       Port-id Priority                           Priority Key
------------------------------------------------------------------------------
1/1/1      lag1(mc)   0       0        IE      00:00:00:00:00:00 0        0
1/1/2      lag2(mc)   0       0        IE      00:00:00:00:00:00 0        0


Remote Actor details of all interfaces:
------------------------------------------------------------------------------
Intf       Aggregate  Port    Port     State   System-ID         System   Aggr
```

```
           name       id      Priority                            Priority Key
-------------------------------------------------------------------------------
1/1/1      lag1(mc)   1001    1        IE      02:01:00:00:01:00 65534    1


Remote Partner details of all interfaces:
-------------------------------------------------------------------------------
Intf       Aggregate  Partner Port     State   System-ID         System   Aggr
           name       Port-id Priority                           Priority Key
-------------------------------------------------------------------------------
1/1/1      lag1(mc)   0       0        IE      00:00:00:00:00:00 0        0
```

Question: What is different in this output compared to the basic **show lacp interfaces** command?

Answer: This command also includes '**Remote Actor**' and '**Remote Partner**' output. These sections show the status on the VSX peer. Since you have executed this command on the sw-agg1, the Remote Actor shows the 'actor' state of sw-agg2.

Question: Is the Remote Actor using a different System-ID compared to the Actor System-ID?

Answer: No, for MCLAGs, both VSX switches will identify themselves using the VSX System MAC, so the adjacent device believes it is connected to the same device.

LACP Fallback

A successful LACP connection, such as the LAG256 has LACP State ALFNCD.

8. Check the output of **show lacp interfaces**.

```
show lacp interfaces
```

```
sw-agg1(config)# show lacp interfaces

State abbreviations :
A - Active        P - Passive      F - Aggregable I - Individual
S - Short-timeout L - Long-timeout N - InSync     O - OutofSync
C - Collecting    D - Distributing
X - State m/c expired              E - Default neighbor state

Actor details of all interfaces:
-------------------------------------------------------------------------------
Intf       Aggr       Port  Port  State   System-ID         System Aggr Forwarding
           Name       Id    Pri                             Pri    Key  State
-------------------------------------------------------------------------------
1/1/1      lag1(mc)   1     1     IE      02:01:00:00:01:00 65534  1    up
1/1/2      lag2(mc)   2     1     IE      02:01:00:00:01:00 65534  2    up
...
```

Question: What is the LACP state for lag1(mc)?

Answer: The state is **IE**.

Question: Use the **Flags** legend from the output, what do the IE Flags indicate?

Answer: I = Individual, E = Defaulted Neighbor. This state is shown when the peer does not have LACP enabled and the LACP Fallback feature is enabled. In the next steps you will explore what this LACP Fallback means.

Question: What is the forwarding state?

Answer: The Forwarding State is **up**.

In a previous lab, you configured the LAG between sw-edge1 and sw-edge2. When sw-edge1 was configured with LACP, but sw-edge2 was not configured yet, the Forwarding State was **lacp-blocked**.
In this activity, the peer is not configured yet, but the Forwarding State is **up**.

Let's investigate what is different now.

9. Check the lag1 interface configuration.

```
show running-config interface lag1
```

```
sw-agg1(config)# show running-config interface lag1
interface lag 1 multi-chassis
    no shutdown
    no routing
    vlan trunk native 1
    vlan trunk allowed all
    lacp mode active
    lacp fallback
    exit
```

Question: What new command do you observe in this LAG configuration context?

Answer: The command **lacp fallback** is a new command.

Question: What is the purpose of this command?

Answer: By default, an LACP LAG will block all traffic, unless it can successfully negotiate with an LACP partner. During new switch deployments with ZTP, when a factory default switch is connected, it does not have any LACP configuration yet. This means a factory default switch would never be able to connect to an existing LACP LAG.
With the **lacp fallback** command, you can set the member ports of the LAG to Forwarding when no LACP control frames are received. Instead of the status 'lacp-blocked', the status will be 'up'.
Verify the LACP fallback is working by testing connectivity on sw-edge1.

10. On sw-edge1, verify you can still reach 10.254.1.21.

```
ping 10.254.1.21
```

Optional steps: Test default LACP behavior

Be sure to continue after this optional section!

11. On sw-agg1, disable the LACP fallback feature on LAG1.

```
interface lag 1
 no lacp fallback
 exit
```

```
sw-agg1(config)# interface lag 1
sw-agg1(config-lag-if)#  no lacp fallback
sw-agg1(config-lag-if)# exit
```

12. On sw-agg1, verify the LACP state now changed to **lacp-blocked**.

```
show lacp interfaces
```

```
sw-agg1(config)# show lacp interfaces

State abbreviations :
A - Active        P - Passive      F - Aggregable I - Individual
S - Short-timeout L - Long-timeout N - InSync     O - OutofSync
C - Collecting    D - Distributing
X - State m/c expired              E - Default neighbor state

Actor details of all interfaces:
------------------------------------------------------------------------------
Intf       Aggr       Port  Port  State   System-ID         System Aggr Forwarding
           Name       Id    Pri                             Pri    Key  State
------------------------------------------------------------------------------
1/1/1      lag1(mc)   1     1     ALFOE   02:01:00:00:01:00 65534  1    lacp-block
1/1/2      lag2(mc)   2     1     IE      02:01:00:00:01:00 65534  2    up
```

13. On sw-edge1, verify that you can no longer reach the host 10.254.1.21.

```
ping 10.254.1.21
```

A factory default switch attempts to send a DHCP request on SVI1. VLAN 1 is untagged on all ports by default. When the upstream aggregation switch has the ports in lacp-blocked state, the factory default switch would not be able to get a DHCP based IP Address.

14. On -sw-agg1, enable LACP fallback again on LAG1.

```
interface lag 1
 lacp fallback
 exit
```

```
sw-agg1(config)# interface lag 1
sw-agg1(config-lag-if)# lacp fallback
sw-agg1(config-lag-if)# exit
```

15. On sw-agg1, verify the LACP state is **up** again.

```
show lacp interfaces
```

```
sw-agg1(config)# show lacp interfaces

State abbreviations :
```

```
A - Active         P - Passive      F - Aggregable I - Individual
S - Short-timeout L - Long-timeout N - InSync     O - OutofSync
C - Collecting    D - Distributing
X - State m/c expired              E - Default neighbor state

Actor details of all interfaces:
------------------------------------------------------------------------------
Intf       Aggr       Port  Port  State   System-ID         System Aggr Forwarding
           Name       Id    Pri                             Pri    Key  State
------------------------------------------------------------------------------
1/1/1      lag1(mc)   1     1     IE      02:01:00:00:01:00 65534  1    up
1/1/2      lag2(mc)   2     1     IE      02:01:00:00:01:00 65534  2    up
```

End of optional section, the next steps are required again!

Task 3: Configure access switch with LAG to VSX

Your colleague has asked you to connect the access switch sw-edge1 with a redundant connection (2 links) to the VSX aggregation cluster. In this task you will properly connect sw-edge1 to the VSX aggregation using uplink ports 1/1/27 and 1/1/28 with a LAG.

Objectives

- Configure an access switch with a LAG to connect to a VSX system.
- Verify the LAG status on a VSX peer device.

Steps

1. Use MGMT PC to open an SSH connection to **sw-edge1**.
2. Review the current STP uplink port states for ports 1/1/27 and 1/1/28.

```
show spanning-tree mst
```

```
sw-edge1(config)# show spanning-tree mst
#### MST0
Vlans mapped:  1-4094
Bridge         Address:64:e8:81:dd:81:80    priority:32768
Regional Root
Operational    Hello time(in seconds): 2  Forward delay(in seconds):15  Max-age(in
seconds):20  txHoldCount(in pps): 6
Configured     Hello time(in seconds): 2  Forward delay(in seconds):15  Max-age(in
seconds):20  Max-Hops:20
Root           Address:02:01:00:00:01:00  Priority:0
               Port:1/1/27                Path cost:2000
Regional Root  Address:64:e8:81:dd:81:80  Priority:32768
               Internal cost:0            Rem Hops:20

Port           Role           State      Cost       Priority   Type              BPDU-
Tx    BPDU-Rx    TCN-Tx     TCN-Rx
-------------- -------------- ---------- ---------- ---------- ----------------- -----
----- ...
1/1/27         Root           Forwarding 2000       128        P2P  Bound        112
80235      51         34
1/1/28         Alternate      Blocking   2000       128        P2P  Bound        141
3720       30         33
...
```

Question: What are the uplink port states?

Answer: Ports 1/1/27 is Root Port and 1/1/28 is Blocking.

3. Configure a new LACP LAG. Typically, the highest possible LAG ID is used for the uplink to the aggregation switch, in this example it will be LAG ID 256. Enable the LAG.

```
interface lag 256
 no shutdown
 lacp mode active
```

```
sw-edge1(config)# interface lag 256
sw-edge1(config-lag-if)#  no shutdown
sw-edge1(config-lag-if)#  lacp mode active
```

4. Configure the LAG as VLAN trunk, allow all VLANs.

```
vlan trunk native 1
vlan trunk allowed all
exit
```

```
sw-edge1(config-lag-if)#  vlan trunk native 1
sw-edge1(config-lag-if)#  vlan trunk allowed all
sw-edge1(config-lag-if)#  exit
```

5. Assign ports 1/1/27 and 1/1/28 to the LAG 256 and enable them.

```
interface 1/1/27,1/1/28
 lag 256
 no shutdown
 exit
```

```
sw-edge1(config)# interface 1/1/27,1/1/28
sw-edge1(config-if-<1/1/27,1/1/28>)#  lag 256
sw-edge1(config-if-<1/1/27,1/1/28>)#  no shutdown
sw-edge1(config-if-<1/1/27,1/1/28>)#  exit
```

6. Review the LLDP neighbors.

```
show lldp neighbor-info
```

```
sw-edge1(config)# show lldp neighbor-info

LLDP Neighbor Information
=========================

Total Neighbor Entries          : 5
Total Neighbor Entries Deleted  : 7
Total Neighbor Entries Dropped  : 0
Total Neighbor Entries Aged-Out : 7

LOCAL-PORT  CHASSIS-ID         PORT-ID                    PORT-DESC
TTL      SYS-NAME
-------------------------------------------------------------------------------
-----
...
1/1/27      b8:d4:e7:da:80:00  1/1/1                      1/1/1
120      sw-agg1
1/1/28      b8:d4:e7:d9:e5:00  1/1/1                      1/1/1
120      sw-agg2
...
```

Question: What is the Chassis-ID on the ports 1/1/27 and 1/1/28? Are they the same or different?

Answer: LLDP uses the base system MAC of sw-agg1 and sw-agg2. Each aggregation switch can be uniquely identified based on the LLDP information.

7. Review the LACP interface status, check the **Partner** section of the output.

```
show lacp interfaces
```

```
sw-edge1(config)# show lacp interfaces

State abbreviations :
A - Active        P - Passive      F - Aggregable I - Individual
S - Short-timeout L - Long-timeout N - InSync     O - OutofSync
C - Collecting    D - Distributing
X - State m/c expired              E - Default neighbor state

Actor details of all interfaces:
------------------------------------------------------------------------------
Intf       Aggr       Port  Port  State   System-ID         System Aggr Forwarding
           Name       Id    Pri                             Pri    Key  State
------------------------------------------------------------------------------
1/1/25     lag50                                                      down
1/1/26     lag50                                                      down
1/1/27     lag256     28    1     ALFNCD  64:e8:81:dd:81:80 65534  256  up
1/1/28     lag256     29    1     ALFNCD  64:e8:81:dd:81:80 65534  256  up


Partner details of all interfaces:
------------------------------------------------------------------------------
Intf       Aggr       Port  Port  State   System-ID         System Aggr
           Name       Id    Pri                             Pri    Key
------------------------------------------------------------------------------
1/1/25     lag50
1/1/26     lag50
1/1/27     lag256     1     1     ALFNCD  02:01:00:00:01:00 65534  1
1/1/28     lag256     1001  1     ALFNCD  02:01:00:00:01:00 65534  1
```

Question: What is the LACP System-ID on the ports 1/1/27 and 1/1/28? Are they the same or different?

Answer: They are the same. Since the Aggregation switches have been configured with a multi-chassis LAG, sw-agg1 and sw-agg2 will use the VSX system MAC in the LACP communication. Sw-edge1 believes it is connected to the same LACP partner on both ports.

Question: What is the Partner Port ID on the ports 1/1/27 and 1/1/28?

Answer: Port 1/1/27 is connected to port 1/1/1 on sw-agg1, Port 1/1/28 is connected to port 1/1/1 on sw-agg2. Since a unique port ID is required for the different LACP port members, the ports of the secondary switch (sw-agg2) start with 1000 + the port number.
Therefore port 1/1/28 is reported to be connected to port 1001 (base 1000 + port 1).

Task 4: Review VSX Active Gateway configuration and status

In this task you will review VSX Active Gateway configuration. This is the VSX active-active alternative to VRRP.

Objectives

- Review the Active Gateway configuration.
- Verify the Active Gateway operation.

Steps

1. Open an SSH connection to sw-agg1.
2. Review the configuration of SVI 11.

```
show running-config interface vlan 11
```

```
sw-agg1(config)# show running-config interface vlan 11
interface vlan11
    vsx-sync active-gateways
    ip address 10.1.11.2/24
    active-gateway ip mac 12:01:00:00:01:00
    active-gateway ip 10.1.11.1
    ip ospf 1 area 0.0.0.0
    ip helper-address 10.254.1.21
    exit
```

Question: What is the interface IP address?

Answer: 10.1.11.2/24

Question: What is the configured Active Gateway IP?

Answer: 10.1.11.1

Question: What is the configured Active Gateway MAC?

Answer: 12:01:00:00:01:00

VSX allows the administrator to easily run commands on the 'peer' VSX node.

3. While connected to sw-agg1, repeat the previous command, but add 'vsx-peer' to the end. This will execute the requested command on the VSX peer device, in this case on sw-agg2.

```
show running-config interface vlan 11 vsx-peer
```

```
sw-agg1(config)# show running-config interface vlan 11 vsx-peer
interface vlan11
    vsx-sync active-gateways
    ip address 10.1.11.3/24
    active-gateway ip mac 12:01:00:00:01:00
    active-gateway ip 10.1.11.1
```

```
    ip ospf 1 area 0.0.0.0
    ip helper-address 10.254.1.21
    exit
```

Question: What is the interface IP address?

Answer: 10.1.11.3/24. On sw-agg2, this is indeed the configured IP address.

Question: Are there any differences in the Active Gateway configuration between sw-agg1 and sw-agg2?

Answer: No, they both have the same virtual MAC and IP configured. Active Gateway is not an active-standby protocol, such as VRRP. With Active Gateway, both aggregation switches actively forward traffic that is sent to the virtual MAC.

For Active Gateway, there is no 'standard' MAC range available like what you saw for VRRP. Therefore, you manually configure a private MAC address as the Active Gateway MAC.

4. On PC1, check the ARP cache.

```
arp -a 10.1.11.1
```

```
C:\Users\student>arp -a 10.1.11.1

Interface: 10.1.11.50 --- 0xe
  Internet Address      Physical Address      Type
  10.1.11.1             12-01-00-00-01-00     dynamic
```

Question: What is the MAC address of the default gateway IP 10.1.11.1?

Answer: 12:01:00:00:01:00. This is the configured Active Gateway MAC of the aggregation switches.

5. Save the configuration on all switches.

```
write memory
```

Optional Task 5: Verify the VSX failover

The customer is pleased with the active-active VSX cluster configuration, but they would like to see how a switch failure is handled by the VSX system. In this task you will reboot each of the aggregation switches and verify the network connectivity using PC1.
This will allow you to demonstrate:

- The MCLAG failover for Layer2 traffic.
- The Active Gateway failover for the Layer3 traffic.
- The OSPF routing failover for the uplink routed traffic.

Objectives

- Verify the VSX system failover

Steps

1. On sw-edge1, verify the LACP state of the LAG256, both member ports should be up.

```
show lacp interfaces
```

```
sw-edge1(config)# show lacp interfaces

State abbreviations :
A - Active        P - Passive      F - Aggregable I - Individual
S - Short-timeout L - Long-timeout N - InSync     O - OutofSync
C - Collecting    D - Distributing
X - State m/c expired              E - Default neighbor state

Actor details of all interfaces:
------------------------------------------------------------------------------
Intf       Aggr       Port  Port  State   System-ID         System Aggr Forwarding
           Name       Id    Pri                             Pri    Key  State
------------------------------------------------------------------------------
1/1/25     lag50                                                      down
1/1/26     lag50                                                      down
1/1/27     lag256     28    1     ALFNCD  64:e8:81:dd:81:80 65534  256  up
1/1/28     lag256     29    1     ALFNCD  64:e8:81:dd:81:80 65534  256  up


Partner details of all interfaces:
------------------------------------------------------------------------------
Intf       Aggr       Port  Port  State   System-ID         System Aggr
           Name       Id    Pri                             Pri    Key
------------------------------------------------------------------------------
1/1/25     lag50
1/1/26     lag50
1/1/27     lag256     1     1     ALFNCD  02:01:00:00:01:00 65534  1
1/1/28     lag256     1001  1     ALFNCD  02:01:00:00:01:00 65534  1
```

2. On sw-edge1, enable terminal-monitor. Make sure your are connected using MGMT PC and an SSH session, not the console connection!

```
terminal-monitor
```

```
sw-edge1(config)# terminal-monitor
Terminal-monitor is enabled successfully
```

3. On PC1, start a continuous ping to 10.254.1.21.

```
ping 10.254.1.21 -t
```

4. On sw-agg1, save the configuration.

```
write mem
```

5. Reboot the switch, confirm the reboot with **y**.

```
boot system
```

6. On PC1, verify the ping continues.
7. On sw-edge1, review the events displayed on your terminal. You should see STP root failover occurs to sw-agg2 (priority 4096) and the LAG failover (due to port 1/1/27 link status down).

```
2022-10-17T05:36:10.586174-0400 hpe-mstpd[3243] <INFO>
Event|2008|LOG_INFO|CDTR|1|CIST starved for a BPDU Rx on port lag256 from 0:020100-
000100
2022-10-17T05:36:10.586566-0400 hpe-mstpd[3243] <INFO> Event|2006|LOG_INFO|CDTR|1|CST
- Root changed from 0: 02:01:00:00:01:00 to 32768: 64:e8:81:dd:81:80
2022-10-17T05:36:15.978772-0400 intfd[718] <INFO> Event|404|LOG_INFO|UKWN|1|Link
status for interface 1/1/27 is down
2022-10-17T05:36:15.998591-0400 lacpd[3150] <INFO> Event|1321|LOG_INFO|CDTR|1|LAG 256
State change for interface 1/1/27: Actor state: ALFO, Partner state ALFNCD
2022-10-17T05:36:16.205704-0400 hpe-mstpd[3243] <INFO> Event|2006|LOG_INFO|CDTR|1|CST
- Root changed from 32768: 64:e8:81:dd:81:80 to 4096: 02:01:00:00:01:00
```

8. On sw-edge1, review the events during the sw-agg1 boot process.

```
2022-10-17T05:37:36.539034-0400 hpe-mstpd[3243] <INFO> Event|2006|LOG_INFO|CDTR|1|CST
- Root changed from 4096: 02:01:00:00:01:00 to 0: 02:01:00:00:01:00
2022-10-17T05:38:07.161997-0400 lldpd[3153] <INFO> Event|106|LOG_INFO|CDTR|1|LLDP
neighbor b8:d4:e7:da:80:00 deleted on 1/1/27

2022-10-17T05:38:49.809988-0400 intfd[718] <INFO> Event|403|LOG_INFO|UKWN|1|Link
status for interface 1/1/27 is up
2022-10-17T05:38:49.827259-0400 lacpd[3150] <WARN> Event|1310|LOG_WARN|CDTR|1|Partner
is out of sync for interface 1/1/27 LAG sport: 2. Actor state: ALFOX, partner state
PSFO
2022-10-17T05:38:49.828560-0400 lacpd[3150] <INFO> Event|1321|LOG_INFO|CDTR|1|LAG 256
State change for interface 1/1/27: Actor state: ALIO, Partner state PLIO
2022-10-17T05:38:49.833016-0400 lacpd[3150] <INFO> Event|1321|LOG_INFO|CDTR|1|LAG 256
State change for interface 1/1/27: Actor state: ALFO, Partner state PLIO
2022-10-17T05:38:49.836836-0400 lacpd[3150] <INFO> Event|1321|LOG_INFO|CDTR|1|LAG 256
State change for interface 1/1/27: Actor state: ALFO, Partner state PLFO
2022-10-17T05:38:49.839875-0400 lacpd[3150] <INFO> Event|1321|LOG_INFO|CDTR|1|LAG 256
State change for interface 1/1/27: Actor state: ALFOE, Partner state PLFO
2022-10-17T05:38:49.843065-0400 lacpd[3150] <INFO> Event|1321|LOG_INFO|CDTR|1|LAG 256
State change for interface 1/1/27: Actor state: ALFOE, Partner state PSFO
2022-10-17T05:38:49.846235-0400 lacpd[3150] <INFO> Event|1321|LOG_INFO|CDTR|1|LAG 256
State change for interface 1/1/27: Actor state: ALFOX, Partner state PSFO
```

```
2022-10-17T05:38:49.853395-0400 lacpd[3150] <INFO> Event|1309|LOG_INFO|CDTR|1|Partner
is detected for interface 1/1/27 LAG 256 : 65534,02:01:00:00:01:00. Actor state:
ALFOX, partner state ALFOX
2022-10-17T05:38:49.853564-0400 lacpd[3150] <INFO> Event|1321|LOG_INFO|CDTR|1|LAG 256
State change for interface 1/1/27: Actor state: ALFOX, Partner state ALFOX
2022-10-17T05:38:49.856348-0400 lacpd[3150] <INFO> Event|1321|LOG_INFO|CDTR|1|LAG 256
State change for interface 1/1/27: Actor state: ALFO, Partner state ALFOX
2022-10-17T05:38:49.860439-0400 lacpd[3150] <INFO> Event|1321|LOG_INFO|CDTR|1|LAG 256
State change for interface 1/1/27: Actor state: ALFO, Partner state ALFO
2022-10-17T05:38:51.454255-0400 lacpd[3150] <INFO> Event|1321|LOG_INFO|CDTR|1|LAG 256
State change for interface 1/1/27: Actor state: ALFN, Partner state ALFO
2022-10-17T05:38:51.533436-0400 lacpd[3150] <INFO> Event|1321|LOG_INFO|CDTR|1|LAG 256
State change for interface 1/1/27: Actor state: ALFN, Partner state ALFN
2022-10-17T05:38:51.541340-0400 lacpd[3150] <INFO> Event|1321|LOG_INFO|CDTR|1|LAG 256
State change for interface 1/1/27: Actor state: ALFNC, Partner state ALFN
2022-10-17T05:38:51.548912-0400 lacpd[3150] <INFO> Event|1321|LOG_INFO|CDTR|1|LAG 256
State change for interface 1/1/27: Actor state: ALFNC, Partner state ALFNC
2022-10-17T05:38:51.555263-0400 lacpd[3150] <INFO> Event|1321|LOG_INFO|CDTR|1|LAG 256
State change for interface 1/1/27: Actor state: ALFNCD, Partner state ALFNC
2022-10-17T05:38:51.565517-0400 lacpd[3150] <INFO> Event|1321|LOG_INFO|CDTR|1|LAG 256
State change for interface 1/1/27: Actor state: ALFNCD, Partner state ALFNCD
2022-10-17T05:38:51.868979-0400 lldpd[3153] <INFO> Event|104|LOG_INFO|CDTR|1|LLDP
neighbor b8:d4:e7:da:80:00 added on 1/1/27
```

9. On sw-edge1, disable the terminal monitor logging.

```
no terminal-monitor
```

NOTE: The remote lab uses the CX Simulator for the Core Router. Therefore, there is no direct, physical connection between the sw-agg1 and the core router. This means that the link down of the sw-agg1 (due to the reboot), will not be observed immediately by the core router and it needs to rely on OSPF timers to expire. Therefore, the failover time in this lab will be higher than a production environment with directly connected devices.

You have completed the lab!

Lab 05 Stacking with VSF

Overview

In the previous lab, your senior network colleague prepared the aggregation switches and you explored how VSX is used on the aggregation layer to build a redundant network. Now your colleague wants you to configure the access layer switches into a VSF stack and connect them with a LAG to the VSX aggregation layer.

Using a VSF stack in the access layer has several advantages:

Reduced number of management points

When a customer requires four switches on a floor, they must remember to connect to the correct switch management IP when some configuration change is required. This may require an additional lookup to find the switch that is connected to a wall outlet. When the four switches on a floor are connected into a VSF stack, the number of management points is reduced by a factor four.

Reduced number of uplink ports

A VSF stack of four switches operates as one switch. They can use two physical uplinks to the aggregation layer. If four individual switches are used, each switch needs its own set of two uplinks or Spanning-Tree must be used to offer redundancy.

Easy device replacement

In the unfortunate event of a switch failure, using a single switch approach means a new switch must be deployed and the configuration restore procedure must be followed.
On the other hand, if a VSF stack is used, the new switch only requires adding to the existing VSF stack. The configuration of the ports will automatically be assigned to the new member.

During this lab you will initially lose the configuration of sw-edge2 since it will join sw-edge1 VSF stack.

Later in the lab, you will also test the VSF automatic stacking. At that point, you will again erase the switch configuration. This is fine since you will be connecting the switches to Aruba Central in the next lab. You will then use Aruba Central to push to correct configuration to the switches again.

Objectives

- Configure VSF.
- Understand auto-VSF.
- Verify VSF operation.
- Configure split-brain protection with VSF.

Task 1: Prepare the topology

In this task you will apply the correct start-up configuration for this lab on sw-agg1 and sw-agg2.

In the upcoming tasks, you will configure the two edge switches into a single VSF stack. Once you do this, the switches can form a LAG with four member ports to the VSX cluster.

In this task you will prepare the VSX cluster, so all ports to the edge switches belong to the same LAG (LAG 1).

You will make the sw-edge2 factory default in the next tasks. However, in the remote lab, a Zero Touch Provisioning (ZTP) DHCP server provides the initial configuration file. This would break the auto-join VSF activity. Therefore, you will disable ZTP for the sw-edge2.

Objectives

- Reconfigure the VSX system to use a LAG with 4 member ports.
- Disable the ZTP for sw-edg2.

Steps

Configure Aggregation VSX

The configuration in the next steps will assign ports 1/1/1 and 1/1/2 on *both* sw-agg1 and sw-agg2 to LAG1.

1. Use MGMT PC to open an SSH connection to **sw-agg1**.
2. Configure port 1/1/2 to become member of LAG1 and enable the port.

```
interface 1/1/2
 no lag 2
 lag 1
 no shutdown
 exit
```

```
sw-agg1(config)# interface 1/1/2
sw-agg1(config-if)#  no lag 2
sw-agg1(config-if)#  lag 1
sw-agg1(config-if)#  no shutdown
sw-agg1(config-if)#  exit
```

3. Use MGMT PC to open an SSH connection to **sw-agg2**.
4. Configure port 1/1/2 to become member of LAG1 and enable the port.

```
interface 1/1/2
 no lag 2
 lag 1
 no shutdown
 exit
```

```
sw-agg2(config)# interface 1/1/2
sw-agg2(config-if)#  no lag 2
```

```
sw-agg2(config-if)#  lag 1
sw-agg2(config-if)#  no shutdown
sw-agg2(config-if)#  exit
```

Disable ZTP for sw-edge2

On MGMT PC you will rename the sw-edge2 ZTP file in the TFTP folder. This will prevent the sw-edge2 from loading the start configuration.

1. On MGMT PC open the TFTP folder on the desktop.
2. Rename *cx-ztp-sw-edge2.cfg* file to ***no-****cx-ztp-sw-edge2.cfg*. When sw-edge2 attempts to download the file, it will not find it anymore. The ZTP process will not be performed, and the switch will automatically join the VSF stack in this lab.

Task 2: Create a VSF Stack using Automatic Join of New Member

In this task sw-edge1 and sw-edge2 will become part of the same VSF stack. Sw-edge1 will be used as VSF member 1, while sw-edge2 will be configured to become VSF member 2.

Every AOS-CX access switch is by default part of a VSF stack (of 1 switch) and assigned member ID 1.
The original member 1 (sw-edge1 in this setup) will keep its configuration in this task. The result will be that sw-edge1 simply stays in the network with the same configuration, but it will receive more ports than it previously had. This shows how an existing customer can extend any existing (VSF capable) AOS-CX switch to a stack without impacting the existing configuration.

You will use the auto-join function to connect sw-edge2 to the VSF stack. This means a new member can be added to an existing stack without having to access the new member device. Factory default switches are enabled for the Auto-join function by default. You only need to configure the VSF links on the existing VSF system. The existing VSF system will start sending VSF hello messages on these VSF links. The factory default switch is automatically listening for these special VSF hello messages on the predefined VSF ports and it will automatically reboot and join the VSF stack.

Objectives

- Understand the concept of VSF member IDs.
- Understand the concept of VSF links.
- Understand the auto-join VSF feature.
- Verify the auto-join VSF.
- Verify the VSF operation.

Steps

In the next steps you will review the default VSF status and configure sw-edge1 with a VSF link to support new member devices.

1. Open an SSH connection to sw-edge1.
2. Review the current VSF status.

```
show vsf
```

```
sw-edge1# show vsf

Force Autojoin              : Disabled
Autojoin Eligibility Status: Not Eligible
MAC Address                 : 64:e8:81:dd:81:80
Secondary                   :
Topology                    : Standalone
Status                      : No Split
Split Detection Method      : None


Mbr Mac Address         type           Status
ID
--- ------------------ -------------- ---------------
```

```
1   64:e8:81:dd:81:80   JL666A          Conductor
```

Question: What is the current member ID (Mbr ID)?

Answer: 1

Question: What is the current device Status?

Answer: Conductor.

Question: What does the secondary option mean in the output?

Answer: In a VSF Stack, one switch (the primary/conductor) controls the control plane and management plane processes, such as LACP, STP, or the SSH server. When the conductor goes offline, due to a power failure for example, the secondary switch should resume the conductor state. This other switch is configured using the secondary option. Without a configured secondary, the VSF stack would be offline when the conductor is not active, and therefore, you should always have a secondary switch!

Question: Has any member been configured as the Secondary?

Answer: No, by default no secondary is configured.

Question: What is the default *Split Detection Method*?

Answer: None. If there is a split stack condition, both fragments would remain active.

3. Review the default VSF running-configuration.

```
show running-config vsf
```

Example output, your output may be different.

```
sw-edge1# show running-config vsf
vsf member 1
    type jl666a
```

Question: How many VSF members do you see in the configuration?

Answer: Only one.

Question: What does the JLxxx number in the line with “type jl..” represent?

Answer: Each Aruba AOS-CX switch model has a unique product code. It indicates whether the switch has 24 or 48 ports, Power over Ethernet (PoE) ports, Smart Rate ports etc. By setting the product code, AOS-CX knows how many ports to provision and which features, such as PoE, are available on the VSF member.

Visit https://www.arubanetworks.com for a complete list of products.
Visit https://www.arubanetworks.com/assets/ds/DS_6300Series.pdf for a list of the 6300 switch series.

4. On the **sw-edge1**, configure the *future* member 2 to become the **secondary** (standby) control plane in the VSF stack. Confirm the message with **y**.

> **NOTE:** This step has no impact on your current setup, since there is no member 2 yet. This command will only take effect when member 2 joins the VSF stack!

```
vsf secondary-member 2
```

```
sw-edge1(config)# vsf secondary-member 2
This will save the configuration and reboot the specified switch.
Do you want to continue (y/n)? y
```

5. Configure the split stack detection to use the OOBM mgmt interface.

```
vsf split-detect mgmt
```

```
sw-edge1(config)# vsf split-detect mgmt
```

> **NOTE:** If there is no OOBM network available, you can connect a cable directly from the commander to the standby member OOBM port to support split stack detection.

6. Verify the configuration changes. A secondary should be present and the split method should be set to *mgmt* now.

```
show vsf
```

```
sw-edge1(config)# show vsf

Force Autojoin              : Disabled
Autojoin Eligibility Status: Not Eligible
MAC Address                 : 64:e8:81:dd:81:80
Secondary                   : 2
Topology                    : Standalone
Status                      : Active Fragment
Split Detection Method      : mgmt


Mbr Mac Address         type           Status
ID
--- ------------------ -------------- ---------------
1   64:e8:81:dd:81:80   JL666A         Conductor
```

7. Enter the VSF member 1 context.

```
vsf member 1
```

```
sw-edge1(config)# vsf member 1
```

8. Configure VSF link 1 with port 1/1/26.

```
link 1 1/1/26
```

```
sw-edge1(vsf-member-1)# link 1 1/1/26
```

9. Configure VSF link 2 with port 1/1/25.

```
link 2 1/1/25
```

```
sw-edge1(vsf-member-1)# link 2 1/1/25
```

10. Exit the VSF context.

```
exit
```

NOTE:
You will use auto-VSF to automatically establish the VSF system. Auto-VSF requires dedicated VSF link-to-port mappings. Make sure you have mapped the ports correctly:
VSF link 1 to port 1/1/26
VSF link 2 to port 1/1/25

11. Verify the configured VSF link to port assignment. Previously ports 1/1/25 and 1/1/26 were configured as member ports of LAG 50.

```
show vsf link
```

```
sw-edge1(config)# show vsf link

 VSF Member 1

     Link       Peer    Peer
Link State      Member  Link   Interfaces
---- ---------- ------- ------ --------------------------
1    down       0       0      1/1/26
2    down       0       0      1/1/25
```

12. Review the port 1/1/25 configuration.

```
show running-config interface 1/1/25
```

```
sw-edge1(config)# show running-config interface 1/1/25
interface 1/1/25
    no shutdown
    exit
```

Question: What happened with the port configuration?

Answer: When a port is assigned to a VSF link, all the existing settings are removed.

13. Review the interface brief output.

```
show interface brief
```

```
sw-edge1(config)# show interface brief
--------------------------------------------------------------------------------
-------------------------
Port      Native  Mode   Type          Enabled Status  Reason                Speed
Description
          VLAN                                                                (Mb/s)
--------------------------------------------------------------------------------
-------------------------
1/1/1     11      access 1GbT          yes     up                            1000
pc1
...
1/1/24    1       access 1GbT          yes     down    Waiting for link      --
--
1/1/25    --      VSF    SFP+DAC1      yes     up                            10000
sw-edge2
1/1/26    --      VSF    SFP+DAC1      yes     up                            10000
--
1/1/27    1       trunk  SFP+DAC1      yes     up                            10000
sw-agg1
1/1/28    1       trunk  SFP+DAC1      yes     up                            10000
sw-agg2
...
```

Question: What are the port modes you have seen so far?

Answer: Access, trunk, and routed.

Question: What is the new port mode for the VSF linked ports, such as 1/1/25 and 1/1/26?

Answer: They are assigned a new port mode: *VSF*. This indicates the port is used for a stacking link.

14. Review the current VSF links.

```
show vsf link
```

```
sw-edge1(config)# show vsf link

 VSF Member 1

     Link       Peer    Peer
Link State      Member  Link   Interfaces
---- ---------- ------- ------ --------------------------
1    down       0       0      1/1/26
2    down       0       0      1/1/25
```

15. Review the current VSF topology with a single switch.

```
show vsf topology
```

```
sw-edge1(config)# show vsf topology
```

```
 Conductor
 +-------+
 |   1   |
 +-------+
```

Perform factory reset of sw-edge2

Sw-edge2 will be factory reset. In the default state, it will listen for VSF frames and it can automatically join the VSF stack.

16. Open a console connection to sw-edge2 and login using admin / Aruba123!
17. Zeroize the sw-edge2 and confirm the operation.

```
erase all zeroize
```

```
sw-edge2# 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.
sw-edge2#
 Oct 17 06:50:47 hpe-mgmtmd[3776113]: RebootLibPh1: Reboot reason: Reboot requested
by user
```

The switch will now reboot and remove all locally stored configurations and log files.

Review the VSF Join on sw-edge1

While sw-edge2 performs the zeroize operation, you can review the status on sw-edge1.

18. On sw-edge1, review the VSF status.

```
show vsf
```

19. Start a repeat sequence. This will automatically repeat the previous command with the configured delay. This way, you don't need to repeat the command yourself while the sw-edge2 is rebooting and auto-joining the VSF system.

```
repeat delay 10
```

```
sw-edge1(config)# repeat delay 10

**************************************
Iteration : 1 Command : show vsf
**************************************

Force Autojoin             : Disabled
Autojoin Eligibility Status: Not Eligible
MAC Address                : 64:e8:81:dd:81:80
Secondary                  : 2
```

```
Topology                    : Standalone
Status                      : Active Fragment
Split Detection Method      : mgmt


Mbr Mac Address         type           Status
ID
--- ------------------ -------------- ---------------
1   64:e8:81:dd:81:80   JL666A         Conductor
...
```

```
...
*************************************
Iteration : 26 Command : show vsf
*************************************

Force Autojoin              : Disabled
Autojoin Eligibility Status: Not Eligible
MAC Address                 : 64:e8:81:dd:81:80
Secondary                   : 2
Topology                    : Standalone
Status                      : Active Fragment
Split Detection Method      : mgmt


Mbr Mac Address         type           Status
ID
--- ------------------ -------------- ---------------
1   64:e8:81:dd:81:80   JL666A         Conductor
2                       JL666A         Not Present
...
```

```
...
*************************************
Iteration : 38 Command : show vsf
*************************************

Force Autojoin              : Disabled
Autojoin Eligibility Status: Not Eligible
MAC Address                 : 64:e8:81:dd:81:80
Secondary                   : 2
Topology                    : Ring
Status                      : Active Fragment
Split Detection Method      : mgmt


Mbr Mac Address         type           Status
ID
--- ------------------ -------------- ---------------
1   64:e8:81:dd:81:80   JL666A         Conductor
2   8c:85:c1:49:20:c0   JL666A         Standby Booting


*************************************
```

```
Iteration : 39 Command : show vsf
************************************

Force Autojoin              : Disabled
Autojoin Eligibility Status: Not Eligible
MAC Address                 : 64:e8:81:dd:81:80
Secondary                   : 2
Topology                    : Ring
Status                      : No Split
Split Detection Method      : mgmt


Mbr Mac Address         type           Status
ID
--- ------------------ -------------- ---------------
1   64:e8:81:dd:81:80   JL666A         Conductor
2   8c:85:c1:49:20:c0   JL666A         Standby
```

20. After a few minutes, the sw-edge2 should have completed the join. When you see two entries in the show vsf repeated output, you can stop the repeat using the CTRL-C key sequence

```
<CTRL>-c
```

21. On sw-edge1, review the VSF running configuration.

```
show run vsf
```

```
sw-edge1(config)# show run vsf
vsf split-detect mgmt
vsf secondary-member 2
vsf member 1
    type jl666a
    link 1 1/1/26
    link 2 1/1/25
vsf member 2
    type jl666a
    link 1 2/1/25
    link 2 2/1/26
```

Question: Do you see the member 2 in the configuration?

Answer: Yes, member 2 was automatically added with the correct type to the configuration.

Question: Did you have to configure VSF links for member 2?

Answer: No. These member 2 VSF links were automatically added to the configuration. Only the VSF links on the existing stack (the conductor in this case) had to be configured.

22. Check the VSF topology with member 2 active.

```
show vsf topology
```

```
sw-edge1(config)# show vsf topology
```

```
 Conductor    Standby
 +-------+    +-------+
 |   1   |1==1|   2   |
 +-------+    +-------+
     2            2
     +============+
```

23. Check the VSF links that have been configured.

```
show vsf link
```

```
sw-edge1(config)# show vsf link

 VSF Member 1

     Link       Peer    Peer
Link State      Member  Link   Interfaces
---- ---------- ------- ------ --------------------------
1    up         2       1      1/1/26
2    up         2       2      1/1/25


 VSF Member 2

     Link       Peer    Peer
Link State      Member  Link   Interfaces
---- ---------- ------- ------ --------------------------
1    up         1       1      2/1/25
2    up         1       2      2/1/26
```

24. Use the **show vsf detail** command to see serial and MAC details of the member devices.

```
show vsf detail
```

```
sw-edge1(config)# show vsf detail
VSF Stack
        MAC Address                  : 64:e8:81:dd:81:80
        Secondary                    : 2
        Topology                     : ring
        Status                       : No Split
        Split Detection Method       : mgmt
        Software Version             : FL.10.09.1040
        Force Autojoin               : Disabled
        Autojoin Eligibility Status  : Not Eligible
        Autojoin Ineligibility Reason: Configuration changes detected
        Name                         : Aruba-VSF-6300
        Contact                      :
        Location                     :

Member ID                            : 1
        MAC Address                  : 64:e8:81:dd:81:80
        Type                         : JL666A
```

```
        Model                         : 6300F 24-port 1GbE Class 4 PoE and 4-port
SFP56 Switch
        Status                        : Conductor
        ROM Version                   : FL.01.09.0002
        Serial Number                 : SG09KN500B
        Uptime                        : 1 day, 22 hours, 43 minutes
        CPU Utilization               : 3%
        Memory Utilization            : 22%
        VSF Link 1                    : Up, connected to peer member 2, link 1
        VSF Link 2                    : Up, connected to peer member 2, link 2

Member ID                             : 2
        MAC Address                   : 8c:85:c1:49:20:c0
        Type                          : JL666A
        Model                         : 6300F 24-port 1GbE Class 4 PoE and 4-port
SFP56 Switch
        Status                        : Standby
        ROM Version                   : FL.01.09.0002
        Serial Number                 : SG11KN501V
        Uptime                        : 2 minutes
        CPU Utilization               : 5%
        Memory Utilization            : 11%
        VSF Link 1                    : Up, connected to peer member 1, link 1
        VSF Link 2                    : Up, connected to peer member 1, link 2
```

Console Access to VSF Member sw-edge2

In the next steps you will explore the console connection of the VSF stack members.

25. Use the lab dashboard to open a console connection to sw-edge2. Notice how the hostname prompt shows sw-edge1, even though you are connected to the console of the sw-edge2.

```
sw-edge1 login:
```

26. Login with admin/Aruba123!

```
sw-edge1 login: admin
Password:

standby#
```

Question: Do you see the hostname in the prompt? What does the prompt show on the standby switch?

Answer: No, the hostname is not shown, only the *standby* prompt.

27. Use ? to see a list of commands.

```
?
```

```
standby# ?
```

```
clear              Reset functions
diagnostics        Change diagnostic commands availability
exit               Exit current mode and change to previous mode
list               Print command list
member             VSF member selection
no                 Negate a command or set its defaults
page               Enable page break
show               Show running system information
start-shell        Start Bash shell
terminal-monitor   Enables Terminal-monitor
top                Top command
vsf-factory-reset  Erase all customer data and reset the switch to factory
                   defaults.
```

Question: Do you have full CLI access on the standby device?

Answer: No, only minimal diagnostics and show commands are available.

You may leave the console connection to sw-edge2 open for now.

This concludes the member auto join VSF.

Task 3: Connect VSF Access Stack to VSX Aggregation with MCLAG

In this task you will connect the VSF stack with a LAG of four uplinks to the VSX aggregation switches. In Task 1 you have prepared the VSX system with a single MCLAG (LAG1 on the VSX) of four member ports.

Typically, an Access VSF stack would only have two uplink ports: one uplink port on the Conductor and one uplink port on the Secondary. In this lab environment, each access switch connects with two ports to the aggregation layer. This is why four uplink ports are bundled into the link-aggregation.

Sw-edge1 and sw-edge2 have their port 1/1/27 connected to sw-agg1, on ports 1/1/1 and 1/1/2 respectively.
Sw-edge1 and sw-edge2 have their port 1/1/28 connected to sw-agg2, on ports 1/1/1 and 1/1/2 respectively.

Objectives

- Configure a VSF stack with LAG to VSX aggregation.
- Use a LAG with member ports of different VSF members.

Steps

1. Open an SSH connection to sw-edge1.
2. Review the interface brief list.

```
show interface brief
```

```
sw-edge1(config)# show interface brief
--------------------------------------------------------------------------------
-------------------------
Port      Native  Mode   Type          Enabled Status  Reason                Speed
Description
          VLAN                                                                 (Mb/s)
--------------------------------------------------------------------------------
-------------------------
1/1/1     11      access 1GbT          yes     up                            1000
pc1
...
1/1/24    1       access 1GbT          yes     down    Waiting for link      --
--
1/1/25    --      VSF    SFP+DAC1      yes     up                            10000
sw-edge2
1/1/26    --      VSF    SFP+DAC1      yes     up                            10000
--
1/1/27    1       trunk  SFP+DAC1      yes     up                            10000
sw-agg1
1/1/28    1       trunk  SFP+DAC1      yes     up                            10000
sw-agg2
2/1/1     1       access 1GbT          yes     down    Waiting for link      --
--
```

```
...
2/1/24    1       access 1GbT           yes     down    Waiting for link       --
--
2/1/25    --      VSF    SFP+DAC1       yes     up                            10000
--
2/1/26    --      VSF    SFP+DAC1       yes     up                            10000
--
2/1/27    1       access SFP+DAC1       yes     up                            10000
--
2/1/28    1       access SFP+DAC1       yes     up                            10000
--
vlan1     --              --            no      down    Administratively down --
--
vlan3     --              --            yes     up                            --
--
lag50     1       trunk  --             no      down    --                    auto
--
lag256    1       trunk  --             yes     up      --                    20000
--
```

Question: In the output, do you notice ports that start with 2/1/ ?

Answer: Yes, these are the ports of the member 2 that are now under the management of this VSF system.

3. Check the LLDP neighbors.

```
show lldp neighbor-info
```

```
sw-edge1(config)# show lldp neighbor-info

LLDP Neighbor Information
=========================

Total Neighbor Entries          : 10
Total Neighbor Entries Deleted  : 2
Total Neighbor Entries Dropped  : 0
Total Neighbor Entries Aged-Out : 2

LOCAL-PORT  CHASSIS-ID         PORT-ID      PORT-DESC                    TTL
SYS-NAME
-------------------------------------------------------------------------------
-----
...
1/1/27      b8:d4:e7:da:80:00  1/1/1        1/1/1                        120      sw-
agg1
1/1/28      b8:d4:e7:d9:e5:00  1/1/1        1/1/1                        120      sw-
agg2
...
2/1/27      b8:d4:e7:da:80:00  1/1/2        1/1/2                        120      sw-
agg1
2/1/28      b8:d4:e7:d9:e5:00  1/1/2        1/1/2                        120      sw-
agg2
...
```

Question: On what ports do you see sw-agg1 and sw-agg2 as LLDP neighbors?

Answer: On ports 1/1/27 and 1/1/28, as well as 2/1/27 and 2/1/28.

4. Review the current LACP interface status.

```
show lacp interfaces
```

```
sw-edge1(config)# show lacp interfaces

State abbreviations :
A - Active        P - Passive      F - Aggregable I - Individual
S - Short-timeout L - Long-timeout N - InSync     O - OutofSync
C - Collecting    D - Distributing
X - State m/c expired              E - Default neighbor state

Actor details of all interfaces:
------------------------------------------------------------------------------
Intf       Aggr       Port  Port  State   System-ID         System Aggr Forwarding
           Name       Id    Pri                             Pri    Key  State
------------------------------------------------------------------------------
1/1/27     lag256     28    1     ALFNCD  64:e8:81:dd:81:80 65534  256  up
1/1/28     lag256     29    1     ALFNCD  64:e8:81:dd:81:80 65534  256  up


Partner details of all interfaces:
------------------------------------------------------------------------------
Intf       Aggr       Port  Port  State   System-ID         System Aggr
           Name       Id    Pri                             Pri    Key
------------------------------------------------------------------------------
1/1/27     lag256     1     1     ALFNCD  02:01:00:00:01:00 65534  1
1/1/28     lag256     1001  1     ALFNCD  02:01:00:00:01:00 65534  1
```

Question: What ports are missing from the LAG256 configuration?

Answer: The new ports 2/1/27 and 2/1/28 of member 2 are missing from the LAG256.

5. Now add ports 2/1/27 and 2/1/28 to the LAG256 and enable them.

```
interface 2/1/27,2/1/28
 lag 256
 exit
```

```
sw-edge1(config)# interface 2/1/27,2/1/28
sw-edge1(config-if-<2/1/27,2/1/28>)#  lag 256
sw-edge1(config-if-<2/1/27,2/1/28>)#  exit
```

6. Verify the LACP interface status. The Forwarding State for all four ports should be *up* now.

```
show lacp interfaces
```

```
sw-edge1(config)# show lacp interfaces

State abbreviations :
```

```
A - Active         P - Passive      F - Aggregable I - Individual
S - Short-timeout  L - Long-timeout N - InSync     O - OutofSync
C - Collecting     D - Distributing
X - State m/c expired               E - Default neighbor state

Actor details of all interfaces:
------------------------------------------------------------------------------
Intf       Aggr       Port  Port  State   System-ID         System Aggr Forwarding
           Name       Id    Pri                             Pri    Key  State
------------------------------------------------------------------------------
1/1/27     lag256     28    1     ALFNCD  64:e8:81:dd:81:80 65534  256  up
1/1/28     lag256     29    1     ALFNCD  64:e8:81:dd:81:80 65534  256  up
2/1/27     lag256     92    1     ALFNCD  64:e8:81:dd:81:80 65534  256  up
2/1/28     lag256     93    1     ALFNCD  64:e8:81:dd:81:80 65534  256  up


Partner details of all interfaces:
------------------------------------------------------------------------------
Intf       Aggr       Port  Port  State   System-ID         System Aggr
           Name       Id    Pri                             Pri    Key
------------------------------------------------------------------------------
1/1/27     lag256     1     1     ALFNCD  02:01:00:00:01:00 65534  1
1/1/28     lag256     1001  1     ALFNCD  02:01:00:00:01:00 65534  1
2/1/27     lag256     2     1     ALFNCD  02:01:00:00:01:00 65534  1
2/1/28     lag256     1002  1     ALFNCD  02:01:00:00:01:00 65534  1
```

Task 4: Split-brain detection test

Your customer is pleased with the VSF configuration and connectivity, but they have a concern about the VSF stacking cables. They would like to know what happens when one or more VSF links fail, and how that affects the network.

You plan a meeting with your senior colleague and the customer, where your colleague explains the impact:

A VSF stack will typically be connected in a ring topology. The ring topology provides protection against a link failure on any of the VSF links. This ensures there is no single point of failure.

However, in the unlikely event that a second VSF link fails, it is possible that there is a split, or fragmented stack, with a conductor in both fragments.

The member 1 will be the conductor in one fragment, while the configured secondary member will be the conductor in the other fragment.

This may be a challenge for the rest of the network, since there are two different instances of STP/LACP operating, but they are using the same system MAC address.

The solution is to make sure that the fragment with the secondary switch as the conductor will disable its ports; this will effectively remove it from the network. The split detection will provide a dedicated connection between the member 1 and the secondary to support this detection. If there is a fragment without the member 1 or the secondary member, the fragment will not have a conductor. As such there is no split-brain, but it also means the switches in this fragment will not be operational, due to the lack of control plane.

In this task you will explore the configuration and operation of a split-stack situation.

Objectives

- Configure the split stack detection.
- Verify the split stack operation.
- Verify the operation after reconnecting the fragments.

Steps

1. Open an SSH connection to sw-edge1 and review the VSF status.

```
show vsf
```

```
sw-edge1(config)# show vsf

Force Autojoin             : Disabled
Autojoin Eligibility Status: Not Eligible
MAC Address                : 64:e8:81:dd:81:80
Secondary                  : 2
Topology                   : Ring
Status                     : No Split
Split Detection Method     : mgmt
```

```
Mbr Mac Address         type           Status
ID
--- ------------------- -------------- ---------------
1   64:e8:81:dd:81:80   JL666A         Conductor
2   8c:85:c1:49:20:c0   JL666A         Standby
```

2. Enable the terminal monitor to see live event logging.

```
terminal-monitor filter vsf
```

```
sw-edge1(config)# terminal-monitor filter vsf
Terminal-monitor is enabled successfully
```

NOTE: The **terminal-monitor** command is not available on the console connection. If you receive an error when using **terminal-monitor**, make sure to use MGMT PC and open an *SSH* connection to sw-edge1.

3. Disable the VSF links by shutting down the physical member ports.

```
interface 1/1/25,1/1/26
 shutdown
 exit
```

```
sw-edge1(config)# interface 1/1/25,1/1/26
sw-edge1(config-if-vsf-<1/1/25,1/1/26>)#  shutdown
sw-edge1(config-if-vsf-<1/1/25,1/1/26>)# exit
```

4. Review the events that are displayed on your session.

```
2022-10-17T07:06:25.278928-0400 vsfd[713] <INFO> Event|9924|LOG_INFO|CDTR|1|VSF link
2 is down
2022-10-17T07:06:25.286822-0400 vsfd[713] <INFO> Event|9908|LOG_INFO|CDTR|1|Topology
is Standalone
2022-10-17T07:06:25.286981-0400 vsfd[713] <WARN> Event|9913|LOG_WARN|CDTR|1|Lost
member 2 with Loss of communication
2022-10-17T07:06:25.299888-0400 vsfd[713] <INFO> Event|9924|LOG_INFO|CDTR|1|VSF link
1 is down
2022-10-17T07:06:26.459501-0400 vsfd[713] <INFO> Event|9927|LOG_INFO|CDTR|1|Fragment
with conductor 1 is Active
```

5. Check vsf status again.

```
show vsf
```

```
sw-edge1(config)# show vsf

Force Autojoin             : Disabled
Autojoin Eligibility Status: Not Eligible
MAC Address                : 64:e8:81:dd:81:80
Secondary                  : 2
Topology                   : Standalone
Status                     : Active Fragment
Split Detection Method     : mgmt
```

```
Mbr Mac Address         type           Status
ID
--- ------------------- -------------- ---------------
1   64:e8:81:dd:81:80   JL666A         Conductor
2                       JL666A         In Other Fragment
```

Question: What is the status for member 2?

Answer: In Other Fragment.

6. Open or switch to the console connection on sw-edge2, login with admin/Aruba123!

```
sw-edge1 login: admin
Password:

Last login: 2022-10-17 11:00:17 from the console
User "admin" has logged in 1 time in the past 30 days
sw-edge1#
```

Question: Is the hostname prompt still *standby*?

Answer: No, since the secondary lost the conductor, it became the active conductor and assumed the conductor hostname, in this example sw-edge1.

7. On the console of sw-edge2, review the interface list.

```
show interface brief
```

```
sw-edge1# show interface brief
-------------------------------------------------------------------------------
-------------------------
Port      Native  Mode   Type          Enabled Status  Reason                Speed
Description
          VLAN                                                               (Mb/s)
-------------------------------------------------------------------------------
-------------------------
2/1/1     1       access 1GbT          yes     down    Disabled by VSF       --
--
2/1/2     1       access 1GbT          yes     down    Disabled by VSF       --
--
2/1/3     1       access 1GbT          yes     down    Disabled by VSF       --
--
2/1/4     1       access 1GbT          yes     down    Disabled by VSF       --
--
2/1/5     1       access 1GbT          yes     down    Disabled by VSF       --
--
2/1/6     1       access 1GbT          yes     down    Disabled by VSF       --
--
2/1/7     1       access 1GbT          yes     down    Disabled by VSF       --
--
2/1/8     1       access 1GbT          yes     down    Disabled by VSF       --
--
2/1/9     1       access 1GbT          yes     down    Disabled by VSF       --
--
```

```
2/1/10    1       access 1GbT           yes     down    Disabled by VSF        --
--
2/1/11    1       access 1GbT           yes     down    Disabled by VSF        --
--
2/1/12    1       access 1GbT           yes     down    Disabled by VSF        --
--
2/1/13    1       access 1GbT           yes     down    Disabled by VSF        --
--
2/1/14    1       access 1GbT           yes     down    Disabled by VSF        --
--
2/1/15    1       access 1GbT           yes     down    Disabled by VSF        --
--
2/1/16    1       access 1GbT           yes     down    Disabled by VSF        --
--
2/1/17    1       access 1GbT           yes     down    Disabled by VSF        --
--
2/1/18    1       access 1GbT           yes     down    Disabled by VSF        --
--
2/1/19    1       access 1GbT           yes     down    Disabled by VSF        --
--
2/1/20    1       access 1GbT           yes     down    Disabled by VSF        --
--
2/1/21    1       access 1GbT           yes     down    Disabled by VSF        --
--
2/1/22    1       access 1GbT           yes     down    Disabled by VSF        --
--
2/1/23    1       access 1GbT           yes     down    Disabled by VSF        --
--
2/1/24    1       access 1GbT           yes     down    Disabled by VSF        --
--
2/1/25    --      VSF    SFP+DAC1       yes     down    Waiting for link       --
--
2/1/26    --      VSF    SFP+DAC1       yes     down    Waiting for link       --
--
2/1/27    1       trunk  SFP+DAC1       yes     down    Disabled by VSF        --
--
2/1/28    1       trunk  SFP+DAC1       yes     down    Disabled by VSF        --
--
vlan1     --             --             no      down    Administratively down  --
--
vlan3     --             --             yes     down                           --
--
lag50     1       trunk  --             no      down    --                     auto
--
lag256    1       trunk  --             yes     down    --                     auto
--
```

Question: What is the status and reason text for the interfaces?

Answer: Disabled by VSF. This is the split brain detection that has disabled the ports on this fragment.

Question: What is the status for the VSF links?

Answer: Waiting for link.

Question: Why are the VSF links not disabled by VSF?

Answer: When the VSF links would be restored, this fragment will automatically detect that the links come up again and it will automatically rejoin the VSF stack by rebooting the current fragment. If the VSF links would have been disabled, this fragment would not be able to detect that the VSF links would have been restored.

8. On the console of sw-edge2, review VSF status.

```
show vsf
```

```
sw-edge1#  show vsf

Force Autojoin              : Disabled
Autojoin Eligibility Status: Not Eligible
MAC Address                 : 64:e8:81:dd:81:80
Secondary                   : 2
Topology                    : Standalone
Status                      : Inactive Fragment
Split Detection Method      : mgmt


Mbr Mac Address         type           Status
ID
--- ------------------- -------------- ---------------
1                       JL666A         In Other Fragment
2   8c:85:c1:49:20:c0   JL666A         Conductor
```

Question: What is the status of the current system?

Answer: The secondary has detected using the OOBM mgmt interface that the conductor is still active; it has marked itself the *Inactive Fragment*.

9. You may close the console connection to sw-edge2.

Restore the connection.

10. On sw-edge1, enable the ports 1/1/25 and 1/1/26 again.

```
interface 1/1/25,1/1/26
 no shutdown
 exit
```

```
sw-edge1(config)# interface 1/1/25,1/1/26
sw-edge1(config-if-vsf-<1/1/25,1/1/26>)#  no shutdown
sw-edge1(config-if-vsf-<1/1/25,1/1/26>)#  exit
```

11. Review the events on the terminal.

```
2022-10-17T07:14:04.971464-0400 vsfd[713] <INFO> Event|9923|LOG_INFO|CDTR|1|VSF link
2 is up
2022-10-17T07:14:04.979681-0400 vsfd[713] <INFO> Event|9923|LOG_INFO|CDTR|1|VSF link
1 is up
2022-10-17T07:14:07.978152-0400 vsfd[713] <INFO> Event|9916|LOG_INFO|CDTR|1|Member 2
elected as standby reason-Designated standby
```

```
2022-10-17T07:14:08.117673-0400 vsfd[713] <INFO> Event|9924|LOG_INFO|CDTR|1|VSF link
1 is down
2022-10-17T07:14:28.653101-0400 vsfd[713] <INFO> Event|9924|LOG_INFO|CDTR|1|VSF link
2 is down
```

Sw-edge2 will not reboot to join the VSF stack again.

12. On sw-edge1, you may disable the terminal monitor.

```
no terminal-monitor
```

```
sw-edge1(config)# no terminal-monitor
```

13. Wait a few minutes and confirm that member has joined the VSF stack as the standby system.

```
show vsf
```

Task 5: Conductor failover Test

Your customer is very interested in the VSF stacking technology, and they would like to know what happens when the VSF Conductor fails or reboots. In this task you will test this failure. This prepares you to respond to the customer questions.

The member 1 is the by default the active conductor in the VSF stack. If member 1 is down due to a reboot or power loss, the secondary assumes the conductor role. It is important that the secondary member ID is set in the configuration, which must be done manually.

In the following tasks you will explore auto-stacking. With auto-stacking, member 2 is automatically configured as the secondary switch in the VSF stack, so there is no need to configure it manually.

Objectives

- Verify the impact of the Conductor failover.

Prepare test connections

First you will initiate some ping tests from various systems to the VSF system.
There will be an out-of-band ping test and an in-band ping test.

Out-Of-Band

1. On sw-edge1, use **show int mgmt**, and take note of MAC.

```
show interface mgmt
```

Example output, your output may be different.

```
sw-edge1(config)# show interface mgmt
  Address Mode                  : static
  Admin State                   : up
  Link State                    : up
  Mac Address                   : 64:e8:81:dd:81:81
  IPv4 address/subnet-mask      : 10.251.1.4/24
  Default gateway IPv4          :
  IPv6 address/prefix           :
  IPv6 link local address/prefix: fe80::66e8:81ff:fedd:8181/64
  Default gateway IPv6          :
  Primary Nameserver            :
  Secondary Nameserver          :
  Tertiary Nameserver           :
```

OOBM MAC:_____

2. On MGMT PC, perform a ping to the OOBM mgmt interface (10.251.1.4) of the sw-edge1.

```
ping 10.251.1.4
```

3. On MGMT PC, check the arp table

```
arp -a 10.251.1.4
```

```
C:\Users\student>arp -a 10.251.1.4

Interface: 10.251.1.90 --- 0xf
  Internet Address      Physical Address      Type
  10.251.1.4            64-e8-81-dd-81-81     dynamic
```

Question: Does the ARP entry for 10.251.1.4 match the MAC address you have recorded in step 1?

Answer: Yes, this is the same MAC.

4. On MGMT PC, start a continuous ping to the OOBM IP 10.251.1.4.

```
ping 10.251.1.4 -t
```

In-Band

For the in-band test, you will use sw-agg1 and ping to the SVI 3 IP address (the access switch in-band management IP).

5. On sw-agg1 ping to 10.1.3.4.

```
ping 10.1.3.4
```

6. Check the ARP table entry for MAC.

```
show arp | include 10.1.3.4
```

```
sw-agg1(config)# show arp | include 10.1.3.4
10.1.3.4         64:e8:81:dd:81:80  vlan3        lag1                          reachable
```

Question: Take note of the MAC address of the IP 10.1.3.4

_____

7. On sw-agg1, now initiate a series of 100 pings. You need to disable paging to ensure the output does not stop after 1 page.

```
no page
ping 10.1.3.4 repetitions 100
```

```
sw-agg1(config)# no page
sw-agg1(config)# ping 10.1.3.4 repetitions 100
PING 10.1.3.4 (10.1.3.4) 100(128) bytes of data.
108 bytes from 10.1.3.4: icmp_seq=1 ttl=64 time=0.167 ms
108 bytes from 10.1.3.4: icmp_seq=2 ttl=64 time=0.166 ms
...
```

You are now ready to test the conductor failover and observe the impact on the MAC address of the stack.

Reboot conductor

8. On the VSF system (SSH connection to sw-edge1), reboot member 1.

```
vsf member 1 reboot
```

```
sw-edge1(config)# vsf member 1 reboot
The conductor switch will reboot and the standby will become the conductor.
Do you want to continue (y/n)? y
sw-edge1(config)#
 Oct 17 07:19:39 vsfd[713]: RebootLibPh1: Reboot reason: Reboot of conductor due to
Reboot requested by user
```

9. On MGMT PC, verify the ping continues to the OOBM IP using the same mgmt MAC.

```
Reply from 10.251.1.4: bytes=32 time<1ms TTL=64
Reply from 10.251.1.4: bytes=32 time<1ms TTL=64
Reply from 10.251.1.4: bytes=32 time<1ms TTL=64
Request timed out.
Request timed out.
Request timed out.
Reply from 10.251.1.4: bytes=32 time<1ms TTL=64
Reply from 10.251.1.4: bytes=32 time=10ms TTL=64
Reply from 10.251.1.4: bytes=32 time<1ms TTL=64
```

10. On MGMT PC, stop the ping and check the ARP table again.

```
arp -a 10.251.1.4
```

```
C:\Users\student>arp -a 10.251.1.4

Interface: 10.251.1.90 --- 0xf
  Internet Address      Physical Address      Type
  10.251.1.4            64-e8-81-dd-81-81     dynamic
```

11. On sw-agg1, verify the ping continues on the in-band IP using the same SVI MAC.

This demonstrates the Conductor role failover to the Secondary and the stability of the VSF MAC address during the failover.

Restore member 1 as conductor

Once the member 1 comes back online, it will not automatically claim the conductor role. This is fine in a production environment. In the lab environment, the next lab requires member 1 to be the conductor. Therefore, you will reboot the secondary. This will make member 1 the conductor again.

12. Re-open the SSH connection to sw-edge1 (this is the OOBM IP 10.251.1.4, but it is currently handled by the sw-edge2 as this is the current conductor).

13. Review the VSF status.

```
show vsf
```

```
sw-edge1# show vsf

Force Autojoin             : Disabled
Autojoin Eligibility Status: Not Eligible
MAC Address                : 64:e8:81:dd:81:80
Secondary                  : 2
```

```
Topology                    : Ring
Status                      : No Split
Split Detection Method      : mgmt


Mbr Mac Address         type           Status
ID
--- ------------------- -------------- ---------------
1   64:e8:81:dd:81:80   JL666A         Standby
2   8c:85:c1:49:20:c0   JL666A         Conductor
```

14. Reboot member 2 to make member 1 conductor. You may also use the 'vsf switchover' command.

```
vsf member 2 reboot
```

or

```
vsf switchover
```

```
sw-edge1# vsf switchover
This will cause an immediate switchover to the standby
and the conductor will reboot.
Do you want to continue (y/n)? y
sw-edge1#
 Oct 17 07:23:34 vsfd[732]: RebootLibPh1: Reboot reason: Reboot of conductor due to
VSF switchover requested
```

IMPORTANT: Wait for the member 2 to join the VSF stack again before moving to the next task.

Task 6: Automatic Conductor VSF stack formation

Your customer is pleased with the auto-join function to add a factory default switch to an existing stack.
However, the customer is thinking about deploying a lot of switches, and they are wondering if the initial configuration of the complete VSF stack can also be done by their technician when a new location needs to be deployed. This would include the initial conductor configuration, but without accessing the configuration of the factory default member 1. After consulting with your senior colleague, you learn that the auto-stacking feature can also be used to configure to the conductor.

In this task you will see an easy method to form a VSF stack, including the initial assignment of the conductor in the stack. This procedure must be applied to factory default switches. This means the first steps in this task will be to make the access switches factory default again. Auto-stacking requires the use of dedicated ports on the switches. This is an extract from the VSF configuration guide:

https://www.arubanetworks.com/techdocs/AOS-CX/10.09/HTML/vsf/Content/fir-int.htm

Reserved interfaces for auto-stacking

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.
The following table shows the list of reserved auto-stacking interfaces based on the product type and
platform

Reserved Interfaces

Platform	SKU Part Number	Ports reserved for Auto-Stacking
6300	JL658A JL660A JL662A JL664A JL666A JL668A	25 and 26
6300	JL659A JL661A JL663A JL665A JL667A JL762A	49 and 50
6200	JL724A JL725A	25 and 26
6200	JL726A JL727A JL728A	49 and 50

Example connection diagram for four switches:

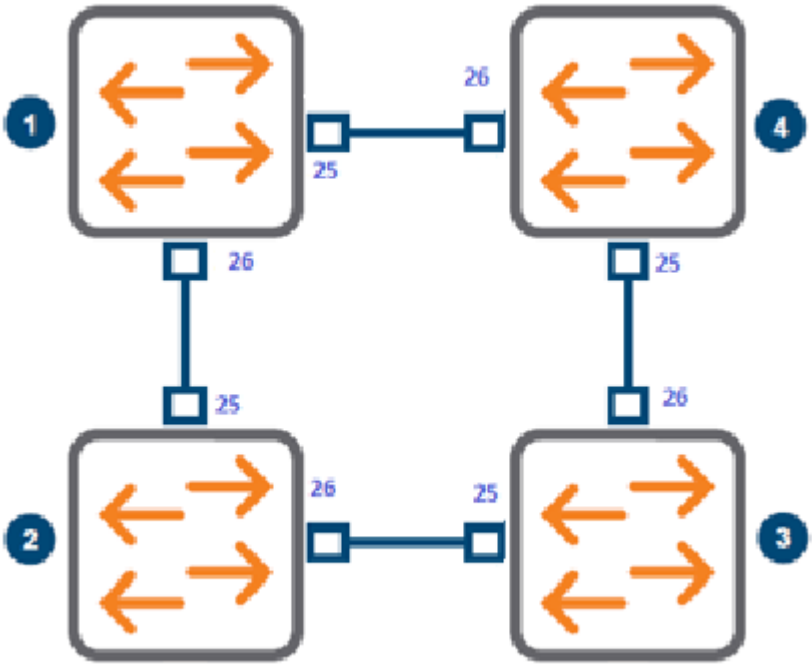

NOTE: Use the higher port number on the member 1 (either 26 or 50) to connect it to the lower port number on the member 2 (either 25 or 49). The rest of the connections should automatically follow this logic of connecting the higher port to the next member lower port.

Objectives

- Understand VSF auto-stacking.
- Enable VSF auto-stacking on the Conductor.
- Verify VSF auto-stacking.

Steps

Factory Reset the devices

1. Open an SSH connection to sw-edge1.
2. Verify that member 2 is active as *Standby* in the VSF stack. It may still be booting after the last task, make sure it has completed the reboot before continuing.

```
show vsf
```

```
sw-edge1# show vsf

Force Autojoin              : Disabled
Autojoin Eligibility Status: Not Eligible
```

```
MAC Address                : 64:e8:81:dd:81:80
Secondary                  : 2
Topology                   : Ring
Status                     : No Split
Split Detection Method     : mgmt


Mbr Mac Address         type           Status
ID
--- ------------------- -------------- ---------------
1   64:e8:81:dd:81:80   JL666A         Conductor
2   8c:85:c1:49:20:c0   JL666A         Standby
```

IMPORTANT: Make sure to wait for both switches to be online (rebooted) after the previous task!

3. Zeroize the configuration.

```
erase all zeroize
```

```
sw-edge1# 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.
sw-edge1#
 Oct 17 07:41:30 hpe-mgmtmd[34933]: RebootLibPh1: Reboot reason: Reboot requested by
user
```

It will take *several minutes* to complete the factory reset.

Prepare the Conductor Switch

NOTE:
This step can also be performed by pushing the LED button on the physical switch. In that case no console cable is required to initiate the creation of a VSF stack.

Here is an excerpt from the AOS-CX 10.09 Virtual Switching Framework (VSF) Guide:

1. Physically connect the switches in a desired topology on the reserved VSF link ports.
2. Press the LED mode button on the conductor until the mode changes to “Stk”. The stack members
reboot one after another and join the stack.
During stacking operation, the port LEDs are displayed in three different states:
- Flashing green—Indicates that the member is the conductor.
- Flashing orange—Indicates that the member is rebooting to join the stack or offline due to error
condition.

- Solid green—Indicates that the member joined the stack and is operational.

In this remote lab, you don't have access to the LED button, so you will use the CLI to initiate the auto-stacking.

4. Open a connection to the console of sw-edge1. Since the switch is now factory default, connect with admin / no password (blank password).
5. When prompted for a new password, leave it blank (just press ENTER).

```
6300 login: admin
Password:


Please configure the 'admin' user account password.
Enter new password:
Confirm new password:
6300#
```

6. Review the default VSF configuration.

```
show vsf
```

```
6300# show vsf

Force Autojoin             : Disabled
Autojoin Eligibility Status: Eligible
MAC Address                : 64:e8:81:dd:81:80
Secondary                  :
Topology                   : Standalone
Status                     : No Split
Split Detection Method     : None


Mbr Mac Address         type           Status
ID
--- ------------------ -------------- ---------------
1   64:e8:81:dd:81:80   JL666A         Conductor
```

Question: Is there any standby member configured?

Answer: No, by default there is no standby member set.

7. Review the VSF running configuration.

```
show running-config vsf
```

```
6300# show running-config vsf
vsf member 1
    type jl666a
```

Question: Is there a member 2 present?

Answer: No, by default only member 1 is present.

Question: Are there any VSF links configured?

Answer: No, by default, no VSF links are configured.

8. Enter the configuration mode and initiate the auto-stacking

```
config
 vsf start-auto-stacking
```

```
6300# config
6300(config)#  vsf start-auto-stacking
This will configure links and secondary on conductor

Do you want to continue (y/n)? y
```

9. Review the updated VSF configuration.

```
show vsf
```

```
6300(config)# show vsf

Force Autojoin              : Disabled
Autojoin Eligibility Status: Not Eligible
MAC Address                 : 64:e8:81:dd:81:80
Secondary                   : 2
Topology                    : Standalone
Status                      : Active Fragment
Split Detection Method      : None


Mbr Mac Address          type           Status
ID
--- ------------------ -------------- ---------------
1   64:e8:81:dd:81:80   JL666A         Conductor
2                       JL666A         Not Present
```

Question: Is there a secondary member configured now?

Answer: Yes, member 2 has been set as the standby.

10. Review the updated VSF running-configuration.

```
show running-config vsf
```

```
6300(config)# show running-config vsf
vsf secondary-member 2
vsf member 1
    type jl666a
    link 1 1/1/26
    link 2 1/1/25
vsf member 2
    type jl666a
```

```
    link 1 2/1/25
    link 2 2/1/26
```

11. Wait for member 2 to complete the reboot in the VSF stack with the role of Secondary.
12. Review the VSF status.

```
show vsf
```

```
6300(config)# show vsf

Force Autojoin              : Disabled
Autojoin Eligibility Status: Not Eligible
MAC Address                 : 64:e8:81:dd:81:80
Secondary                   : 2
Topology                    : Ring
Status                      : No Split
Split Detection Method      : None


Mbr Mac Address         type           Status
ID
--- ------------------- -------------- ---------------
1   64:e8:81:dd:81:80   JL666A         Conductor
2   8c:85:c1:49:20:c0   JL666A         Standby
```

Question: What is the configured Split Detection Method?

Answer: None. Split detection is not automatically configured by the auto-stacking feature.

This concludes the VSF auto-stacking task.

Task 7: Cleanup configuration

In this last task you will:

- revert the aggregation switch configuration to support two access switches with an MCLAG for each.
- factory reset the access switches. After the reset they will operate as two individual switches again.

Revert the aggregation switch configuration

1. On sw-agg1, assign port 1/1/2 back to LAG2.

```
interface 1/1/2
 no lag 1
 lag 2
 no shutdown
 exit
```

```
sw-agg1(config)# interface 1/1/2
sw-agg1(config-if)#  no lag 1
sw-agg1(config-if)#  lag 2
sw-agg1(config-if)#  no shutdown
sw-agg1(config-if)#  exit
```

2. On sw-agg2, assign port 1/1/2 back to LAG2.

```
interface 1/1/2
 no lag 1
 lag 2
 no shutdown
 exit
```

```
sw-agg2(config)# interface 1/1/2
sw-agg2(config-if)#  no lag 1
sw-agg2(config-if)#  lag 2
sw-agg2(config-if)#  no shutdown
sw-agg2(config-if)#  exit
```

Factory Reset the Edge Switches

3. On sw-edge1, zeroize the configuration.

```
erase all zeroize
```

```
6300# 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.
```

The edge switches are left at factory default state at the end of this lab activity.

NOTE: The edge switches are left at factory default state at the end of this lab activity. Since you are *not* using the VSF **start-auto-stacking** command, the two switches will remain in their factory default, single switch state.
This is the expected state for the next lab activity.

You have completed this Lab!

Lab 07.01 Managing Aggregation Switches with Aruba Central

Overview

Your senior network colleague is very happy with your progress in networking skills. Your next assignment is to learn to use the Aruba Central cloud-based management platform to manage and monitor the switches.

This lab will first introduce you to the basic access and navigation in Aruba Central. This will include the configuration of sites and labels.

Next you will migrate the aggregation switches to Aruba Central using a template configuration that has been prepared for you by your senior colleague.

Objectives

- Configure Aruba Central sites and labels.
- Understand the requirements for the cloud connection.
- Assign devices to a template-based group.
- Verify the deployment of a template configuration.

Task 1: Prepare the topology

In this task you will configure the aggregation switches with an IP Helper address on SVI1. This ensures that the factory default edge switches can receive a DHCP address and attempt to connect to Aruba Central for their cloud management connection.

Objectives

- Configure IP helper on the aggregation switches.

Steps

1. Use MGMT PC to open an SSH connection to **sw-agg1**.
2. Configure SVI1 with an IP Helper address of 10.254.1.21.

```
interface vlan 1
 ip helper-address 10.254.1.21
 exit
```

```
sw-agg1(config)# interface vlan 1
sw-agg1(config-if-vlan)#  ip helper-address 10.254.1.21
sw-agg1(config-if-vlan)#  exit
```

Task 2: Essential Aruba Central configuration

In this task you will perform some essential steps when configuring an Aruba Central environment for the first time.

Objectives

- Access Aruba Central using HPE GLCP.
- Understand the core navigation flow in Aruba Central.
- Configure sites and labels in Aruba Central.

Steps

Access HPE GreenLake Cloud Platform (GLCP)

1. On your local system, use a browser to open a connection to the HPE GreenLake Cloud Platform

https://common.cloud.hpe.com

2. Click on **Sign In with SSO**. Enter the email address provided by your instructor and click Next.

You are now redirected to the Aruba Training Labs SAML host.

3. Enter the email address and password provided by your instructor and click Login.

You are now logged in to HPE GreenLake Cloud Platform.

Launch Aruba Central application

4. On the **Aruba Central** tile, click **Launch**.

Aruba Central UI Navigation Instructions

In the next steps you will be introduced to the core navigation structure in Aruba Central: using the context, navigation, and top areas. This will be used in the remainder of the labs to guide you to the correct Aruba Central screen.

Context Filter

The context filter is used to narrow the scope of the Aruba Central UI.

5. By default, the Global context is selected.

6. Click on the context area to see a pop-up with context options. Currently, only the default group is listed in the context filter.

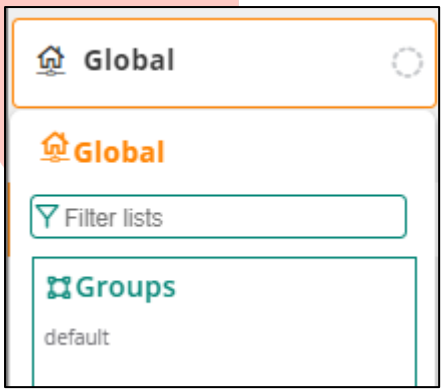

7. As you progress with the lab activities, additional groups, sites, and labels will show up in this context selection list.

   Example: Context filter with Groups, Sites and Labels:

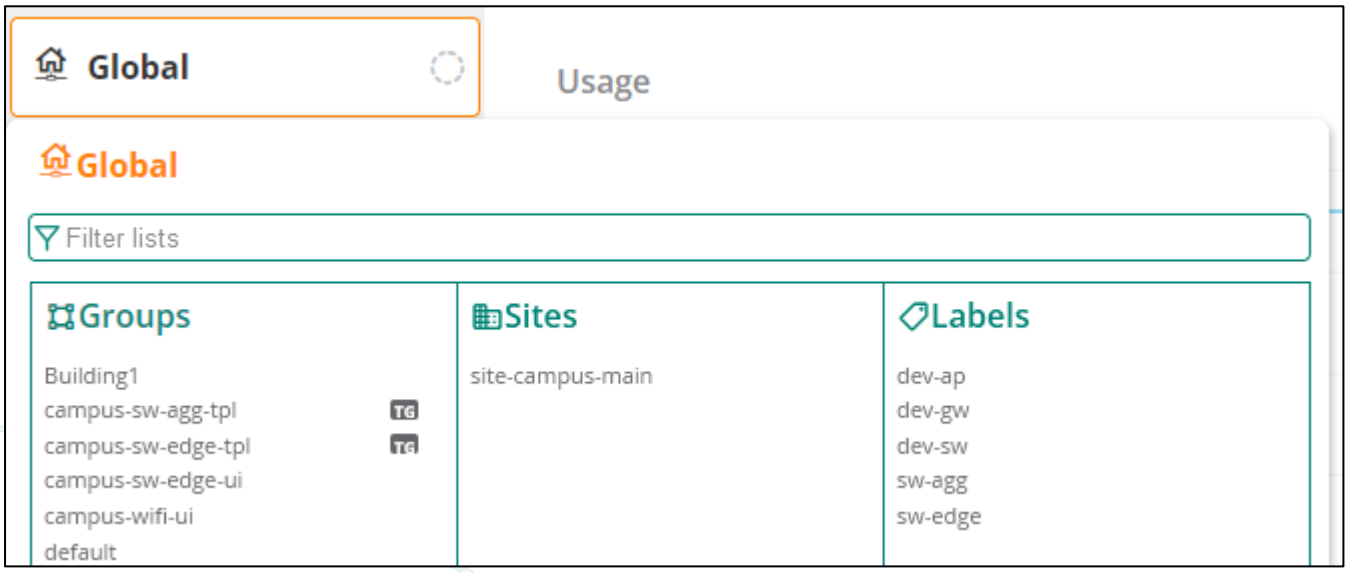

In the lab guide, you will find instructions such as:

Navigate to Context: **Groups / default**

8. You should now click on the **Context** button, under the Groups heading, then click on **default**.

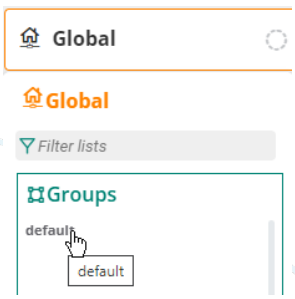

9. The context filter will now show the *group* icon  and the name of the group: *default*.

Navigation Menu

On the left side of the UI, you see the navigation menu. This menu provides access to various status and configuration screens for a given context.

10. This is an example of the Navigation menu with Manage, Analyze, and Maintain sections:

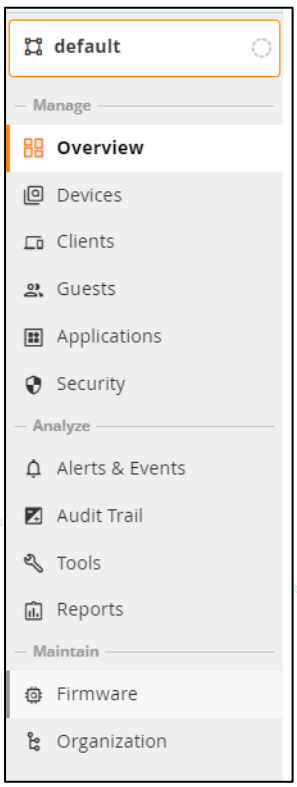

11. You may see an instruction such as:

Navigate to:
Context: **Groups / default** > Navigation: **Devices**.

12. You should first verify that you are at the correct context filter level (or change the context filter), then use the Navigation menu on the left side and click on **Devices.**

    Example:

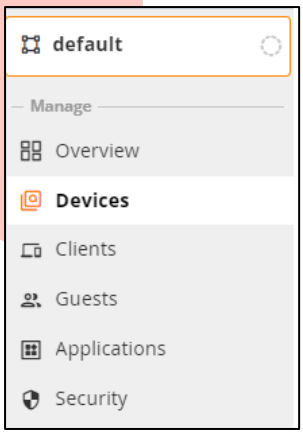

Top Menu

Each of the navigation menu's Aruba Central screens have one or more tabs at the top of the screen.

13. In this example (Context: **Group / default** > Navigation: **Devices**), the Top options are **Access Points** and **Switches**.

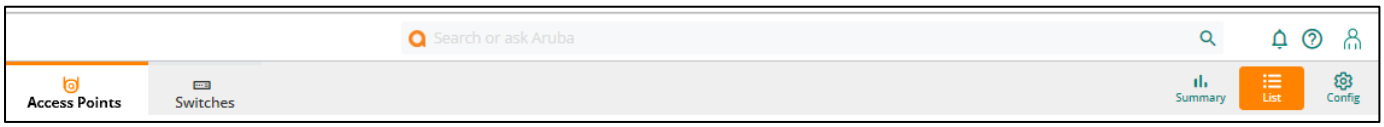

The example instruction to reach this page would be:
Context: **Groups / default** > Navigation: **Devices** > Top: **Access Points**.

Top Right

Many of the Aruba Central pages will have similar options at the Top right of the screen.

14. In this example (Context: **Group / default** > Navigation: **Devices** > Top: **Access Points**), the Top Right options are **Summary, List** and **Config**.

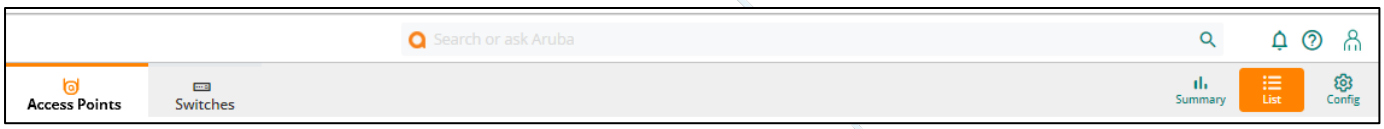

To access this example page, the instruction would be:

Context: **Group / default** > Navigation: **Devices** > Top: **Access Points** > Top right: **List**

This concludes the introduction to the Aruba Central navigation.

Sites and Labels

Aruba Central uses sites to map devices to their physical locations. In Aruba Central, all devices in the same location should be mapped to the same site. A device can only belong to one site.

15. In Aruba Central, navigate to:

Context: **Global** > Navigation: **Organization** > Top: **Network Structure** > **Sites**

Example screenshot:

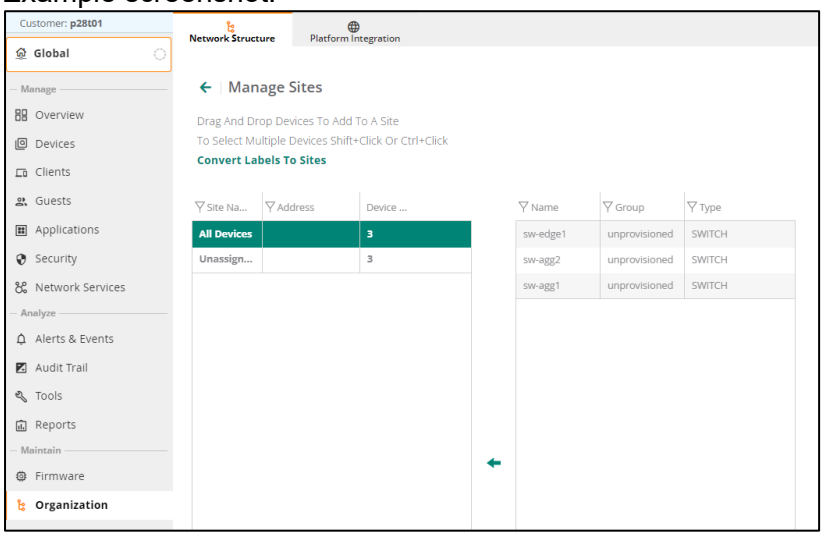

16. At the bottom of the page, click on the plus sign to create a **New Site**.

> **NOTE:** This guide uses a fictitious example site address. Feel free to enter your own site address information, but please use the site name **site-campus-main**.

Field	Value
Name	site-campus-main
Street Address	Main Street
City	Oranjestad
Country	Aruba
State	Aruba
ZIP	0000

17. Click **Add** to save the site.

Device Labels

Labels can be used to assign logical markers to devices. Unlike a site, it is possible to assign multiple labels to a single device.

18. In Aruba Central, navigate to:

Context: **Global** > Navigation: **Organization** > Top: **Network Structure** > **Labels**

19. At the bottom of the page, click the plus sign to create a **New Label**. Set the lab name to **sw-agg**
20. Add another label and set the lab name to **sw-edge**.

Verify site and label

21. In Aruba Central, click on the Context filter. The new site and labels will be listed in the Sites and Labels columns. This makes it convenient to jump to a site or location.

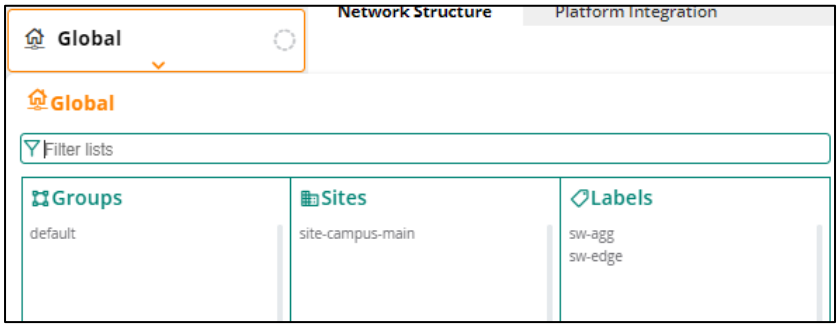

Task 3: Provide Aggregation Switches with Static Cloud Connection

In this task you will connect the existing aggregation switches to Aruba Central. The goal of this task is to show the connection requirements and how the connection is established for a switch that has an existing configuration.

Aruba Central uses groups to organize device configurations. At this point in the lab, no group has been created. This also means that the aggregation devices will not receive a configuration from Aruba Central yet. This will be done in the upcoming tasks. For right now, only the connection to Aruba Central will be verified.

Objectives

- Connect an AOS-CX switch to Aruba Central.

Steps

1. Open a connection to Aruba Central.
2. Navigate to

Context: **Global** > Navigation: **Devices** > Top: **Switches**

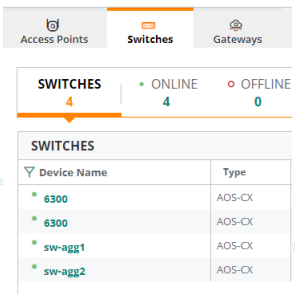

Question: Are there any switches listed?

Answer: Yes, your edge and aggregation switches are connected.

Question: What is the color of the small dot in front of the sw-agg1 device name?

Answer: This is a green dot; it represents an active connection.

Sw-agg1

3. Use MGMT PC to open an SSH Connection to sw-agg1.
4. Review the Aruba Central connection status.

```
show aruba-central
```

Question: What is the admin state?

Answer: Aruba Central support is enabled. This is the default state.

Question: What is the connection status?

Answer: The reported connection status is connected.

Aruba Central requires internet access and DNS functionality. You will now review these requirements.

5. Attempt to reach an IP address on the internet to verify IP reachability.

```
ping 8.8.8.8
```

Question: Was the ping successful?

Answer: Yes, in previous lab you have configured OSPF adjacency with the core router. The core router provides a default route to the internet. Behind the core router, there is router that provides NAT services for your lab.

6. Attempt to reach an internet system based on a DNS FQDN. You can use the Aruba Path Quality Monitoring (PQM) host for this test.

```
ping pqm.arubanetworks.com
```

Question: What do you observe?

Answer: The name cannot be resolved. This indicates something is missing in the DNS configuration.

7. Review the DNS configuration.

```
show ip dns
```

Question: Are there any DNS servers configured?

Answer: No, a DNS server is missing from the configuration.

Question: How is it possible that the connection with Aruba Central is connected when no DNS servers have been configured?

Answer: The Aruba Central client process on AOS-CX includes by default the DNS server 8.8.8.8. Only the Aruba Central client component on the switch uses this default DNS server. Other utilities, such as ping, use the global DNS configuration.

8. To override the default DNS, configure the host 10.254.1.21 as the DNS server address.

```
ip dns server-address 10.254.1.21
```

9. Attempt to validate the change with a ping to pqm.arubanetworks.com. The ping should now be successful.

```
ping pqm.arubanetworks.com
```

10. Check the status of the Aruba Central connection again, it should still show as connected.

```
show aruba-central
```

NOTE: If something was wrong with the connection, you can speed up the connection retry process by disabling and enabling the Aruba Central feature on the switch.

```
aruba-central
 disable
 enable
 exit
```

Question: What is the Aruba Central location?

Answer: This is the URL of the Aruba Central region your account is hosted on.

```
sw-agg1(config)# show aruba-central
Central admin state                     : enabled

Central location                        : device-uswest4-d2.central.arubanetworks.com
VRF for connection                      : default
Shared Token                            : N/A
Central connection status               : connected

Central source                          : activate
Central source connection status        : connected
Central source last connected on        : Tue Sep 27 16:02:10 UTC 2022
System time synchronized from Activate  : True

Activate Server URL                     : devices-v2.arubanetworks.com
CLI location                            : N/A
CLI VRF                                 : N/A

Source IP                               : 10.254.101.2
Source IP Overridden                    : False

Central support mode                    : disabled
```

Assign the site

Now that the aggregation switches are connected to Central, you can assign them to a site. A device can only belong to a single site.

11. In Aruba Central, navigate to

Context: **Global** > Navigation: **Organization**> Top: **Network Structure** > **Sites**

12. In the left pane, select **Unassigned**. The two aggregation switches should be listed in the pane on the right side.
13. Select both aggregation switches (you can use the control or command key to select multiple entries), then drag them to the site named **site-campus-main**.

14. Confirm the action with **Yes**.

Assign a label

Next you can assign a label. Up to five labels can be assigned to a device.

15. In Aruba Central, navigate to

Context: **Global** > Navigation: **Organization**> Top: **Network Structure** > **Labels**

16. In the left pane, select **All Devices**. The two aggregation switches should be listed in the pane on the right side.
17. Select both aggregation switches, then drag them to the label named **sw-agg**.
18. Confirm the action with **Yes**.

Verify Site and Label

19. In Aruba Central, click the **Context** menu.
20. While this menu is open, your cursor is in the '**Filter lists'** field. You can start typing here to filter the context menu. Type **main.**

Question: What do you notice?

Answer: Only the groups, sites, and labels that contain the string *main* will be displayed now. This is convenient when a customer has many groups, sites, or labels.

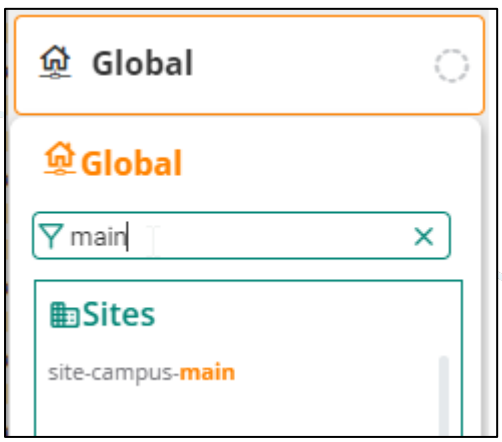

Task 4: Aggregation switches Template Configuration Group

Your senior colleague wants to ensure that the aggregation switches receive their full configuration from Aruba Central. This should be based on a configuration that has been prepared by your senior colleague.

In this task you will prepare a template group in Aruba Central for the aggregation switches and upload the prepared configuration.

Next you will assign the aggregation switches to this group, this will make Aruba Central push the configuration to the switches. When Aruba Central is managing the configuration using a template, network administrators should not make any changes to the switch configuration anymore.

In this task you will see how the configuration will become read-only, and how you can still make exceptions to this to provide troubleshooting and support.

Objectives

- Create a switch template group.
- Assign a template to a group.
- Move a device to the group.
- Understand the configuration lockout feature.
- Understand how the support mode feature can be used.

Steps

Create Template Group

First you will create the switch template group.

1. In Aruba Central, navigate to

Context: **Global** > Navigation: **Organization**> Top: **Network Structure** > **Groups**

2. At the right-top, click the Plus sign to add a **New Group**.
3. For the name, enter '**campus-sw-agg-tpl**'.
4. For the value "Group will contain", only select **switches**.
5. Configure using **templates**: move the slider to the right (**enabled**). The checkbox for switches is automatically selected.
6. Click **Next**.
7. Types of switches used: Select **AOS-CX only**.

> **NOTE:** You don't need to select 'Make these the preferred settings'. This would make the current selection the default selection for future group additions.

8. Click **Add**.

9. Verify that the group **campus-sw-agg-tpl** is now listed.

Question: Why is there a TG label in front of the group name?

Answer: This is how Central shows that the group configuration is based on a template group (TG).

Question: What icon type do you see in the second column?

Answer: Central shows what type of devices are enabled for the group. Since you only enabled switches, only a switch icon will be displayed for this group.

Assign a template to a group

You will now apply a configuration to this group. This configuration was prepared by your senior colleague and has been saved in the ACAF Student files folder on the MGMT PC Desktop .

10. Navigate to Context: **Groups / campus-sw-agg-tpl** > Navigation: **Devices**> Top: **Switches** > **Config** (gear icon).

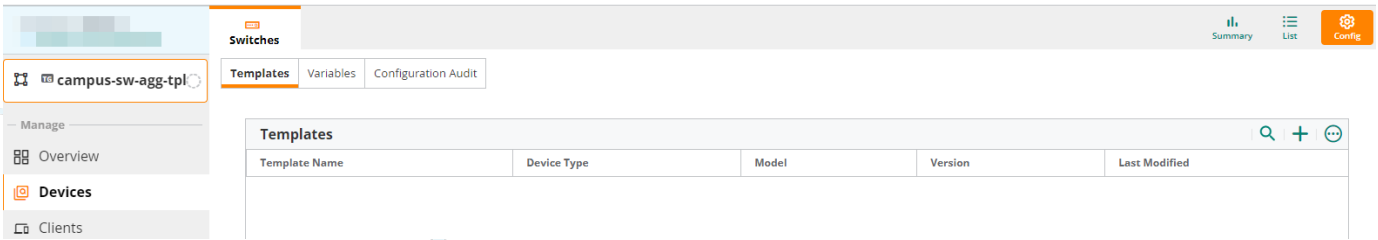

The Top navigation shows **Templates, Variables,** and **Configuration Audit**.

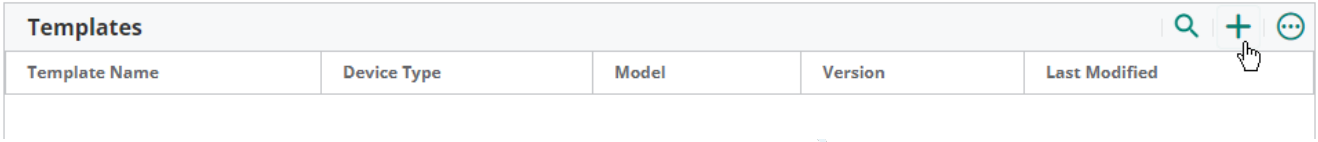

11. Under **Templates**, at the right-top, use the + icon to add a template. The Add Template window appears.
12. For template name, enter **sw-agg**.
13. For **device type**, verify **Aruba CX** is selected.(default based on the group settings)
14. You can leave the other fields at default and click **Next**.

15. From the MGMT PC, navigate to Desktop > ACAF student files and open **lab07.01-sw-agg-tpl.txt**.
16. Copy all the contents of the template text file.
17. Paste the contents of the file in the Template screen.
18. Click **Save** to complete this step. The template will now be uploaded to Aruba Central.
19. Verify that the new template is now listed in the Templates list.
20. To view the contents of an existing template, move the mouse cursor over the **sw-agg** template entry. At the end of the line, a **pen** icon and **delete** icon will show up.
21. Click the pen icon to edit the template. The template contents will now be displayed.
22. Navigate to line 24 in the template text.

Question: What is the text you see on the lines 24-30?

Answer:

```
%if _sys_hostname=sw-agg1%
spanning-tree priority 0
    %endif%
%if _sys_hostname=sw-agg2%
spanning-tree priority 1
    %endif%
```

Question: What does this mean?

Answer: There is a condition:
If a variable named '_sys_hostname' has the value 'sw-agg1', then the command 'spanning-tree priority 0' will be applied.
For '_sys_hostname' sw-agg2 the command 'spanning-tree priority 1' will be applied.

Question: What would happen when the hostname would be sw-agg3?

Answer: Since the hostname does not match any condition, none of these 2 example commands would be applied. This template relies on the correct hostname of the switches. Next, you will double-check the hostname of your aggregation switches before moving them to this group.

23. Look for the interface 1/1/1 in the template.

```
######## per switch interface config
### sw-agg1
%if _sys_hostname=sw-agg1%
interface 1/1/1
    description sw-edge1-g1/1/27
interface 1/1/2
    description sw-edge2-g1/1/27
interface 1/1/5
    description gw1-g0/0/1
interface 1/1/8
    no shutdown
    description rtr-core1-1/1/1
    ip address 10.254.101.2/24
```

```
    ip ospf 1 area 0.0.0.0
    no ip ospf passive
    ip ospf dead-interval 4
    ip ospf hello-interval 1
    ip ospf network point-to-point
interface 1/1/10
    description gw2-g0/0/1
interface 1/1/46
    description sw-agg2-g1/1/46
interface 1/1/47
    description sw-agg2-g1/1/47

%endif%

### sw-agg2
%if _sys_hostname=sw-agg2%
interface 1/1/1
    description sw-edge1-g1/1/28
interface 1/1/2
    description sw-edge2-g1/1/28
interface 1/1/5
    description gw1-g0/0/2
interface 1/1/8
    no shutdown
    description rtr-core1-1/1/2
    ip address 10.254.102.3/24
    ip ospf 1 area 0.0.0.0
    no ip ospf passive
    ip ospf dead-interval 4
    ip ospf hello-interval 1
    ip ospf network point-to-point
interface 1/1/10
    description gw2-g0/0/2
interface 1/1/46
    description sw-agg1-g1/1/46
interface 1/1/47
    description sw-agg1-g1/1/47

%endif%
```

Question: What do you observe?

Answer: Interface 1/1/1 is configured two times in a conditional section for the hostnames sw-agg1 and sw-agg2. On each of the aggregation switches, a unique interface description is applied.

Verify the aggregation switch hostnames

24. In Aruba Central, click **Cancel** to close the template details.
25. In Aruba Central, navigate to Context: **Global** > Navigation: **Devices**> Top: **Switches** > Right-top: **List**
26. Verify in the Device Name column that the listed name is **sw-agg1** and **sw-agg2**.
27. If the hostname is not correct, the next steps can be used to correct it.

Only perform these steps if the system would have an incorrect hostname

In the remote lab, the aggregation switches have initially been configured using a Zero Touch Provisioning file, this includes the correct hostnames sw-agg1 and sw-agg2.

The next steps are **only** required if you would have changed the hostname yourself to some other hostname.

28. Using MGMT PC, open an SSH connection to sw-agg1 or sw-agg2.
29. Verify that the hostname is exactly **sw-agg1**.

```
show hostname
```

IMPORTANT:
If the hostname is not exactly sw-agg1, change it with the command

```
config
 hostname sw-agg1
 end
```

End of steps to correct the hostname (if applicable).

Move a device to the group

In the next steps you will start by moving *only* sw-agg1 to the target group, then you will verify the applied configuration. *After* these validation steps have been completed, you will move sw-agg2.

30. In Aruba Central, navigate to Context: **Global** > Navigation: **Organization**> Top: **Network Structure**> **Groups**
31. Expand “All connected devices” by clicking on the “>” icon.
32. Under All connected devices, select sw-agg1.

NOTE: Pay attention. This screen is a multi-select enabled by default; clicking on multiple devices will select all of them. To un-select, click again on the device.

33. On the right-hand side, a popup will be displayed with the **Move Devices** action button.

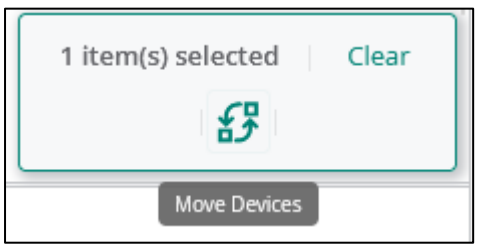

34. Click on the **Move Devices** button.
35. Click on the **Destination Group** field.

36. You can either select the group **campus-sw-agg-tpl** OR start typing a substring of the group name, such as **agg**, then select the group from the filtered list. This can be convenient when a lot of groups exist in Central.
37. Click on the **Move** button to continue. With this action, Aruba Central will be in control of the sw-agg1 configuration. The existing configuration of sw-agg1 will be completely replaced with the templated that you have uploaded to Central.

Verify the configuration sync status

In these steps you will verify the status of the configuration sync. This can be done at the device or group level.
In the next steps you will check the configuration sync at the device level. This level also allows you to check the 'Attempted' configuration and the 'Running' configuration of the device.

The *Attempted* configuration is the configuration that Central attempts to push to the device, while the running configuration is the configuration that is currently running on the device. There may be a difference in the syntax (for example port ranges) and command order between these two versions.

38. In Aruba Central, navigate to the AI search bar:

Top : Search or ask Aruba

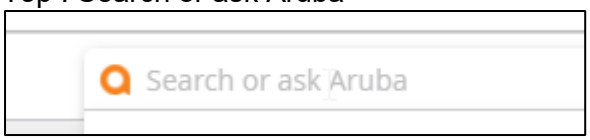

39. Enter **sw-agg1** in the search box and press **ENTER**. After a few seconds, the results will be displayed in a popup, this will include sw-agg1.

40. Click the name **sw-agg1** in the popup window. This will take you to the device level homepage.

> **NOTE:** Several shortcuts are available in the popup window. Just click the sw-agg1 text next to the green dot.

Notice that the context filter has been set to **sw-agg1** now.

41. In the Navigation bar, click **Device**, under Top: **Switch**, select **Configuration Audit**.
42. Check the status in the Configuration Status box.

Question: What is the status in the box?

Answer: The status should be **Config In Sync**.

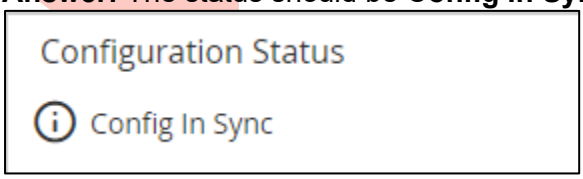

Check attempted and running configuration

43. Under this box, click on the '**View Details'** link. A new window will appear with configuration options.

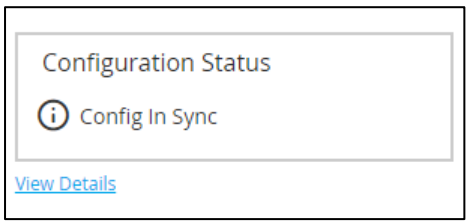

44. To see the attempted configuration, click the **Attempted Configuration** link.

Question: Scroll through the attempted configuration. What is the spanning tree priority command you see?

Answer: The attempted configuration shows the command 'spanning-tree priority 0'.

Question: Do you see the %if .. conditional logic that was present in the template that you have uploaded?

Answer: No, this is the final (the attempted) configuration that Aruba Central has generated based on the template that you have uploaded, combined with variables (such as the hostname) and the conditional logic. This is a good place to detect any command issues with the template for a given device.

45. To see the actual device running configuration, click the **Device Running Configuration** link.
46. Click **Close** to close the configuration window.
47. Use MGMT PC to open an SSH connection to **sw-agg1**.
48. Check the show interface brief output.

```
show interface brief
```

Question: Do you see an interface description for interface 1/1/1 and 1/1/2?

Answer: Yes, the template has pushed the updated configuration, this includes interface descriptions.

Move sw-agg2 to the campus-sw-agg-tpl group

After the successful migration of sw-agg1, you will now move sw-agg2 to the campus-sw-agg group.

You have learned in the previous section how to move a device to a new group. Try to move the device sw-agg2 to the group campus-sw-agg-tpl yourself. Refer to the previous section if you are unsure about the steps.

In the next section you will verify the configuration sync.

Verify sw-agg2 configuration sync

Previously you have used the device level configuration audit. This time you will use the group-level configuration audit. This is convenient to monitor a list of devices in a single operation.

49. In Aruba Central, navigate to Context: **Groups / campus-sw-agg-tpl** > Navigation: **Devices**> Top: **Switches** > Right-top: **Config**
50. Click on Configuration Audit.
51. In the middle of the page, a **Configuration Status** box is shown.

Question: What is the count for **Not In Sync** devices?

Answer: The count will be 0 when the configuration was successfully applied. It may take a few moments (less than a minute) to push the configuration. If the count would be 1, refresh the page until it is 0.

52. Use MGMT PC to open an SSH connection to sw-agg2.
53. Check the show interface brief command output.

```
show interface brief
```

Question: Are the interface descriptions updated for ports 1/1/1 and 1/1/2?

Answer: Yes, the interface descriptions have been updated.

Congratulations, you have successfully migrated the aggregation switches to Aruba Central.

You have completed the Lab!

Lab 07.02 Managing Edge Switches with Aruba Central

Overview

Now that you have migrated the aggregation switches, you will explore options to deploy access switches using Zero Touch Provisioning (ZTP). When the access switches have an active connection, you will explore how Central UI and MultiEdit can be used to configure the access switches.

The last step is to apply a prepared template configuration to the access switches. This ensures that they have a 'known-good' configuration, it makes them ready to for the rest of the lab activities.

Objectives

- Understand the ZTP process.
- Assign devices to a UI-based configuration group.
- Use MultiEdit to configure switches.
- Apply a template to edge switches and verify the template deployment.

Task 1: Review Edge Switch Cloud Connection with ZTP

In the previous lab, you connected the aggregation switches to Aruba Central using a static IP address and a template configuration. In this task you will use the edge switches. You will review how the factory default access switches can be onboarded with Zero Touch Provisioning. This is convenient when performing a deployment of new devices.

Objectives

- Understand the ZTP process.
- Understand the infrastructure requirements for ZTP.

Steps

ZTP Untagged Native VLAN

To support ZTP, the access switches must get a DHCP IP address. By default, AOS-CX switches have a DHCP client enabled on their SVI1 (VLAN1). This is untagged on all ports.

When connecting the factory default switch to the network aggregation switches, the aggregation switches must be provisioned with an untagged native VLAN on their ports connecting to the access switches. This does not need to be VLAN1 on the aggregation side.

When you loaded the template to the aggregation switches, the native VLAN on the LAG to the access switches was set to VLAN1. This VLAN is used as the ZTP device onboarding VLAN in the lab.

1. Use **MGMT PC** to open an SSH connection to **sw-agg1**.
2. Review the current configuration of LAG1. This LAG connects to sw-edge1.

```
show running-config interface lag1
```

```
interface lag 1 multi-chassis
    no shutdown
    description sw-edge1
    no routing
    vlan trunk native 1
    vlan trunk allowed 1,3-4,11-15
    lacp mode active
    lacp fallback
    exit
```

Question: What is the configured trunk native VLAN?

Answer: The trunk native VLAN is set to VLAN 1.

3. Review the SVI1 configuration.

```
show running-config interface vlan1
```

```
sw-agg1# show running-config interface vlan1
interface vlan1
    ip address 10.1.1.2/24
```

```
    active-gateway ip mac 12:01:00:00:01:00
    active-gateway ip 10.1.1.1
    ip ospf 1 area 0.0.0.0
    ip helper-address 10.254.1.21
    exit
```

Question: Is there an IP helper address configured?

Answer: Yes, the IP helper points to 10.254.1.21. Any device connecting on SVI1 could receive a DHCP address.

Now check if there are any IP devices active on VLAN 1. The DHCP server scope IP range is 10.1.1.50-10.1.1.150.

4. Check the ARP table for VLAN1 entries.

```
show arp | include vlan1
```

```
sw-agg1# show arp | include vlan1
10.1.1.51        64:e8:81:dd:81:80  vlan1        lag1                      reachable
10.1.1.52        20:4c:03:c6:09:50  vlan1        lag1                      reachable
10.1.1.54        20:4c:03:b1:d1:9a  vlan1        lag256                    reachable
10.1.1.55        20:4c:03:b7:70:1a  vlan1        lag256                    reachable
```

Question: Are there any IP DHCP devices active?

Answer: Yes. There are four devices active.

Question: Why four devices?

Answer: The APs are connected to the factory default access switches. Since all ports on the access switches belong to VLAN1 (in the factory default state), the APs are also getting an IP address from SVI1.

Optional: You can check the DHCP client status and the DHCP options, such as default gateway and DNS Servers.

5. **Optional Step**: Open the console of **sw-edge1** and check the DHCP client of the switch.

```
show ip dhcp
```

```
6300# show ip dhcp
INTERFACE-NAME ADDRESS       DEFAULT_GATEWAY  DOMAIN_NAME          VRF        DNS-
SERVERS
-------------------------------------------------------------------------------
-----
vlan1          10.1.1.51/24   10.1.1.1        aruba-training.com   default
10.254.1.21
6300#
```

ZTP LACP Fallback

Typically access switches connect via an LACP LAG to the aggregation switches. When a factory default switch connects to an LACP LAG, the aggregation switch does not receive any LACP communication from the access switch and the ports are set to *lacp-blocked* state.

To support ZTP with factory default switches, the aggregation switch should enable LACP fallback on the LAG. This is enabled on the ports when no LACP communication is received.

6. On **sw-agg1**, review the running configuration of LAG1.

```
show running-config interface lag1
```

```
sw-agg1# show running-config interface lag1
interface lag 1 multi-chassis
    no shutdown
    description sw-edge1
    no routing
    vlan trunk native 1
    vlan trunk allowed 1,3-4,11-15
    lacp mode active
    lacp fallback
    exit
```

Question: What LACP commands to you see?

Answer: The LACP commands are:

```
lacp mode active
lacp fallback
```

7. Review the LACP interface state.

```
show lacp interface
```

```
sw-agg1# show lacp interfaces

State abbreviations :
A - Active        P - Passive      F - Aggregable I - Individual
S - Short-timeout L - Long-timeout N - InSync     O - OutofSync
C - Collecting    D - Distributing
X - State m/c expired              E - Default neighbor state

Actor details of all interfaces:
------------------------------------------------------------------------------
Intf       Aggr       Port  Port  State   System-ID         System Aggr Forwarding
           Name       Id    Pri                             Pri    Key  State
------------------------------------------------------------------------------
1/1/1      lag1(mc)   1     1     IE      02:01:00:00:01:00 65534  1    up
...


Partner details of all interfaces:
```

```
--------------------------------------------------------------------------------
Intf       Aggr       Port  Port  State   System-ID         System Aggr
           Name       Id    Pri                             Pri    Key
--------------------------------------------------------------------------------
1/1/1      lag1(mc)   0     0     IE      00:00:00:00:00:00 0      0
...
```

Question: What is the lag1(mc) LACP state?

Answer: Since the access switch is not sending any LACP communication, the state is IE (default).

Question: What is the forwarding state of the lag1(mc)?

Answer: The port is *up* instead of lacp-blocked. This was the result of the LACP fallback command.

Site Assignment

Assign both access switches to the site **site-campus-main** and label **sw-edge**.

8. In Aruba Central, navigate to Context: **Global**> Navigation: **Organization**> Top: **Network Structure**> **Sites**
9. Select both access switches 6300, drag them to the site **site-campus-main**.
10. Confirm the action with **Yes**.

Label Assignment

11. In Aruba Central, navigate to

Context: **Global**> Navigation: **Organization**> Top: **Network Structure**> **Labels**

12. Select both access switches 6300, drag them to the label **sw-edge**.
13. Confirm the action with **Yes**.

Task 2: Configure Access Switches with Central UI Group

In this task you will use an Aruba Central UI group to configure the access switches. When compare your previously built configurations for each access switch, you notice many similar configuration items. These can be considered **common configuration settings**:

- Administrator password
- Global VLANs
- Every switch will have an uplink LAG256 with LACP
- DNS/NTP/Timezone configuration
- Default Gateway IP address

There are also several **unique configuration settings** for each switch:

- Hostname
- The uplink ports that are assigned to the uplink LAG. This may be the same, but it could be different if you have some 24 port and some 48 port switches, or different combinations in a VSF stack.
- Switch IP address
- Access port configurations

You can apply common configuration settings at the group level. These settings are automatically inherited by every switch assigned to the group. The unique per-switch settings can then be applied to each individual switch.

Conflicts
Typically, the set of configuration items at the group and device level will be different, so there should not be a conflict between the 2 levels. If an item is configured at both group and device level, the device level configuration is applied.

Objectives

- Use an Aruba Central UI Group
- Apply settings at group and device level
- Verify the configuration synchronization

Steps

Create Switch UI Group and Apply Base configuration

1. In Aruba Central, navigate to Context: **Global**> Navigation: **Organization**> Top: **Network Structure**> **Groups**
2. Click on the + icon to create a new group.
3. For name, enter **campus-sw-edge-ui**.
4. Select **Switches**, do not select the template option.
5. Click **Next**.
6. Select **AOS-CX only** and click **Add**.
7. Verify the new group is listed under Groups.

Apply Basic Group Configuration

8. In Aruba Central, navigate to Context: **Groups / campus-sw-edge-ui** > Navigation: **Devices**> Top: **Switches** > Right top: **config**

For a new group, a popup is displayed with an administrator password prompt.

9. Set the password to **Aruba123!** and click **Save**.

Configure System Properties

10. Click the **System** > **Properties** link.
11. Set the timezone to Detroit(UTC -05:00).
12. Set **DNS** servers to **10.254.1.21.**
13. Set **NTP** servers to **10.254.1.21.**

> **NOTE:** You just configured the admin password with the popup. If required, the group level admin password can be changed in this screen.

14. Click **Save**.

Configure the uplink LAG

15. Click **Interfaces > Ports & Link Aggregations**.
16. At the right-top, click + sign to add a port.

> **NOTE:** You may need to scroll to the right to see the + icon.

17. In the Add LAG screen enter these values:

- Name: **lag256**
- Do **not** configure Port Members. This will be done at the device level.
- VLAN Mode: **trunk**
- Native VLAN: **1**
- Allowed VLANs: **all**
- Admin Up: **Yes** (checked)
- Aggregation Mode: **LACP active**

18. Click **Add** to save the LAG.
19. Click the **Back** button (left arrow) to return to the main menu.

Configure the global VLANs.

20. Click **Bridging > VLANs**.
21. At the right-top, click + sign to add a VLAN.
22. In the Add VLAN screen enter these values:

- ID: **3**
- Name: **v3-mgmt**
- Leave admin Up enabled

23. Click **Add**.
24. Add another VLAN with these values:

- ID: **11**
- Name: **v11-employee**

25. Click **Add**.
26. Verify the new VLANs 3 and 11 are displayed in the VLANs list.
27. Click the **Back** button (left arrow) to return to the main configuration menu.

Configure default gateway

28. Click **Routing > Static Routing**.
29. Click + icon to add a new entry

- Destination: **0.0.0.0/0**
- Next Hop: **10.1.3.1**

30. Click **Save**.
31. Click the **Back** button (left arrow) to return to the main configuration menu.

You have now completed the basic configuration for all switches that will be assigned to this group. Now you are ready to move the first switch to the group.

Move two access switches to the group

In this section you will move the two access switches to the new group. When you move the devices, the group level settings (as defined in the previous steps) are applied to both switches.

32. In Aruba Central, navigate to Context: **Global**> Navigation: **Organization**> Top: **Network Structure**> **Groups**
33. Expand Unprovisioned devices.
34. Select **both** Aruba CX switches.(multi-select is automatically enabled in this screen)
35. Click on the **Move Devices** button.

36. In the destination group field, enter or select **campus-sw-edge-ui**, then click **Move**.

> **NOTE:** Do **not** check the 'Retain CX-switch configuration'. The switches are factory default at this point in the lab. This option can be used to import existing switch configurations (Brownfield deployment) or when a switch is moved from a template to a UI Group.

Configure the Per-Switch Settings for sw-edge1

In this section you will configure the **per-device** settings, such as the hostname, LAG port members and a static IP address for the switch.

37. In Aruba Central, navigate to Context: **Global**> Navigation: **Devices**> Top: **Switches** > Right top: **List**

38. In the column header, use the 3 dots in to enable the **Serial/Stack ID** column.

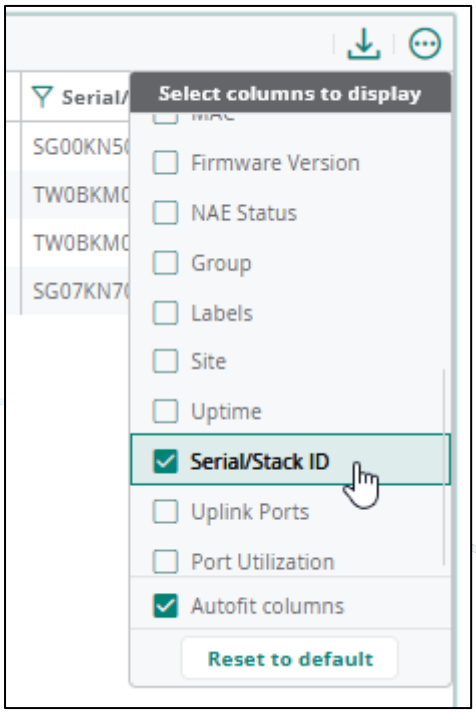

39. Use the **Serial column** (you may need to scroll to the right) to identify your sw-edge1. In Lab 1 you noted the serial numbers for the edge switches.

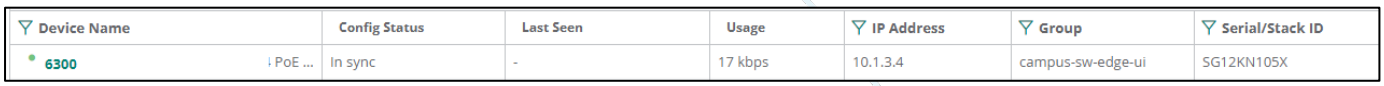

40. Click the 6300 device name at the start of the record. This takes you to the device context.

41. In the Navigation menu click **Device** to access the device level configuration.

Configure System Properties

1. Click the **System** > **Properties** link
2. Set the name to **sw-edge1**.
3. Verify that time zone, DNS, and NTP were received from the group level configuration.
4. Click **Save**.

Configure the uplink LAG

5. Click **Interfaces** > **Ports & Link Aggregations**. You see the interface list, including the lag256.
6. At the right end of the lag256 line, click the 'pencil' icon to edit the lag256.

7. In the Edit lag256 window, select **Port Members**. Check ports **1/1/27 and 1/1/28**.
8. Click **Save**.
9. Click the **Back** button to return to the main configuration screen.

Uplink Assignment (for Monitoring purpose only)

Aruba Central can provide dedicated statistics monitoring for the uplink ports of a device. In these steps you will set the uplink LAG member ports to 1/1/27 and 1/1/28.

10. In Aruba Central, navigate to Context: **Global**> Navigation: **Devices**> Top: **Switches** > Right top: **List**
11. Move your mouse over the line with sw-edge1. On the right-end, a link UPLINKS will appear.
12. Click the link **UPLINKS**.
13. In the **Assigned Uplink Ports** dropdown, select the ports **1/1/27 and 1/1/28**. These are the member ports of the uplink LAG.
14. Click **Assign**.

You have now completed the sw-edge1 onboarding configuration after ZTP was used for the initial setup.
Sw-edge1 is now connected with an LACP LAG with 2 member ports, the LAG is configured as a VLAN trunk and it supports VLANs 3 and 11.

Practice with sw-edge2

You may practice the per-device configuration for sw-edge2.

15. Configure these values at the device level:

- Hostname: **sw-edge2**
- LAG256 ports: **1/1/27, 1/1/28**

16. Configure the Aruba Central **uplink ports** with LAG256.

Use the previous section if you are unsure about the steps.

Configure PC Access Ports

Now that both access switches have been onboarded successfully into a UI group, you can use the UI to make configuration changes, such as the VLAN access port settings. In this section you will configure the sw-edge1 port for PC1 (1/1/1) to become an access port in VLAN 11.

17. In Aruba Central, navigate to Context: **Global**> Navigation: **Devices**> Top: **Switches** > Right top: **List**
18. Click the **sw-edge1**.
19. In the Navigation menu, click **Device**.
20. Click **Interfaces** > **Ports & Link Aggregations**. You will see the interface list.
21. Click port **1/1/1**, use the pencil icon at the right-hand side to **edit** the record.

> **NOTE:** You may need to scroll to the right.

22. Apply these settings to the port 1/1/1:

- Description : **pc1**
- VLAN Mode : **Access**
- Access VLAN : **11**

23. Click **Save**.

Verify the configuration.

24. In the Navigation menu, click **LAN**.
25. In the PORTS list, the port 1/1/1 should now be listed with the name pc1 and Native VLAN 11.

Port	Name	STATUS	LAG	MAC Address	VLAN	Native VLAN
1/1/1	pc1	Up		8c:85:c1:65:0e:e7		11

NOTE: It may take several minutes for Aruba Central to discover the ports on the switch. If you don't see the ports yet, you may continue the lab and check the ports at a later time.

26. On PC1, bounce the Lab NIC (disable/enabled) and verify the PC has received an IP address in VLAN 11.
27. On PC1, open a command prompt (cmd) and check the IP address using **ipconfig**.

```
ipconfig
```

```
C:\Users\student>ipconfig

Windows IP Configuration
...

Ethernet adapter Lab NIC - Edge-1 Port 1:

   Connection-specific DNS Suffix  . : aruba-training.com
   IPv4 Address. . . . . . . . . . . : 10.1.11.50
   Subnet Mask . . . . . . . . . . . : 255.255.255.0
   Default Gateway . . . . . . . . . : 10.1.11.1
```

Task 3: Configure Access Switches with Central UI Group MultiEdit

In this task you will use MultiEdit to make a configuration change to the access switches. While the UI configuration tiles are convenient for common features, not all configurations are available with the tiles.

Aruba Central provides a CLI-like interface to manage the AOS-CX switches: MultiEdit, which is part of a UI configuration group. In this task you use MultiEdit to make an example change to the devices.

Objectives

- Use Central MultiEdit to change the switch configuration.
- Verify the configuration synchronization status.

Steps

1. In Aruba Central, navigate to Context: **Groups / campus-sw-edge-ui** > Navigation: **Devices**> Top: **Switches** > Right top: **Config**
2. At the top of the screen, the MultiEdit slider is displayed.

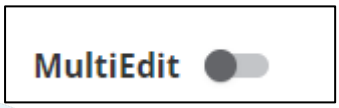

3. Move the slider to the right to access MultiEdit.

View CLI-like switch configuration - single switch

4. In the list of switches, select the **sw-edge1**.
5. At the right-bottom, click '**VIEW CONFIG**' in the option screen.

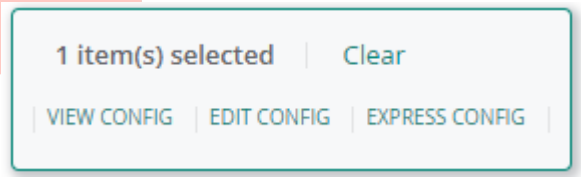

6. After a few moments, you will see the current switch configuration. In this screen you could easily select the text-based configuration and copy the text to save or compare the configuration.

Question: What is the description and VLAN set on interface 1/1/1?

Answer: Interface 1/1/1 was previously assigned description *pc1* and is an access port in VLAN 11.

```
26 interface 1/1/1
27     no shutdown
28     no routing
29     description pc1
30     vlan access 11
```

7. Click the **Back** button to return to the MultiEdit switch list.

View CLI-like switch configuration – multiple switches

When you need to compare two or more switch configurations, the MultiEdit interface can be used to quickly detect what values or settings have been applied to what switch.

8. Select *both* **sw-edge1 and sw-edge2** by clicking on the first, then on the second.

> **NOTE:** To unselect, click again on the device.

9. At the right-bottom, notice the text *2 items* selected, since you have two switches selected.
10. Click '**VIEW CONFIG**' in the option screen.
11. In the configuration screen, notice the first line:

```
hostname HOSTNAME
```

Question: Why is the HOSTNAME in all-capitals?

Answer: MultiEdit uses this all-capitals notation for values that are different on the selected switches.

12. Right-click on the HOSTNAME value. On the right-hand side, the "**View Parameter**" window is displayed.

Question: What is the value for hostname set on sw-edge1 and sw-edge2?

Answer: sw-edge1 and sw-edge2.

13. Click the **'X'** to close the View Parameter window.
14. Scroll down to interface 1/1/1.

Question: How can you see that the line 'description pc1' only exists on 1 switch?

Answer: At the end of the line, sw-edge1 is shown. This indicates that the line only exists on this switch. When there is no text at the end of the line, the line exists on all selected switches.

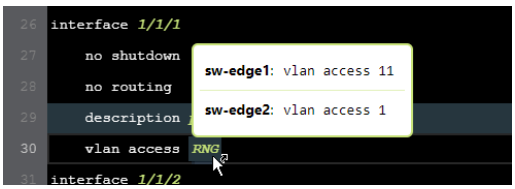

Question: What do you see when you right-click **RNG** in the line **vlan access RNG**? (You may also hoover with the mouse over the value RNG)

Answer: sw-edge1 has port 1/1/1 assigned to VLAN 11, while on sw-edge2 it is still assigned to VLAN 1

15. Click the **Back** button to return to the MultiEdit switch list.

Change switch configuration: multiple switches

You can also use the MultiEdit interface to make configuration changes. This can be convenient to push changes to multiple devices. With the parameter concept, it is still possible to push per-device settings as well in the MultiEdit interface.

In this section you will create the AP management VLAN 4 and assign port 1/1/2 on both sw-edge1 and sw-edge2 to this VLAN. This is a change that should be made on both switches. Next you will set the port description on sw-edge1 to ap1, and on sw-edge2 to ap2.

16. Select *both* **sw-edge1 and sw-edge2** by clicking on the first, then on the second.
17. Click **EDIT CONFIG** in the option screen.

Create VLAN 4

18. Scroll down to the line with ‘vlan 1’. Put the cursor after the ‘vlan 1’ and press <Enter>.

> **NOTE:** This will put you on a new blank line. Pay attention that the cursor is now left-justified - at the start of the line. This represents that you are currently at the global configuration context.

> **NOTE:** MultiEdit is a CLI-like interface, but it is not a traditional CLI like the switch CLI. Since you see the entire configuration, you will need to *assist* the CLI by entering a context with a space if you need to enter a sub-context.

19. Type ‘**vlan 4**’ and press <Enter>.

Question: What happened with the line position and the cursor?

Answer: MultiEdit automatically detected the new VLAN should have been placed after VLAN 3 and moved the line. The cursor is back at the global configuration level.

20. Press the <Space Bar>, this is the *hint* for MultiEdit that you want to *enter* the VLAN context.
21. Type **name v4-ap** and press <Enter>.

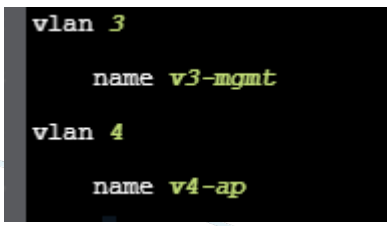

Question: What happened with the cursor?

Answer: MultiEdit detected that the cursor was inside the VLAN context, it kept the cursor at this context level.

Create VLANs 12-15

VLANs 11-15 are used to support the lab’s bridged WLANs. You have previously created the VLAN 11 in the UI. You will now create the VLANs 12-15 in the MultiEdit interface.

22. Move the cursor to the start of the line (global configuration) and configure VLANs 12-15, press <Enter>.

```
vlan 12-15
```

Assign port 1/1/2 to VLAN 4

Now you will assign port 1/1/2 as a VLAN trunk port with native VLAN 4 and assign a unique description on each switch.

23. Scroll down to interface 1/1/2.
24. Under **interface 1/1/2**, change the line *vlan access 1* to **vlan trunk native 4** and press <Enter>. VLAN 4 is the AP management VLAN. When APs boot, they send an untagged DHCP request, which is handled by VLAN 4 on the switch.

```
vlan trunk native 4
```

Since the cursor was under the interface context, it will now remain inside this context.

25. Configure the VLAN trunk allowed list. A VLAN trunk is required for bridged WLANs. The AP will support several WLANs (employee, guest, PSK), and it can tag the wireless traffic with a VLAN tag. These VLAN tags must be accepted (allowed) by the switch. In the lab, the bridged WLANs will use VLAN IDs 11-15.

```
vlan trunk allowed 4,11-15
```

> **IMPORTANT:** Pay attention for the **vlan trunk allowed** command in MultiEdit. The configured list will be the actual, final allowed VLAN list. If you entered **vlan trunk allowed 11-15,** you would not just add those VLANs, but you would also remove VLAN 4 from the list.
>
> This is different from the switch CLI, where the command **vlan trunk allowed 11-15** would simply add those VLANs to the allowed list, but it would not affect the other allowed VLANs on that trunk port.

26. Enter **description ap** on the next line and press <Enter>.

> **NOTE:** You can use command completion with the <Tab> key in MultiEdit.

27. Right-click the **ap** text. This reveals the “Modify Parameter” window on the right side.
28. Apply these values:

- For sw-edge1: **ap1**
- For sw-edge2: **ap2**

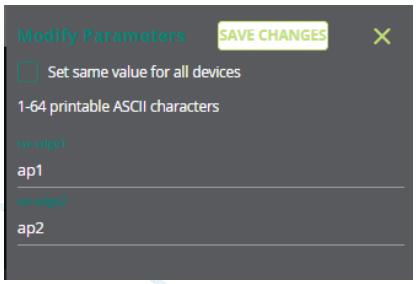

29. Click **Save Changes** to close the Modify Parameters window.

30. Verify in the MultiEdit CLI that port 1/1/2 now has **description DSC**. This indicates there are different values for the selected switches.

```
interface 1/1/2
    no shutdown
    no routing
    description DSC
    vlan trunk native 4
    vlan trunk allowed 4,11-15
```

You are now ready to push the configuration to the two switches.

31. Scroll down and click the **Save** button. Central returns to the switch device list.
32. In the list of switches, check the **Config Status** column, which dynamically updates as the configuration is pushed to the switches.

Devices (2)

Name	Firmware Version	Config Modified	Status	Config Status
sw-edge1	10.09.1040	Oct 20, 2022, 15:04:30	Online	Not in sync (Pushing ...
sw-edge2	10.09.1040	Oct 20, 2022, 15:02:38	Online	Not in sync (Pushing ...

33. Wait for the **Config Status** to be **Sync**. This indicates the configuration on the device is synchronized.

Devices (2)

Name	Firmware Version	Config Modified	Status	Config Status
sw-edge1	10.09.1040	Oct 20, 2022, 15:18:05	Online	Sync
sw-edge2	10.09.1040	Oct 20, 2022, 15:18:04	Online	Sync

You have now configured two switches using MultiEdit.

NOTE: MultiEdit changes the **device level** configuration. These changes will not be shown at the group-level configuration, but they will appear in the device-level configuration tiles in the UI. (For features that are supported by the UI tiles).

NOTE: If you want to switch back to using the UI tiles, at the top of the screen, set the slider for MultiEdit to be disabled.

Task 4: Configure Access Switches with Central Template Group

In this task you will configure a template group for the edge switches. The template has been prepared by your senior colleague; you should be able to configure the group and apply the template to the group.

Objectives

- Configure edge switches with a template group.

Steps

1. Use the lab dashboard to open a console session to sw-edge1.
2. Take note of the serial number.

```
show system
```

Example output, your output may be different.

```
6300# show system
Hostname           : 6300
System Description : FL.10.09.1040
System Contact     :
System Location    :

Vendor             : Aruba
Product Name       : JL666A 6300F 24G CL4 PoE 4SFP56 Sw
Chassis Serial Nbr : SG11KN501V
Base MAC Address   : 8c85c1-4920c0
ArubaOS-CX Version : FL.10.09.1040

Time Zone          : UTC

Up Time            : 1 hour, 13 minutes
CPU Util (%)       : 18
Memory Usage (%)   : 21
```

3. Take note of the serial number.

sw-edge1 serial:

_____

4. Take note of the Product Name part number (JL666A, JL662A or similar)

_____

5. In Aruba Central, navigate to Context: **Global** > Navigation: **Organization** > Top: **Network Structure** > Tile: **Groups**
6. Create new Group with name **campus-sw-edge-tpl**, select **Switches.**
7. Slide the bar to **enable** the **templates** and click **Next**.
8. For Switches in group, select **AOS-CX only** and click **Add**.

9. Under **All connected devices**, select both edge 6300 switches and move them to the new group **campus-sw-edge-tpl**.

Assign per-device hostname variable

10. In Aruba Central, navigate to Context: **Groups / campus-sw-edge-tpl** > Navigation: **Devices**> Top: **Switches** > Right top: **Config**
11. Click **Variables**.

Each device that has been assigned to this group (in this example two access switches) is listed with three variables. The Device Serial Number column contains the serial number.

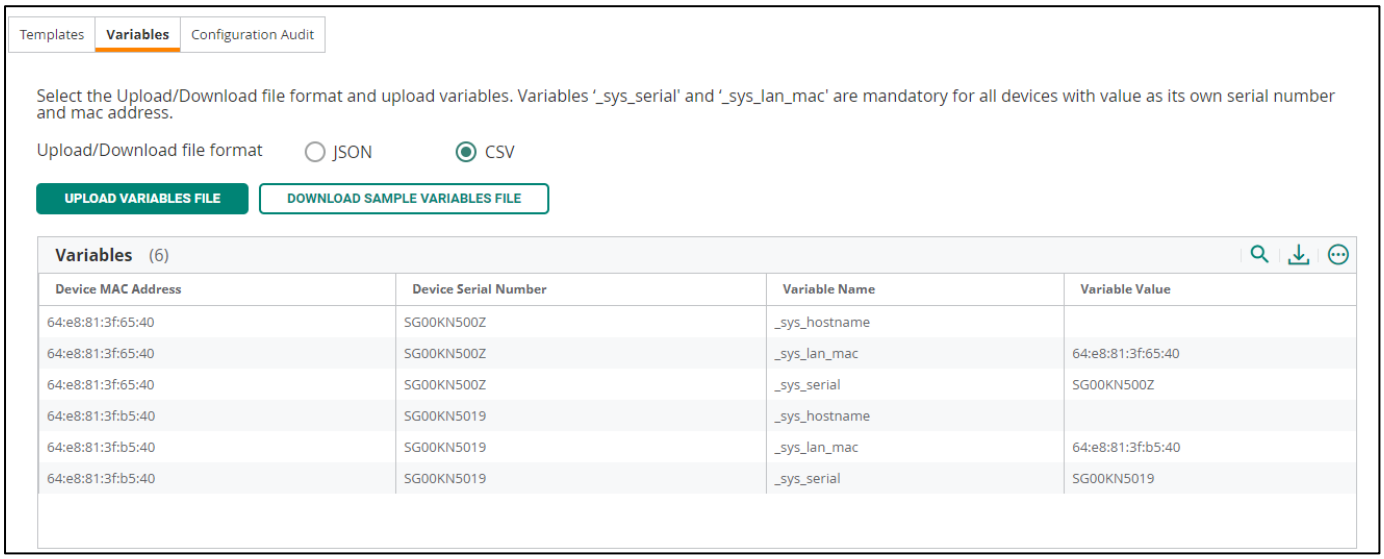

12. Hover with the mouse over the line that contains your sw-edge1 serial and the variable _sys_hostname. At the end of the line, a pencil icon will appear.

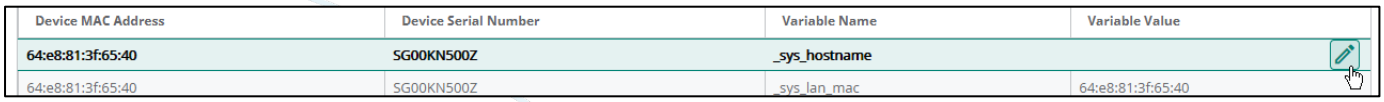

> **NOTE:** Depending on the lab reset status, the _sys_hostname may already have been set to sw-edge1. If this is the case, you can skip the next three steps.

13. Click the pencil icon to edit the value.
14. In the New Value field, enter **sw-edge1**

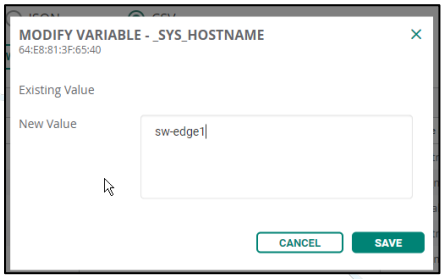

IMPORTANT: The access switch template contains conditions that check on the correct hostname. Please double-check the hostname is exactly **sw-edge1 (use - , not the underscore _)**

15. Click **Save** to submit the value.

The other serial will be your sw-edge2.

16. Verify that the other serial **_sys_hostname** value is set to **sw-edge2** (double-check the name), update the value if necessary.

It takes a few seconds to submit and process the values.

17. Refresh your browser to refresh the variable list.
18. Verify the resulting variable list shows both hostnames **sw-edge1** and **sw-edge2**.

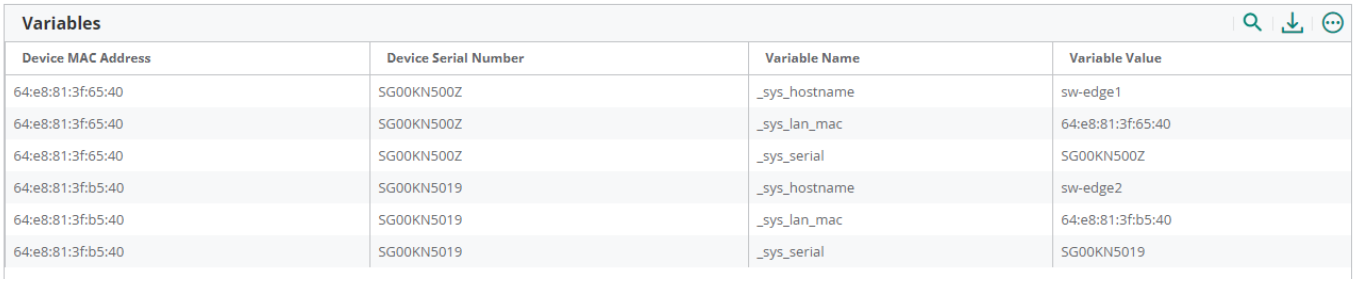

Variables

Device MAC Address	Device Serial Number	Variable Name	Variable Value
64:e8:81:3f:65:40	SG00KN500Z	_sys_hostname	sw-edge1
64:e8:81:3f:65:40	SG00KN500Z	_sys_lan_mac	64:e8:81:3f:65:40
64:e8:81:3f:65:40	SG00KN500Z	_sys_serial	SG00KN500Z
64:e8:81:3f:b5:40	SG00KN5019	_sys_hostname	sw-edge2
64:e8:81:3f:b5:40	SG00KN5019	_sys_lan_mac	64:e8:81:3f:b5:40
64:e8:81:3f:b5:40	SG00KN5019	_sys_serial	SG00KN5019

Template configuration

Now you will upload the prepared template for the edge switches. Based on the hostname, each switch will receive the correct configuration from the template.

19. Click on **Templates**.
20. At the right-top, click the + icon to add a new template.
21. For Template Name use sw-edge.
22. Device Type should be Aruba CX, based on the group configured switch type.
23. Click **Next**.
24. On MGMT PC, in the ACAF Student Files, open **lab07.02-sw-edge-tpl-JLxxxx.txt** for your switch model**.**

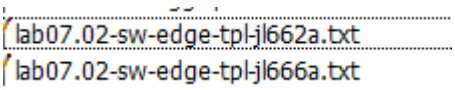

If your switch model was **JL666a**, you should open lab07.02-sw-edge-tpl-**jl666a**.txt

Refer to the **show system** output in step2 of this task if you are unsure about the product number.

25. Copy all the contents.
26. Paste the contents of the file into the template.
27. Click **Save**.

Verification

28. Click **Configuration Audit**. The Configuration status will show 2 devices Not In Sync.

Since the template changes the switch management IP address, the switches will need to re-establish their connection to Aruba Central and check in with the new IP address. This will take 1-2 minutes to complete.

29. After about 2 minutes (when the switches have reconnected to Aruba Central), the status should show 0 devices Not in Sync.

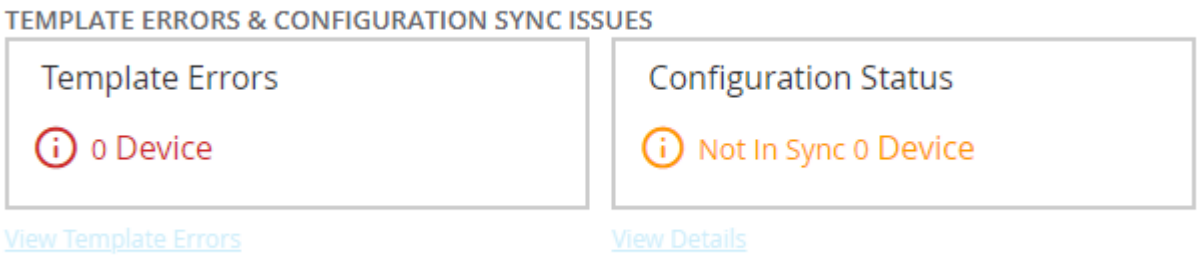

You have completed the Lab!

Lab 08.01 Onboarding APs

Overview

Now that you have completed the wired network rollout, you are ready to deploy the wireless network.

The first phase of the rollout is to make sure:

- The APs are powered up.
- The APs are connected to the correct AP management VLAN.
- The respective switch ports are configured with a VLAN trunk port to support the VLANs for the bridged SSIDs.

Once the APs can reach Aruba Central, you need to prepare a configuration group and apply some basic settings to the APs, such as the country code and the AP names.

In the last section of this lab, you will prepare a floorplan in Aruba Central and position APs on the floorplan.

Objectives

- Verify AP deployment.
- Verify AP connection to Aruba Central.
- Assign the AP name and country code.
- Assign the AP site information.
- Configure a floorplan.

Task 1: Verify AP wired onboarding

Switch port configuration for APs.

The switch port that will support APs with bridged SSIDs needs to support the user VLANs for these SSIDs. This is achieved by configuring the port as a VLAN trunk port.

The AP will send out a DHCP request to obtain a management IP address. This DHCP request is sent untagged by default. This means that the DHCP request will hit the switch port on the configured native untagged VLAN. The switch port trunk native untagged VLAN should be set to the desired AP management VLAN. In this lab environment, VLAN 4 is the AP management VLAN.

You can configure switch ports manually or automatically. Many enterprise customers prefer to use a centralized AAA solution, such as Aruba ClearPass, to assign switch port configurations. In a smaller deployment, you can use device profiles, which you configure locally on the switch. With device profiles, the switch can recognize APs based on either LLDP or MAC address information and assign the correct port configuration. When an AP is removed, the port configuration reverts to the port configuration as set in the running-configuration.

PoE
In a production deployment, the APs will typically be powered by the access switches. In this remote lab environment, the APs are powered in a centralized cage by an external PoE device. This means you will not see any PoE activity on your access switch. It also means that a shutdown of the port on your lab switch will not be noticed by the AP, since it is connected to a transit (intermediate) switch, which provides PoE power to the APs.

For example, if the AP is rebooted, your access switch port will not lose power and move from a down to an up state (i.e., the AP is not directly connected to the switch shown in the lab topology).Please remember this when troubleshooting!

Objectives

- Verify AP connection on the switch.

Steps

PoE

Although our lab APs are not powered by your access switches, it is good to review the PoE commands.

1. Use MGMT PC to open an SSH connection to the **sw-edge1**.
2. Review the overall PoE power.

```
show power-over-ethernet
```

```
sw-edge1# show power-over-ethernet

System Power Status for member 1

  PoE Power Status              : No redundancy
  Operational Power Status      : No redundancy
```

```
  Total Available Power        : 370.00 W
  Total Failover Pwr Avl       :   0.00 W
  Total Redundancy Power       :   0.00 W
  Total Power Drawn            :   0.00 W
  Total Power Reserved         :   0.00 W
  Total Remaining Power        : 370.00 W
  Trap Threshold               : 80 %
  Trap Enabled                 : Yes
  Always-on PoE Enabled        : 1/1
  Quick PoE Enabled            : None


 Internal Power:
        Total Power
  PS    (Watts)        Status
  ----- ------------   --------------------
  1/1   950            OK
```

Question: What is the total watts power available on power supply PS1/1?

Answer: PS 1/1 supports 370Watts for PoE.

3. Review the PoE interface status.

```
show power-over-ethernet brief
```

```
sw-edge1# show power-over-ethernet brief

  Member 1 Power Status
    Available: 370.00 W  Reserved: 0.00 W  Remaining: 370.00 W
    Always-on PoE Enabled:1/1
    Quick PoE Enabled:None

PoE      Pwr Power    Pre-std Alloc         PSE Pwr PD Pwr PoE Port      PD     Cls
Type
Port     Ena Priority Detect  Act           Rsrvd   Draw   Status        Sign
-------- --- -------- ------- ------------- ------- ------ ------------- ------ --- -
---
1/1/1    Yes low      Off     usage         0.0 W   0.0 W  searching     N/A    N/A
N/A
1/1/2    Yes high     Off     usage         0.0 W   0.0 W  searching     N/A    N/A
N/A
1/1/3    Yes low      Off     usage         0.0 W   0.0 W  searching     N/A    N/A
N/A
1/1/4    Yes low      Off     usage         0.0 W   0.0 W  searching     N/A    N/A
N/A
...
```

Question: What does *searching* indicate?

Answer: There is currently no PoE device connected to the port, the switch is searching for a PoE capable device on the port.

Device Profiles

You can use device profiles to automatically assign the correct switch port configuration when an AP is connected to the associated port. Your senior colleague has included the device profile in the switch template configuration.

4. On sw-edge1, review the running-configuration of port 1/1/2, which is connected to AP1.

```
show running-config interface 1/1/2
```

```
sw-edge1# show running-config interface 1/1/2
interface 1/1/2
    no shutdown
    description ap1
    no routing
    vlan access 1
    spanning-tree bpdu-guard
    no aaa authentication port-access allow-lldp-auth
    aaa authentication port-access client-limit 5
    aaa authentication port-access mac-auth
        enable
    exit
```

Question: What is the configured VLAN mode on port 1/1/2?

Answer: Port 1/1/2 is configured as access port in VLAN 1.

5. Review the operational VLAN status for port 1/1/2.

```
show vlan port 1/1/2
```

```
sw-edge1# show vlan port 1/1/2

--------------------------------------------------------------------------------
VLAN  Name                                Mode            Mapping
--------------------------------------------------------------------------------
4     VLAN4                               native-untagged port-access
11    VLAN11                              trunk           port-access
12    VLAN12                              trunk           port-access
13    VLAN13                              trunk           port-access
14    VLAN14                              trunk           port-access
15    VLAN15                              trunk           port-access

Overridden VLAN list: 1
```

Question: What do you notice?

Answer: VLAN 4 is the native-untagged VLAN. Other VLANs are allowed as well on the trunk port.

Question: Why is this different from the port configuration?

Answer: A device profile applied a different configuration to the port since it discovered an AP MAC address. If the AP would disappear, the port would revert to its local port configuration.

6. Review the list of configured device profiles.

```
show port-access device-profile
```

```
sw-edge1# show port-access device-profile

    Profile Name            : aruba_ap
    LLDP Groups             :
    CDP Groups              :
    MAC Groups              : aruba_ap
    Role                    : dev_aruba_ap
    State                   : Enabled
```

7. Review the applied list of device profiles

```
show port-access device-profile interface all
```

```
sw-edge1# show port-access device-profile interface all
  Port 1/1/2, Neighbor-Mac  20:4c:03:c6:09:50
    Profile Name:           : aruba_ap
    LLDP Group:             :
    CDP Group:              :
    MAC Group:              : aruba_ap
    Role:                   : dev_aruba_ap
    State:                  : applied
    Failure Reason:         :
```

Question: What is the configured MAC Group?

Answer: *aruba_ap*

Question: What is the applied role?

Answer: *dev_aruba_ap.* This role contains all the settings that will be applied to the port.

8. Review the running configuration for the port-access section.

```
show running-config port-access
```

```
sw-edge1# show running-config port-access
port-access role aruba_ap
port-access role dev_aruba_ap
    auth-mode device-mode
    poe-priority high
    trust-mode dscp
    vlan trunk native 4
    vlan trunk allowed 4,11-15
port-access device-profile aruba_ap
    enable
    associate role dev_aruba_ap
    associate mac-group aruba_ap
aaa authentication port-access mac-auth
    enable
interface 1/1/2
```

```
    no aaa authentication port-access allow-lldp-auth
    aaa authentication port-access client-limit 5
    aaa authentication port-access mac-auth
        enable
```

9. Review the configured mac-groups.

```
show running-config mac-group
```

```
sw-edge1# show running-config mac-group
mac-group aruba_ap
     seq 10 match mac-oui 20:4c:03
     seq 20 match mac-oui 18:7a:3b
```

Verify the application of the device profile

In this section you will force an access port client *logoff* on interface 1/1/2. This allows you to see events related to the device profile when a matching device is detected.

10. On sw-edge1, review the list of active port access clients.

```
show port-access clients
```

```
sw-edge1# show port-access clients

Port Access Clients

Status Codes: d device-mode, c client-mode, m multi-domain

--------------------------------------------------------------------------------
--------------------
  Port     MAC-Address       Onboarding     Status      Role
Device Type
                             Method
--------------------------------------------------------------------------------
--------------------
d 1/1/2    20:4c:03:c6:09:50 device-profile Success     dev_aruba_ap
```

11. Enable terminal-monitor to see live events in your session.

```
terminal-monitor
```

```
sw-edge1# terminal-monitor
Terminal-monitor is enabled successfully
```

NOTE: Terminal-monitor is not supported on the console. If you received an error, make sure to use MGMT PC and open an SSH connection to the sw-edge1 instead of a console connection.

12. Logoff the client on port 1/1/2. Since the AP almost continuously sends some traffic, it will be back online within a few seconds.

```
port-access log-off client interface 1/1/2
```

```
sw-edge1# port-access log-off client interface 1/1/2
```

NOTE: This command should be executed in the manager context (#), if you access the configuration mode, return to manager mode with the **end** command. If you do want to run manager mode commands from the configuration context, you can prepend the command with the keyword **do**. For example:
do port-access log-off client interface 1/1/2

13. Review the resulting event logs.

```
2022-07-29T10:48:24.346774+0000 port-accessd[4355] <INFO>
Event|10502|LOG_INFO|CDTR|1|Port 1/1/2 is blocked by port-access
2022-07-29T10:48:24.394887+0000 ops-switchd[686] <INFO>
Event|2107|LOG_INFO|CDTR|1|The mode for port 1/1/2 changed from native-untagged to
access on VLAN 1
```

```
2022-07-29T10:48:24.722076+0000 port-accessd[4355] <INFO>
Event|10503|LOG_INFO|CDTR|1|Port 1/1/2 is unblocked by port-access
2022-07-29T10:48:24.743632+0000 ops-switchd[686] <INFO>
Event|2107|LOG_INFO|CDTR|1|The mode for port 1/1/2 changed from access to native-
untagged on VLAN 4
2022-07-29T10:48:24.751112+0000 hpe-mstpd[3446] <INFO>
Event|2012|LOG_INFO|CDTR|1|CIST - Topology Change generated on port 1/1/2 going in to
forwarding
```

14. Stop the terminal monitor.

```
no terminal-monitor
```

Verify the AP Management VLAN MAC addresses

To verify that the APs are online in the correct VLAN, you can check the MAC Address table of VLAN 4.

15. Check the MAC address table of VLAN 4.

```
show mac-address-table vlan 4
```

```
sw-edge1# show mac-address-table vlan 4
MAC age-time            : 300 seconds
Number of MAC addresses : 5

MAC Address          VLAN     Type                        Port
-------------------------------------------------------------
20:4c:03:c6:09:50    4        port-access-security        1/1/2
b8:d4:e7:da:80:00    4        dynamic                     lag256
b8:d4:e7:d9:e5:00    4        dynamic                     lag256
12:01:00:00:01:00    4        dynamic                     lag256
20:4c:03:c6:04:90    4        dynamic                     lag256
```

Question: Do you see a MAC address starting with 20:4c:03 on port 1/1/2?

Answer: Yes.

Task 2: Prepare Aruba Central AP Group

In this task you will prepare a group in Aruba Central for the Access Points.

Objectives

- Create an AP group in Aruba Central.

Steps

1. Open or switch to your Aruba Central session.
2. In Aruba Central, navigate to Context: **Global**> Navigation: **Organization**> Top: **Network Structure**> **Groups**
3. Group name **campus-wifi-ui**, select only APs. Do not enable templates.

> **NOTE:** For AOS10, UI groups are the recommended best practice.

4. Click **Next**.
5. For Architecture, leave **ArubaOS 10.** For Network Role leave **Campus/Branch**.

> **NOTE:** You could choose AOS8 to configure an Instant AP based deployment, but this lab uses the cloud AP deployment with AOS10.

6. Click **Add**.

Configure Basic Group settings

In the next steps, you will apply some essential settings for the new group.

7. In Aruba Central, navigate to Context: **Groups / campus-wifi-ui** > Navigation: **Devices**> Top: **Access Points**> Right top: **Config**

Since this is the first time you access the AP Group configuration, and since no password has been set, the system prompts you to set the admin password. This is the password for the 'admin' user account on the APs that belong to this group.

8. For the password enter **Aruba123!**
9. Confirm the password and click **Set Password**.

By default, the basic configuration UI is displayed. In these labs you will also use some other settings, so you can enable the Advanced View.

10. Click **Show Advanced**. You see additional configuration tabs.
11. Click **System.** The **General** section opens by default.

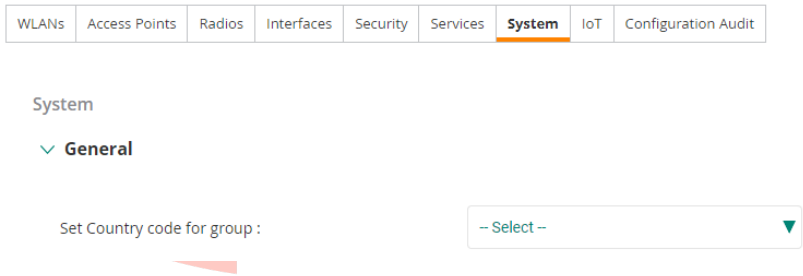

12. Configure these values under **General**:

- Country Code for Group: **US**
- Timezone: **Eastern-Time UTC-05**
- NTP server: **10.254.1.21**

13. Click Save Settings.

14. Navigate to the **System** > **Administrator** section.

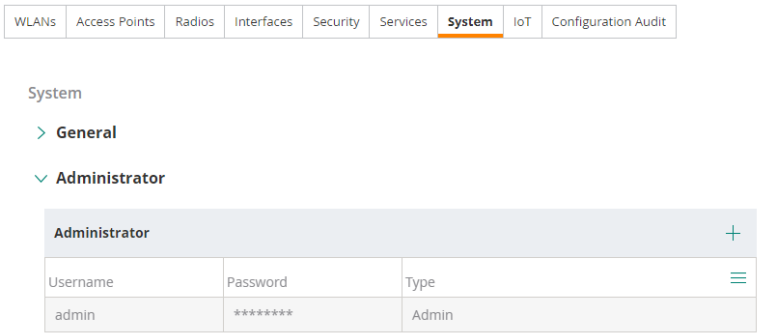

This allows you to change the *admin* password that you previously configured when accessing the group for the first time.

> **NOTE:** Make sure to leave the admin password to **Aruba123!**

Configuration Audit Trail

Use the audit trail to track AP configuration changes in Aruba Central. This helps to verify that the UI was translated to the correct AP CLI commands when Aruba Central sends the commands to the APs.

15. In Aruba Central, navigate to Context: **Groups / campus-wifi-ui** > Navigation: **Audit Trail**> Top: **Audit Trail**

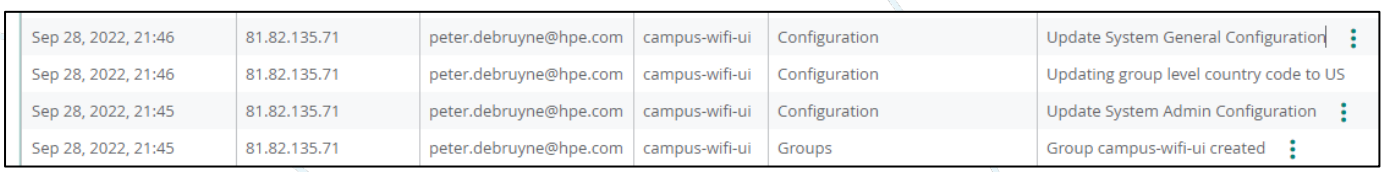

Sep 28, 2022, 21:46	81.82.135.71	peter.debruyne@hpe.com	campus-wifi-ui	Configuration	Update System General Configuration
Sep 28, 2022, 21:46	81.82.135.71	peter.debruyne@hpe.com	campus-wifi-ui	Configuration	Updating group level country code to US
Sep 28, 2022, 21:45	81.82.135.71	peter.debruyne@hpe.com	campus-wifi-ui	Configuration	Update System Admin Configuration
Sep 28, 2022, 21:45	81.82.135.71	peter.debruyne@hpe.com	campus-wifi-ui	Groups	Group campus-wifi-ui created

16. Click the three dots icon at the end of the various records to see the configuration commands that you have generated.

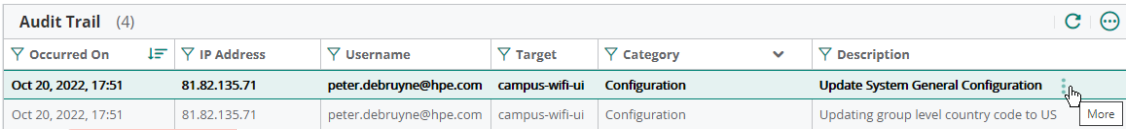

Example for the **System General Configuration** details.

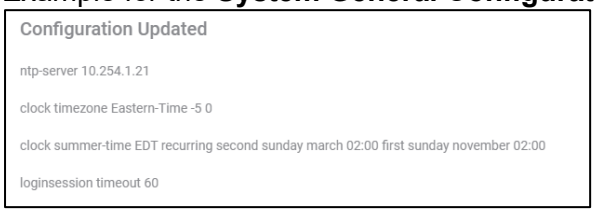

17. **Close** the audit details window.

Task 3: Move APs to campus-wifi-ui Group

In this task you will move the APs to the group that you just prepared. When you move the APs, Central pushes basic settings to the APs – settings that you prepared at the group level, like the admin password and country settings.

Objectives

- Move APs to the correct AP Group.

Steps

Review the APs are online in Central

First, you will verify that APs are connected to Aruba Central. You will make a record of their serial numbers and MAC address, which are used during labs to identify the different APs.

1. In Aruba Central, navigate to Context: **Global**> Navigation: **Devices**> Top: **Access Points**

This is the global list of APs that have ever come online on this account. You should see 2 APs with their MAC address as the device name.

2. Compare the APs MAC and Serial with the inventory information you have collected during Lab 1.

AP1 MAC _____
SERIAL _____

AP2 MAC _____
SERIAL _____

3. Use the three dots in the column header to enable the **Group** column.

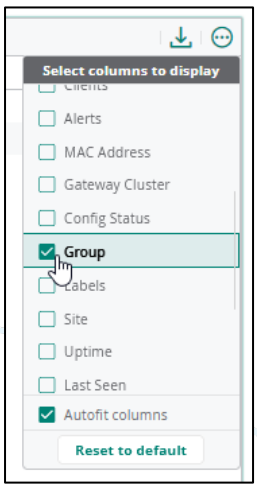

Question: To what group are these APs assigned?

HINT: You may need to scroll to the right to see the group column.

Answer: They are assigned to the group **default**.

Move the APs

4. In Aruba Central, navigate to Context: **Global**> Navigation: **Organization**> Top: **Network Structure**> **Groups**
5. Under **All connected devices**, select both APs.
6. On the right-hand side, you see a popup with the **Move Devices** action button.
7. Click the **Move Devices** button.
8. Click the **Destination Group** field.
9. Select the group **campus-wifi-ui.**
10. Click the **Move** button to continue.

Verify the Move

11. In Aruba Central, navigate to Context: **Global**> Navigation: **Audit Trail**

Question: Do you see records with Moving Device in the description column?

Answer: Yes, there are two entries. If you don't see two entries, wait a few seconds and use the refresh button at the right-top of the page.

Question: What is listed under the Target column for these records?

Answer: The serial number of the AP.

Assign APs to Site

Assign both APs to the site **site-campus-main.** A label is not required for the APs.

12. In Aruba Central, navigate to Context: **Global**> Navigation: **Organization**> Top: **Network Structure**> **Sites**
13. Select both APs (use the CTRL key to select both), drag them to the site **site-campus-main**.
14. Confirm the site assignment with **Yes**.

Per-AP Settings - AP Name

To simplify monitoring and troubleshooting, it is a good idea to give each AP its own, unique name. This is a *per-AP setting*. This means that the value is stored and saved on the AP: even when if you would move the AP to another Aruba Central group, this setting would stay the same.

You will now access the Per-AP settings for the 2 new APs.

15. In Aruba Central, navigate to Context: **Groups / campus-wifi-ui** > Navigation: **Devices**> Top: **Access Points**> Right top: **Config**

> **NOTE:** Although named “Per AP”, these settings are typically still configured using the group context.

> **NOTE:** You collected inventory information in lab 1 (**mfginfo)** showing the MAC address and serial number from the AP consoles. Use this to assign the correct names to the APs.

16. In the configuration menu, click **Access Points**. You see a list with your two APs.
17. Move the mouse over the first AP. Use the **pencil** icon at the end of the line to **EDIT** the record.
18. Set the **name** to **ap1**.
19. Click Save Settings.
20. Repeat this procedure for ap2 to assign the **name ap2**.
21. Check the latest entry in the audit trail (Navigation: **Audit Trail**) for the per-AP change. You may need to refresh the page to see the latest audit entries. It may take a moment for the records to appear.

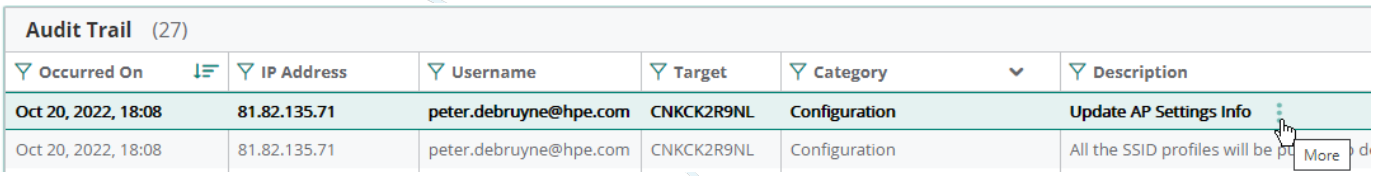

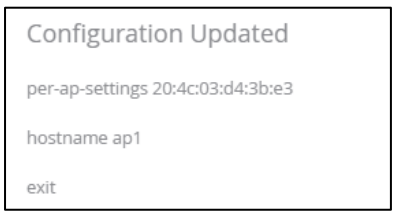

22. **Close** the window with the audit details.

Verify the Aruba Central connection and AP settings

In these steps you will review the AP to Aruba Central (cloud) connection.

23. Use the lab dashboard to open a console session to AP1. Login with admin / Aruba123!
24. Review the Aruba Central cloud connection.

```
show ap debug cloud-server
```

```
ap1# show ap debug cloud-server

IAP mgmt mode               :athena-mgmt
cloud config recved         :TRUE
autojoin mode               :disable
state diff                  :disable
Device Cert status          :SUCCESS
Cert Verify                 :enable
Domain Name Verify          :enable
Device info send            :SUCCESS
Aruba Central server                :device-uswest4.central.arubanetworks.com
Aruba Central server path           :/ws
Aruba Central proxy server          :None
Aruba Central redirect from         :device-uswest4.central.arubanetworks.com
Aruba Central Protocol              :WSS
Aruba Central uptimes               :8m:24s
Aruba Central status                :Login_done

Cloud Debug Statistics
----------------------
Key                         Value
---                         -----
Connect establish success   1(1)
Login done times            1(1)
Connect retry times         1(1)
Device Info send            1(1)
Domain list receive         1(1)
Domain response send        1(1)

Cloud Last connect status
-------------------------
Last connect ID         :1
Last connect time       :2022-10-20 12:04:16
Last connect trigger    :athena redirect

Cloud Last login done status
------------------------
Last connect done       :2022-10-20 12:04:19
```

25. Review the running configuration.

```
show running-config
```

```
ap1# show running-config
version 10.3.1.0-10.3.1
```

```
virtual-controller-country US
virtual-controller-key d079d04f01be93b88362d5b4428638522776f58654d2f73bec
name campus-wifi-ui
terminal-access
loginsession timeout 60
ntp-server 10.254.1.21
clock timezone Eastern-Time -05 00
clock summer-time EDT recurring second sunday march 02:00 first sunday november 02:00
rf-band all
...
extended-ssid

hash-mgmt-password
hash-mgmt-user admin password hash
9a0b99b502a062503b4051ef3d548d660821e55b5cfdbea45034201b3acb6e2bc9f7627fe2
...
```

Question: What is the Virtual-controller country?

Answer: As you have configured on the group, this is set to US.

Question: What is the NTP server?

Answer: 10.254.1.21

Question: Do you see a line with **hash-mgmt-user**?

Answer: Yes, this is the admin password, which Aruba Central applied to the configuration.

26. Review the per-AP settings. These values are stored in AP flash memory, separate from the running/startup configuration.

```
show ap-env
```

```
ap1# show ap-env

Antenna Type:Internal
Need USB field:Yes
name:ap1
```

Question: Do you see the AP name?

Answer: Yes, the name value is set to *ap1*.

27. Close the console connection.

Task 4: Configure the Site Floorplan with APs

In this task you will create a floorplan for the campus-site. In the floorplan, you can place the two deployed APs. This will allow Aruba Central to provide context for associated wireless clients and visualize their location.

Objectives

- Create a floorplan for a site in Aruba Central
- Assign APs to a floorplan

Steps

Verify AP Site Assignment

Before adding the floorplan, you will verify that your APs have been assigned to the correct site. This should have been completed in an earlier lab, so this is only a verification.

1. In Aruba Central, navigate to Context: **Sites / site-campus-main** > Navigation: **Devices**> Top: **Access Points**
2. Verify that both **ap1** and **ap2** are listed.

> **NOTE:** If one of the APs is not listed, refer to the previous task in this lab “Move APs to *campus-wifi-ui* group”. This is where you assigned APs to the correct site.

Configure the Site Floorplan

Floorplans are bound to a physical building or location. Thus, the floorplan configuration is situated under an Aruba Central *site* object.

3. In Aruba Central, navigate to Context: **Sites / site-campus-main** > Navigation: **Overview**> Top: **Floor Plan.**

By default, there is no floorplan defined for a site.

4. Click **Add Floor** to access the floor plan Edit Mode.

Measurement System

5. On the right-hand site, check the Ceiling Height value.

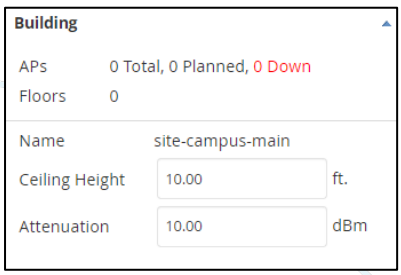

Question: Is this value in feet or meters?

Answer: By default, the imperial system is applied: height is reported in feet (ft).

Optional: Change to Metric System

Make sure to continue the lab after this optional section!

If you are based in a region that uses the metric system, you can follow the next steps. Other students can skip this optional section and move to *Create New Floor Plan* section on the next page.

To change to the metric system, you will need to add a floorplan once, save it, and exit. You will be able to change the measurement system in the Floorplan view.

The next steps only apply if you want to change to the metric system!

6. Right-click the area and click **New Floorplan.**

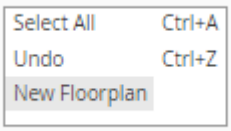

7. For Floorplan file, click **Choose File**.
8. From the student files, select the file **lab08.01-floorplan.png**.

> **NOTE:** You can access Aruba Central using the MGMT PC to upload the floorplan or your instructor may provide you access to the student files.

9. You may leave the floor name and number at their default values.
10. Click **Save**. This uploads the file and starts the new floor wizard, which takes a few seconds to complete.
11. You don't need to complete the steps of the wizard for this lab, click **Finish**.
12. Next click **Exit Edit Mode** to return to the view mode.
13. At the right-top, click the **view** icon to open the View Settings.

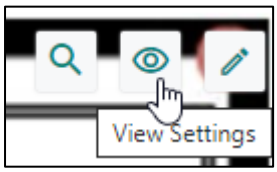

14. At the bottom, there is a toggle for Meter and Feet. Click **Meter** radio button to set the metric system.

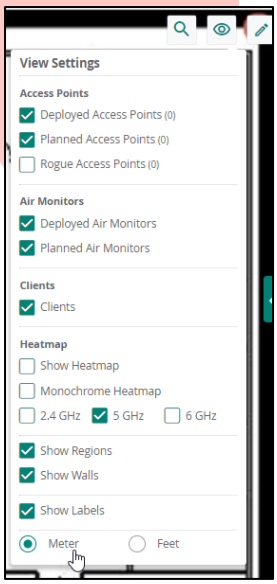

15. At the top-right, click **List** to see a list of floors. Your temporary floor is shown.
16. At the end of the record, use the **trashcan** icon to remove the floor.

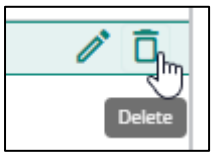

17. Confirm by clicking **Delete**.
18. At the right-top, click **Summary** to switch the view.

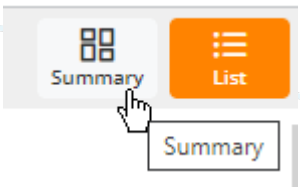

19. Click **Add Floor** to access the floor plan *Edit Mode*.

End of the optional section (Change to the metric system).

Make sure to continue the lab from this point!

Create New FloorPlan

These steps apply to all students again (both metric and imperial system).

20. Right-click the gray background area and click **New Floorplan.**

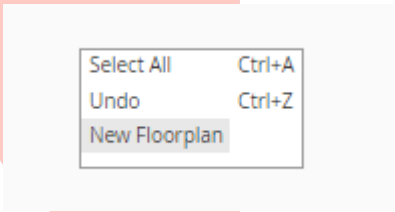

21. For Floorplan file, click **Choose File**.
22. From the ACAF Student Files folder, select the file **lab08.01-floorplan.png**.
23. You may leave the floor name and number to their default values.
24. Click **Save** to upload the file and start the new floor wizard.

Set Scale

You must set a known distance on the floorplan, this will be used by Aruba Central to perform distance calculations on the map, such as the locations of devices on the RF heatmap.

25. Click the **Measure** button. The mouse pointer changes to a plus sign.
26. Click and drag from the top-left corner of the building to the top-right corner.
27. When you release the mouse button, the Enter Distance window will popup.
28. Enter **100 feet** or **30 meters** (depending on your measurement configuration)
29. Click **OK**.
30. Click **Next.** This takes you to the **Region** step.

> **NOTE:** If you need to perform AP deployment planning, this can be configured using the Planning Regions feature. In this lab you will focus on monitoring the active APs, so there is no need to create planning regions.

Complete the Floor Plan wizard

31. Click **Next** to move to CAD Layer step. If a CAD drawing is available for the floor, you can upload it here. There is no need to change this in the lab, however.
32. Click **Next** to move to the **Access Points** step.
33. Click **Add deployed APs**. Both your APs should be shown.

> **NOTE:** If your APs are not listed, they have not been assigned to this site. Only APs that have been assigned to this site can be added to the floorplan of this site.

34. Drag and drop each AP to the map.

> **NOTE:** Since this is a lab environment, the actual position on the map does not matter.

35. Click **Finish** to complete the new floor wizard.

Edit an existing Floor

Once a floor has been added, you can still add or change the APs or floor properties.

> **NOTE:** By default, the location of APs is locked. This prevents accidentally moving an AP on the map.

36. To move an AP, you must first unlock the dragging function. Click the **lock** icon.

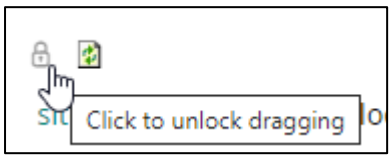

37. Now you can **drag and drop an AP** to a new location on the map.

Heatmap

Aruba Central can display the wireless heatmap based on AP locations. This is based on a calculation; therefore, it is important to correctly position APs on the map.

38. Click somewhere on the map to activate the floor Properties view. On the right side, the **Floor 1** properties should be displayed.

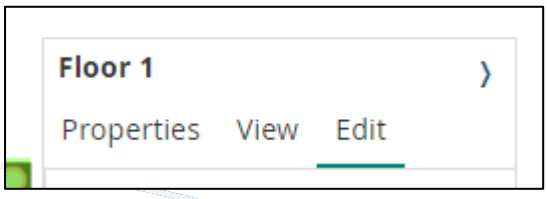

39. Click on **View**.
40. Clicking the **Heatmap** label enables or disables the heatmap view. Make sure the heatmap is *enabled*.

> **NOTE:** The heatmap can only be displayed when Aruba Central has received the AP Channel and Transmit power from the AP. As the AP has not been configured with a WLAN yet, the radios are still disabled, and thus you will not see a heatmap yet. You may just continue with the lab, however.

This is an example of how the heatmap will appear once the AP radios are enabled. This will be done in the next lab activity.

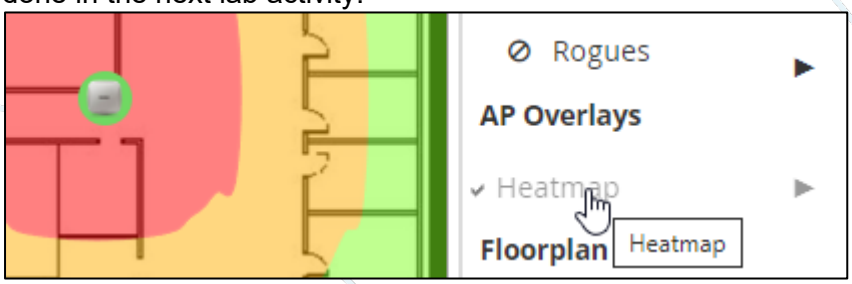

41. Expand the **Heatmap** section.
42. This section allows you to choose the desired band for the heatmap – 2.4, 5, or 6GHz.

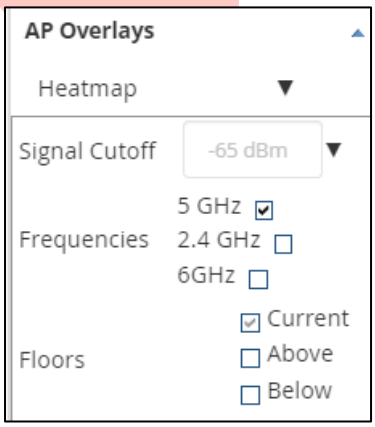

43. After you add or move an AP, you can manually refresh the heatmap display.

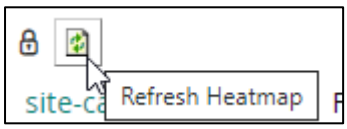

44. Hover your mouse over an AP to see a popup with current channel information.

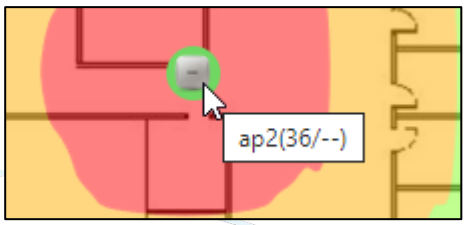

NOTE: This channel view option will not be available yet, as the AP radios are still disabled.

45. Click an AP to an AP-level window.

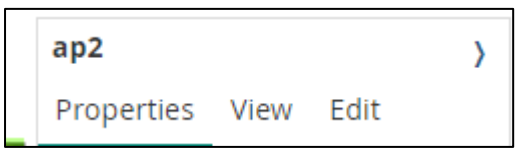

46. Click the **Properties** link to view the channel and transmit power.

Example:

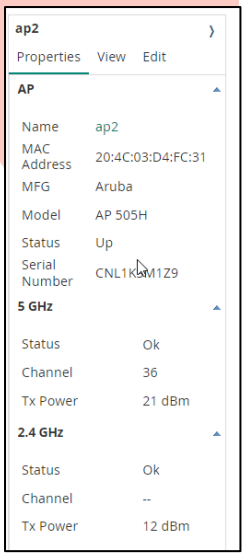

47. Click **Exit Edit Mode** to switch to the View Mode.

View Mode

In the default view mode of the floor plan, several View Settings can be configured.

48. On the right-top, click the **View Settings** icon.

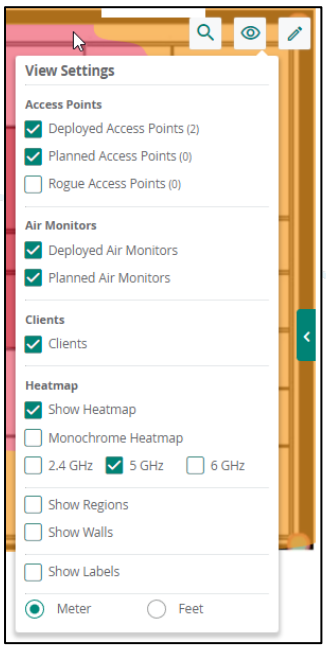

49. The View Settings window allows you to choose whether clients, APs, and/or Air Monitors should be displayed on the map.

This concludes the Floorplan configuration.

You have completed this Lab!

Lab 09.01 WLAN Fundamentals

Overview

Now that the APs have been deployed in the network, your customer would like to understand what options Aruba Central offers with regards to RF monitoring and radio configuration. In this lab you will explore the configuration options using radio profiles to control the channels and power settings. You will review the centralized channel and power planning offered by AirMatch and how you can see the applied settings from the AP and in Aruba Central. The last section will show how to deploy a WLAN to a limited set of APs in a group.

Objectives

- List current radio profile use.
- Change radio profile
- Review channel and power changes.
- List AirMatch global schedule.
- Restricting a WLAN to a subset of APs in the group.

Task 1: Review Radio Default Channel and Power

In this task you will create an example WLAN to review the default AP radio channel and power settings.
You explore AP BSSID information and see how radio channel utilization and power settings can be visualized in Aruba Central.

Objectives

- Review the default radio profile.
- Explore AP BSSID information.

Steps

Review the Radio Status at the Group Level

First review the current radio settings at the group level.

1. In Aruba Central, navigate to Context: **Groups / campus-wifi-ui** > Navigation: **Devices** > Top: **Access Points** > Right-top: **List** > **Access Points**

Question: Are there two access points connected to Central? How can you see that an AP is connected to Central?

Answer: Yes, both APs are connected. They have status '**Online'** with a green dot next in front of it.

2. Click the **Radios** page.

Question: What is the status for your radios?

Answer All radios are currently offline.

Review the Radio Status at the Device Level

You can also review the radio settings at the individual AP device level.

3. Click **ap1** in the Access point List. Central will take you to the ap1 device-level 'Overview' page. You are now at this location in Central:

Context: **Device / ap1** > Navigation: **Overview** > Top: **Summary**

Question: After the **Device** and **Network** applet, the **Radios** applet is displayed. What is the status for the radios 2.4 and 5 Ghz?

Answer: Both radios are DISABLED now.

Question: Why would the radios currently be disabled?

Answer: There are no WLAN SSIDs configured for the APs, the radios remain disabled.

Configure WLAN

You will now configure a very basic, open SSID to enable the radios. Later you will apply more configuration to this WLAN.

4. In Aruba Central, navigate to Context: **Groups / campus-wifi-ui** > Navigation: **Devices**> Top: **Access Points**> Right top: **Config**
5. Under the **WLANs** tab, click **Add SSID** to start the New Network Wizard.
6. On the General tab, configure name **p#tx-guest**

> **NOTE:** Make sure to replace
> # with your Pod number
> x with your Table number (01-14)
>
> For example, if you are using Table 07 in Pod 28, your WLAN name will be
>
> **p28t07-guest**
>
> This indicates Pod 28, Table 07.
> Please check with your instructor if you are not sure about the Pod and Table number. This procedure ensures that you will be able to recognize your own SSID in the WLAN list, even when multiple pods (classrooms) are active at the same time.

7. Click **Next**.
8. On the **VLANs** page, leave the traffic forwarding mode to bridge mode (default).

> **NOTE:**
> The use of tunnel and mixed forwarding requires an Aruba gateway in the deployment. In this Associate class only the bridge mode option is used. In the Professional level Aruba Campus Access course you learn about Aruba gateways.

9. For the Client VLAN Assignment, select Static.
10. Set **VLAN ID** to **15**.
11. **Remove** VLAN 1.

> **IMPORTANT:**
> Make sure to remove the default VLAN 1 from the VLAN list. If VLAN 1 were listed together with VLAN 15, the wireless clients would be distributed over both VLANs.

12. Click **Next**.
13. On the Security tab, move the slider to **open**. Leave other values to default.

> **NOTE:**
> By default, the Enhanced Open option is enabled. This provides automatic encryption on an open WLAN for clients that support Enhanced Open. Clients that do not support this feature will still connect without wireless encryption to this WLAN. This feature uses an additional hidden SSID for the encrypted traffic. The feature is known as Opportunistic Wireless Encryption (OWE), which is used in the extra, hidden SSID.

14. Click **Next**.
15. On the Access tab, leave the default to **unrestricted.**

> **NOTE:** The Access page controls the access control to the network, such as the firewall rules. The unrestricted setting translates to a default rule 'Allow any traffic to all destinations'.

16. Click **Next**.
17. To complete the configuration, click **Finish**.
18. Wait for the deployment wizard to complete and confirm by clicking **OK**.

Verify Configuration Push

After making a configuration change, you can verify the that the updated configuration is pushed (synchronized) to the devices in the group.

19. Click on the **Configuration Audit** tab.

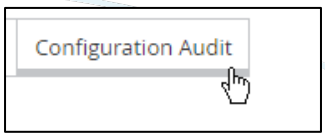

20. The **Configuration Status** tile should show: **Not in Sync 0 Device**. This indicates that all APs have synced the configuration successfully.

> **NOTE:** It may take a minute for the changes to be queued.
> You may refresh the page a few times to see the Configuration Status change.

In Central verify the status of the new WLAN

Now that the new WLAN has been created and pushed to the APs, you can use Aruba Central to verify the operational status of the WLAN.

21. In Aruba Central, navigate to Context: **Groups / campus-wifi-ui** > Navigation: **Overview** > Top: **Summary**

22. Scroll down to the WLANs table.

Question: Do you see the new WLAN in the list?

Answer: Yes, the configured guest WLAN is listed.

> **NOTE:** It may take a few minutes for the WLAN to show up in the list. If any SetMeUp WLANs are listed, they will eventually disappear.

Review the base Radio MAC

Each AP-advertised WLAN is advertised with a user-friendly name - the Extended Service Set Identifier (ESSID), and a unique MAC address for that specific AP radio - the Basic Service Set (BSS). The endpoint should know the name (ESSID) to which it should connect. Then the endpoint's WLAN NIC queries the APs. Each AP responds with a list of the actual MAC addresses (BSS table) to which the endpoint can connect. Each AP radio has a MAC address range. These addresses are used to advertise configured WLANs.

In these steps you review the base MAC address of your lab environment radios. Each radio (2.4 and 5 GHz in this lab) has its own base MAC address.

23. In Aruba Central, navigate to Context: **Groups / campus-wifi-ui** > Navigation: **Devices**> Top: **Access Points** > Top-Right: **List**
24. Click on the **Radios** page.

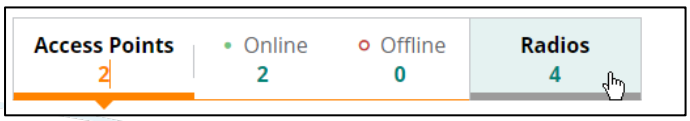

25. A list of AP radios for the APs in this group is displayed.

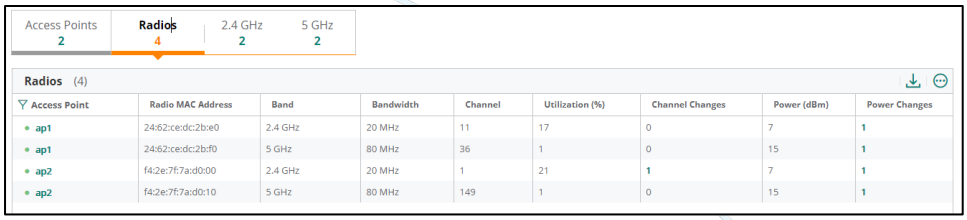

Question: How many radios are listed?

Answer: There are 4 radios, 2 for each AP.

Take note of your ap1 2.4GHz radio MAC address:

_____

Take note of your ap1 5GHz radio MAC address:

_____

TIP: You can also just download a CSV file with the ap names, radio base MAC and band information using the download icon at the right-top:

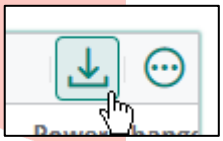

Example of the CSV output:

	A	B	C	D	E	F	G	H	I	J	K
1	ACCESS POINT	AP STATUS	RADIO MAC ADDRESS	BAND	BANDWIDTH	CHANNEL	UTILIZATION (%)	CHANNEL CHANGES	POWER (dBm)	POWER CHANGES	NOISE FLOOR (dBm)
2	ap1	Up	24:62:ce:dc:2e:10	5 GHz	80 MHz	52	0	0	18	0	-95
3	ap1	Up	24:62:ce:dc:2e:00	2.4 GHz	20 MHz	11	21	0	9	0	-94
4	ap2	Up	24:62:ce:dc:2e:d0	5 GHz	80 MHz	36	4	0	18	0	-92
5	ap2	Up	24:62:ce:dc:2e:c0	2.4 GHz	20 MHz	11	15	0	9	0	-93
6											

Question: For each AP, what is the relationship between the radio base MAC address of the 2.4 and 5 GHz?

Answer: The second, least-significant number is incremented with 1. Since MAC addresses are hexadecimal numbers, this is an increment of 16.

Example radio MAC list:

Access Point	Radio MAC Address	Band
ap1	cc:d0:83:ce:0a:00	2.4 GHz
ap1	cc:d0:83:ce:0a:10	5 GHz
ap2	cc:d0:83:d4:5e:40	2.4 GHz
ap2	cc:d0:83:d4:5e:50	5 GHz

Question: What is the significance of this number 16?

Answer: The base radio MAC Address is the start of a range of MAC addresses. Each SSID (WLAN Name) that this radio shall advertise gets the next available MAC address assigned as the BSS MAC.

Review the BSSID MAC used by the WLAN

You recently saw the base radio MAC address. Next you see how these MAC addresses are allocated on the test WLAN that you created.

26. In the list of radios, click on **Access Point ap1**. You see ap1 device details screen. You are now at this location: Context: **device / ap1** > Navigation: **Overview** > Top: **Summary**

27. Scroll down to see the WLANs table at the bottom of this screen.

Question: You have only configured one WLAN. Why do you see two ESS names?

Answer: The p#tx-guest is the open WLAN. Since you have Enhanced Open enabled, an additional, hidden SSID was created to support Opportunistic Wireless Encryption (OWE). These WLANs start with *_owe*. They are effectively using the same settings as the open WLAN, but they have encryption enabled without need for a PSK.

28. Expand the p#tx-guest WLAN to see the BSSID MAC that is used by this AP (ap1).

Example screenshot

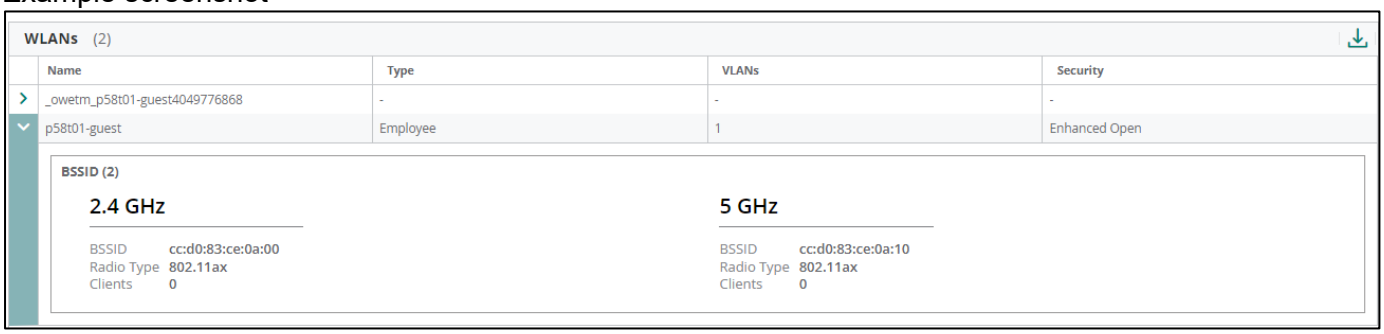

Question: Do you see a relation with the base Radio MAC?

Answer: Yes, since this is the first SSID that is advertised by this radio, the first address from the radio MAC range is used.

29. Expand the WLAN that starts with _owe.

Example screenshot

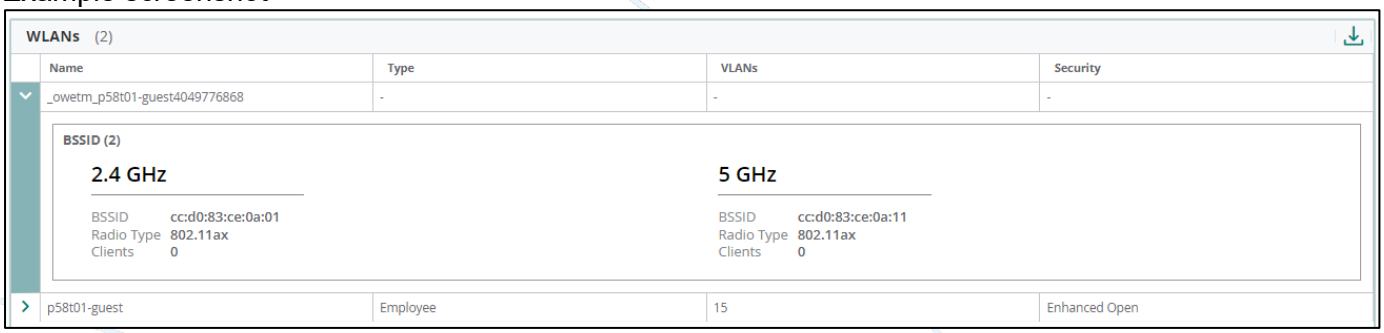

Question: What is the difference between this BSSID and the previous WLAN BSSID?

Answer: Since this is the second SSID that is advertised by the radio, the next BSSID MAC in the range is used.

Connect with Client

Next you use PC1 to connect to your WLAN. Since PC1 has both a wired and wireless NIC, you first check network connections. The wired NIC should be disabled and the wireless NIC must be enabled.

30. Open a desktop session to PC1.
31. Right-click the **Network icon** in the notification bar.
32. Click **Open Network & Internet Settings**.

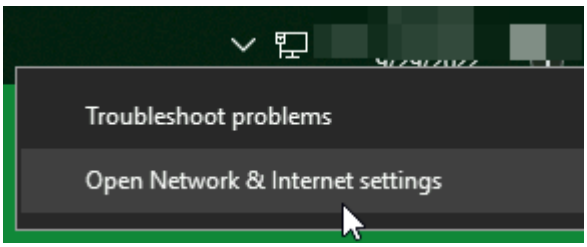

33. Click **Change adapter options**.

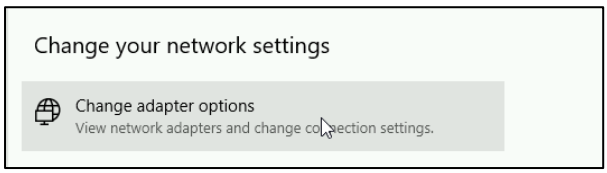

34. In the Network Connections, **disable** the Lab NIC (this is the wired NIC connected to sw-edge1 port 1/1/1). (right-click on the NIC > disable)

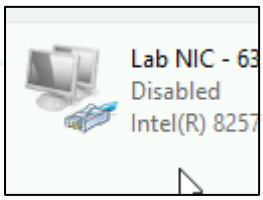

35. Make sure the Wireless NIC is enabled.

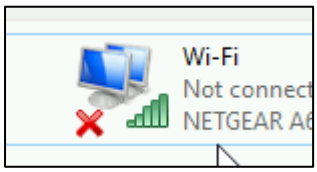

36. Click the **Network** icon in the notification bar to see the current list of WLANS.

37. In the WLAN list, look for your WLAN (P#Tx-guest) and connect to it.

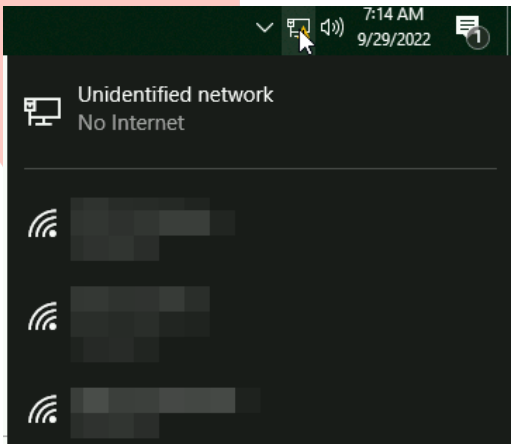

38. Open a command prompt (cmd) and review your IP address using ipconfig.

```
ipconfig
```

```
...
Wireless LAN adapter Wi-Fi Lab:

   Connection-specific DNS Suffix  . : aruba-training.com
   Link-local IPv6 Address . . . . . : fe80::753c:42f5:cdf5:6f7a%19
   IPv4 Address. . . . . . . . . . . : 10.1.15.50
   Subnet Mask . . . . . . . . . . . : 255.255.255.0
   Default Gateway . . . . . . . . . : 10.1.15.1
...
```

Question: What IP address is assigned to the wireless NIC?

Answer: The wireless NIC should have received an IP address in the 10.1.15.0/24 range, for example 10.1.15.50.

39. To generate some traffic, start a continuous ping to 10.254.1.21 (your lab DHCP/DNS server). The ping should be successful.

```
ping 10.254.1.21 -t
```

Review the WLAN on the AP

You can review these BSS entries using the **show ap bss-table** command on the AP. Using Aruba Central, it is easy to send common CLI diagnostics commands to one or more devices.

40. In Aruba Central, navigate to Context: **Groups / campus-wifi-ui** > Navigation: **Tools** > Top: **Commands**
41. In the **Device Type** dropdown select **Access Point**.
42. Click **Available Devices** and select both **ap1** and **ap2**.
43. In the **Categories** list, select **All Category**. This will show all available commands.
44. In the **Commands** list, click the **filter** icon and type '**bss**'.

45. Select the **show ap bss-table** command, then click **Add**.

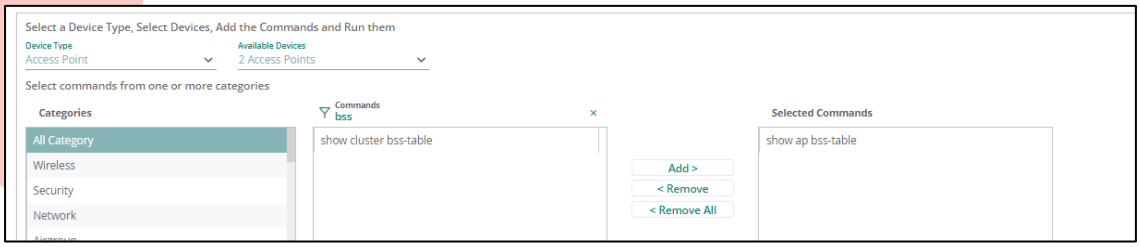

46. Click **Run** to initiate the request. You will now see rotating dots while Central collects command output on both ap1 and ap2

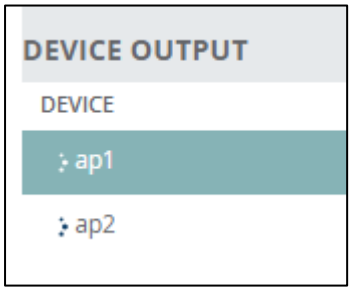

47. After a few moments, the output will be collected. You can detect this based on the confirmation icon next to the device names.

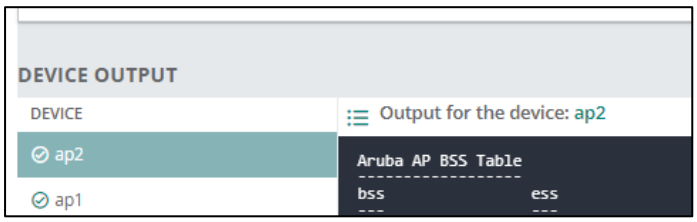

48. Under **Device Output**, click the **ap1** entry to see the results for ap1.

TIP: You can maximize the command output using the maximize button:

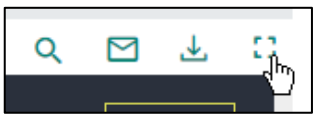

To resize, use the close button:

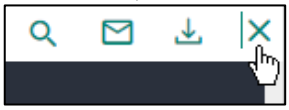

.

Question: How many entries do you see in the BSS list?

Answer: Four.

Question: What do these BSS MAC addresses represent?

Answer: These are the MAC address for this AP for this ESSID, reported as ESS in the output. This is the same BSS information you have seen in the AP device details screen.

Question: What **"phy"** is reported?

Answer: This depends on the AP models used in your lab setup. For example, **a**-HE or **a**-VHT, which represents the 5 GHz band, and **g**-HE or **g**-HT, which represents the 2.4 GHz band.

> **NOTE:** The **a** and **g** values are based on the bands used by the original wireless standards (802.11a and 802.11g), but it does not represent the operational WLAN standard.

Question: What do HE, HT or VHT represent?

Answer: This is the operational WLAN standard in use by the radio. This table shows the WLAN standards and how they are displayed on the AP **phy**.

Standard	Phy
802.11n	HT (High Throughput)
802.11ac	VHT (Very High Throughput)
802.11ax	HE (High Efficiency)

This allows you to determine the standard used by the radio.

Question: What is **ch/eirp** column?

Answer: This represents the Channel and the Effective Isotropic Radiated Power (EIRP), which represents the AP's effective transmit power.

Take note of the channel and power for the 2.4GHz:

_____

Take note of the channel and power for the 5GHz:

_____

49. Check the output for ap2.

Take note of the channel and power for the 2.4GHz:

_____

Take note of the channel and power for the 5GHz radio:

_____

Aruba Central Radio information

Now you will review the same radio information from the Central UI. While the CLI works fine for a single AP, Central provides an easy overview and dashboard for a group of AP radios.

50. In Aruba Central, navigate to Context: **Groups / campus-wifi-ui** > Navigation: **Devices** > Top: **Access Points** > Right-top: **List**
51. The top shows Access Points and Radios. Click the **Radios** link.

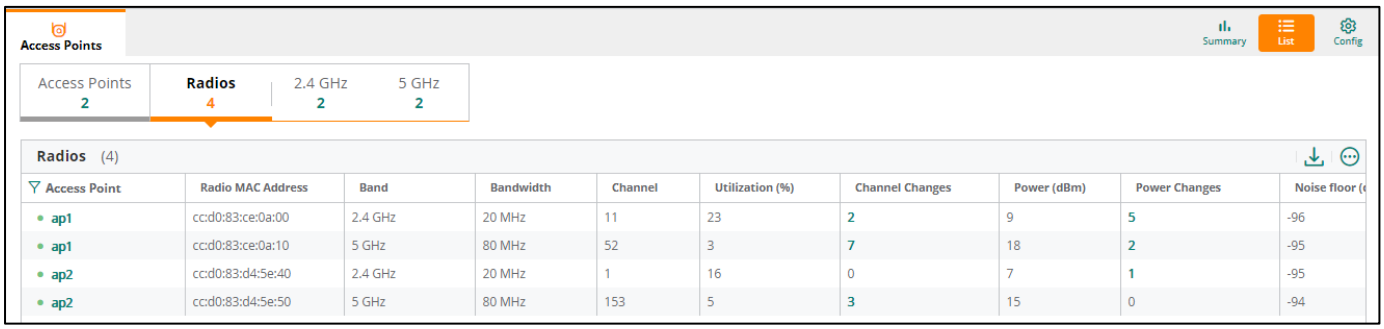

Access Point	Radio MAC Address	Band	Bandwidth	Channel	Utilization (%)	Channel Changes	Power (dBm)	Power Changes	Noise floor
ap1	cc:d0:83:ce:0a:00	2.4 GHz	20 MHz	11	23	2	9	5	-96
ap1	cc:d0:83:ce:0a:10	5 GHz	80 MHz	52	3	7	18	2	-95
ap2	cc:d0:83:d4:5e:40	2.4 GHz	20 MHz	1	16	0	7	1	-95
ap2	cc:d0:83:d4:5e:50	5 GHz	80 MHz	153	5	3	15	0	-94

Question: Compare the channel and power to the command output you previously captured. Do you see the same information?

Answer: Yes, the AP reports the actual channel, power, and utilization to Aruba Central, which is known as telemetry data.

Visualize Channel Distibution

Aruba Central also provides a convenient dashboard overview with the channels and bands that are in use by the selected APs.

52. At the right-top, click **Summary**.
53. At the top, click **Radios**.

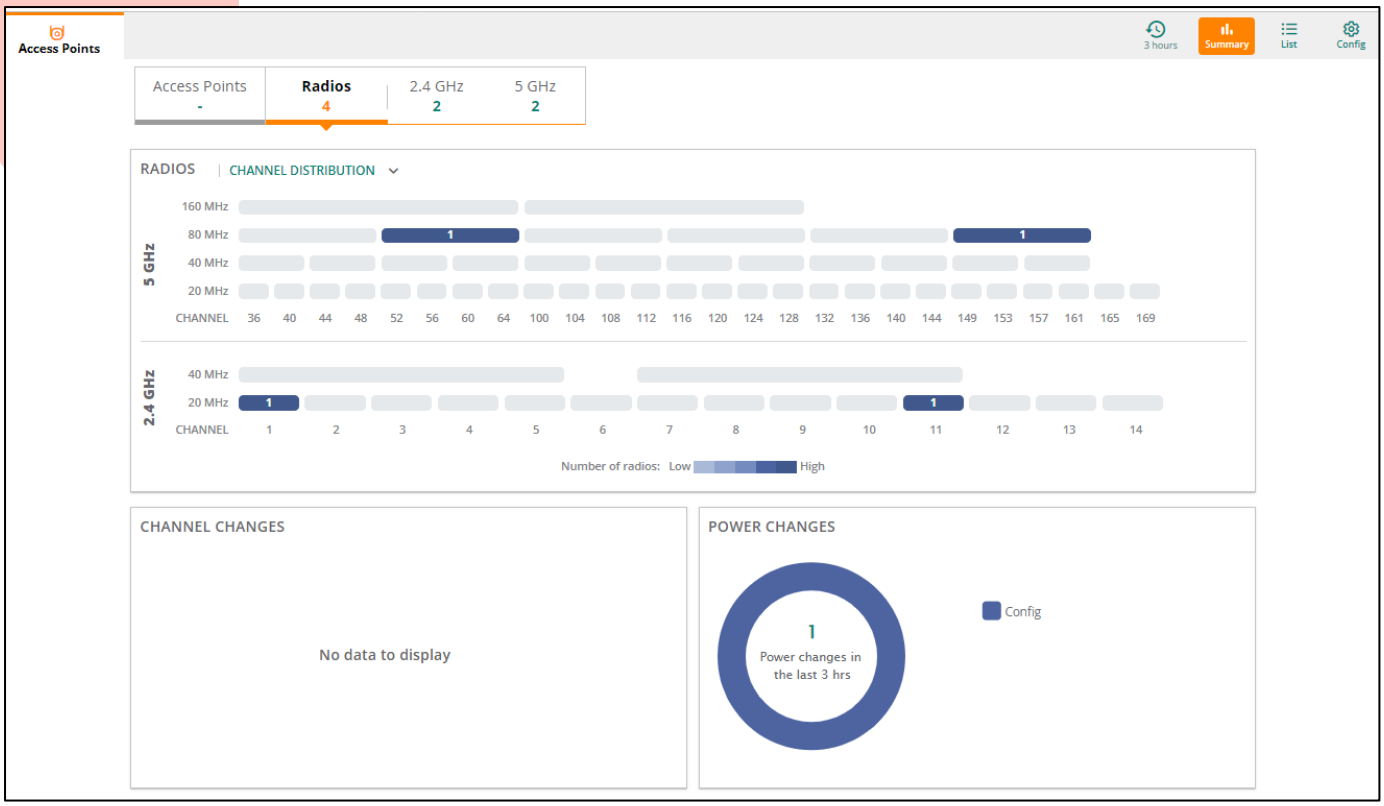

This provides a graphical representation of the number of radios that are using a channel, as well as the configured channel width.

Visualize Power Distribution

Like the channel distribution, Aruba Central can also provide a dashboard for the radio power settings.

54. At the top, click **Channel Distribution** and select **Power Distribution** from the list.

Example screenshot, your output may be different.

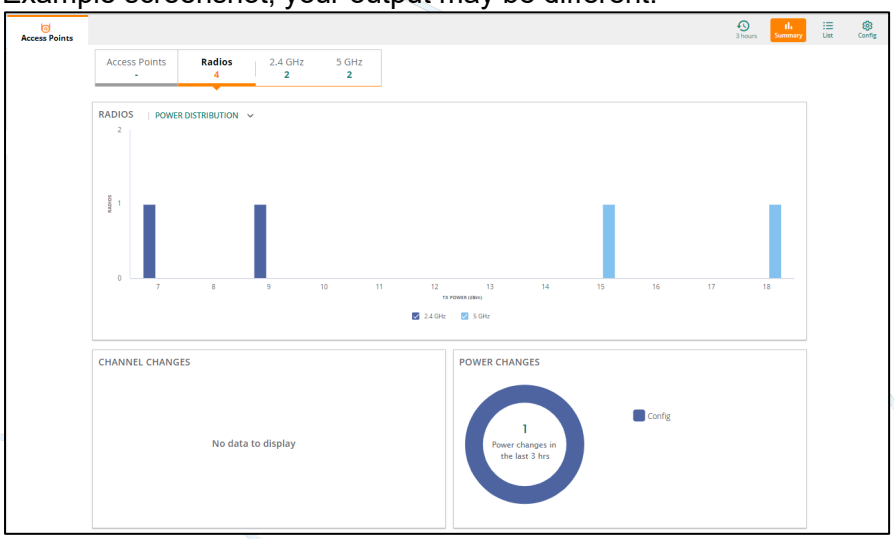

Now you have observed the default channel and power. In the next section you will use a radio profile to control the channel and power calculation.

Task 2: Configuring Radio Profiles

Your customer would like to know if they can control the use of some channels and the transmit power of the radios. After contact with your senior colleague, you learn that you can control these settings using a radio profile.

Within an Aruba Central configuration group, multiple AP radio profiles could exist, to handle some APs in special areas that may have special radio tweaking requirements. However, at the AP level, only one radio profile can be applied. This must be one of the radio profiles you have created at the group level. In the next steps you will first create a new radio profile, and then apply it to your APs.

Objectives

- Create AP radio profiles.
- Apply AP radio profiles.
- Verify the AP radio profile.

Steps

Create Radio Profile

In the next steps you will create a new radio profile to control the channel use of the AP radios.

1. In Aruba Central, navigate to Context: **Groups / campus-wifi-ui** > Navigation: **Devices** > Top: **Access Points** > Right top: **Config**
2. Click the **Radios** page.
3. Click Add Profile.
4. Name : campus-table

> **NOTE:** Do not adjust the Transmit Power settings at this point. They will be adjusted in the next task.

5. Click the allowed channel list for the 5 Ghz radio. (Click on the list with the channel numbers).

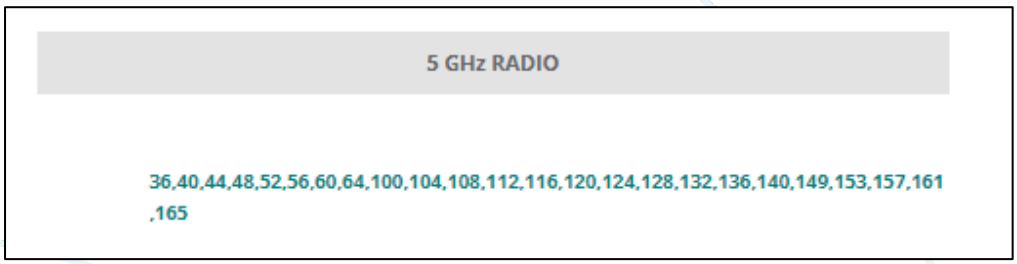

This will open the **Allowed Channels** window.

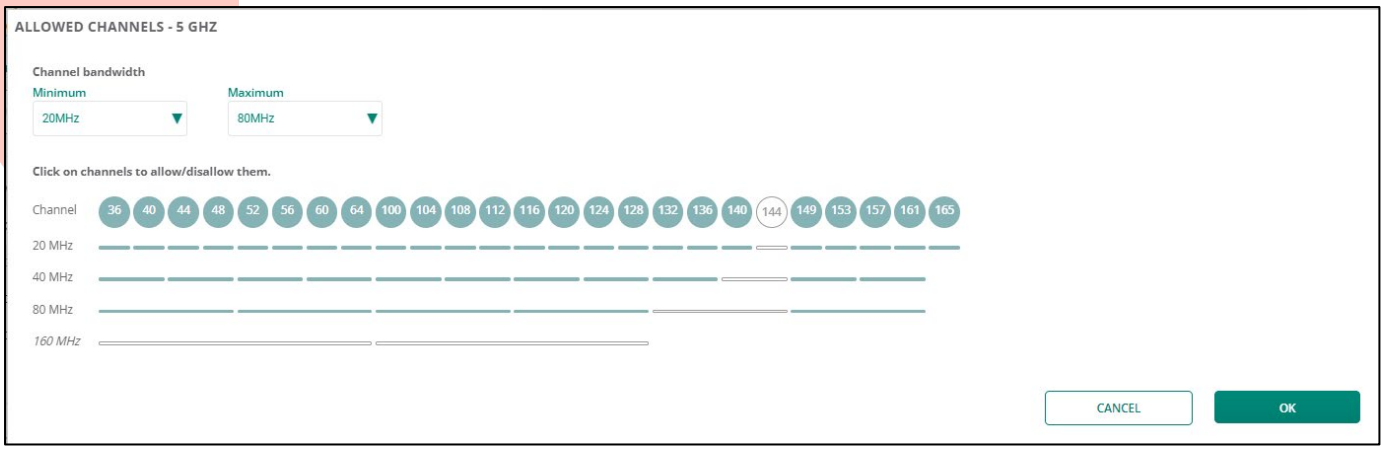

6. Set the 5 Ghz **Maximum** to **20Mhz**. Leave the Minimum to the default 20Mhz.

> **NOTE:** In the remote lab environment, there are MANY APs. Reducing the channel width helps to reduce the interference with other lab environments.

7. For testing purposes, adjust the **Allow channel** list so your profile only contains two 20Mhz channels. You can click a channel number to enable or disable it. Disable all channels, except the two channels from the table below for your assigned table.

Table	Channels
1 / 5 / 9 / 13	36-40
2 / 6 / 10 / 14	44-48
3 / 7 / 11	52-56
4 / 8 / 12	60-64

Example for **table1**. Use the table above to configure the channels for your own table number:

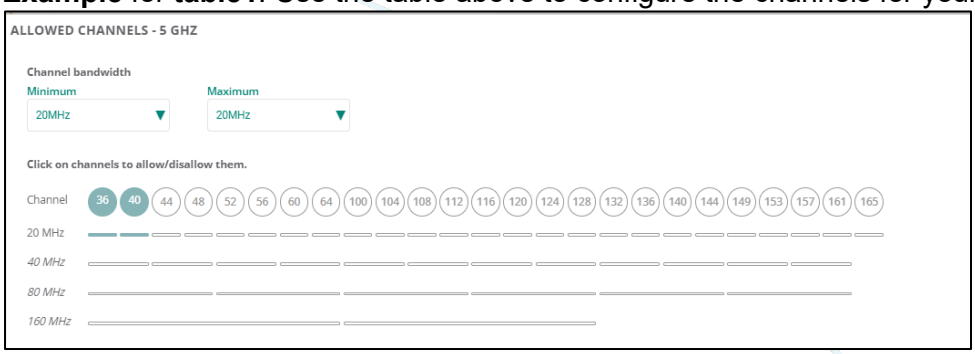

8. Click **OK** to submit the allowed channels.
9. Click **Show advanced Settings**.

Question: What does the **Very High Throughput** option represent?

Answer: VHT represents the radio phy operation, it represents the activation of 802.1ac.

Question: What does the option **Set Second Radio Differently** indicate?

Answer: Some Aruba APs can have 2 5GHz radios. With this option it is possible to apply different radio properties to the second 5GHz radio. By default, the second 5GHz radio will use the same settings as the primary 5GHz radio.

> **NOTE:** This setting only applies to APs with a second 5 GHz radio, it will be ignored by other APs.

10. Click **Save** to commit the new radio profile.
11. Verify that the new radio profile is shown in the list.

Apply the Radio Profile

Now you have created a new radio profile, but it is not yet in use. In the next steps you will activate the new radio profile on your APs. Only one radio profile can be active per radio, this defaults to the **default** radio profile.

12. Click the **Access Points** page.
13. Click *both* **APs**. Both APs should be selected now.
14. At the bottom, click the **pencil** icon to edit the Access Points.
15. In the **Radio Profile** column, select the **campus-table** profile.
16. Enable the **Apply to all access points** checkbox, then click **Apply**.

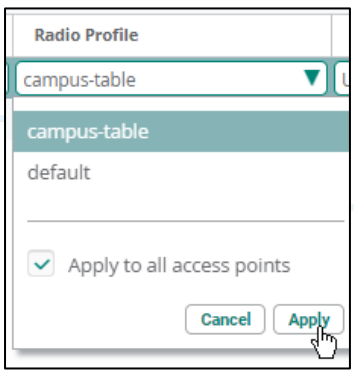

17. Click the **WLANs** column.
18. Enable the **Apply to all access points** checkbox, then click **Apply**.

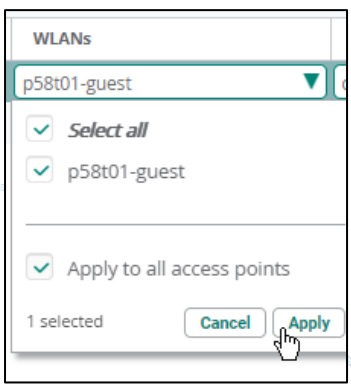

19. Click Save Settings.

Verify the Generated Configuration

After performing a configuration change in Aruba Central, you can review and analyze the generated commands. In the next steps you explore the radio profile configuration commands.

20. In Aruba Central, navigate to Context: **Groups / campus-wifi-ui** > Navigation: **Audit Trail**
21. In the list, look for an entry with Description *Add/Update RF Radio Profile.*
22. Click the three dots at the end of this record to see the details.

23. Review the commands that were generated for the radio profile.

Example output (your table may use different channels).

```
rf dot11g-radio-profile campus-table
 disable-arm-wids-functions off
 exit

rf dot11a-radio-profile campus-table
 disable-arm-wids-functions off
 allowed-channels 44,48
 ch-bw-range 20MHz 20MHz
 exit

rf dot11a-secondary-radio-profile campus-table
 disable-arm-wids-functions off
 allowed-channels 44,48
 ch-bw-range 20MHz 20MHz
 exit
```

Question: What do the dot11g and dot11a represent?

Answer: dot11g represents the 2.4 GHz radio profile, dot11a represents the 5 GHz radio profile.

Question: What is the dot11a-secondary radio profile?

Answer: Some Aruba APs support two 5 GHz radios. This profile would be applied to the second 5 GHz radio. By default, it is the same as the primary 5 GHz radio. Different settings could have been applied in the radio profile with the advanced option *Set Second Radio Differently*. (Not applicable to this lab).

Question: What is controlled by the **ch-bw-range** command?

Answer: This command controls the channel bandwidth. You have configured the maximum channel width to 20Mhz in the radio profile.

Question: Which channels do you see for the allowed-channels command?

Answer: These are the channels you have selected in the radio profile.

Question: Why is the allowed-channels command missing for the dot11g radio profile?

Answer: You only configured the 5 GHz radio profile. No changes were made for the 2.4 GHz radio profile.

24. **Close** the audit entry details window.

Verify the configuration received by the AP

In the previous steps you have seen how Central generates the configuration for the AP based on your UI configuration steps. These commands are then sent to the AP, applied to the running configuration, and saved locally on the AP. Locally on the AP you can also monitor received configuration commands from the cloud (Aruba Central). In the next steps you will open a console connection to the AP and review the received commands.

25. In Aruba Central, navigate to Context: **Groups / campus-wifi-ui** > Navigation: **Tools** > Top: **Commands**
26. In the **Device Type** dropdown select **Access Point**.
27. Click **Available Devices** and select **ap1**.
28. In the **Categories** list, select **All Category** to see all available commands.
29. In the **Commands** list, click the **filter** icon and type '**cloud'**.
30. Select the show ap debug cloud-config-received command, then click Add.
31. Click **Run** to start the request.
32. Wait a few moments until the Output for the device is updated and review the commands.

TIP: Remember you can maximize the output screen:

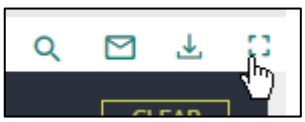

33. Starting from the end, look back until you see a timestamp message. This was the last time configuration changes were received.

Example output:

```
timestamp: 2022-09-06 05:17:12
rf dot11g-radio-profile campus-table: OK
disable-arm-wids-functions off: OK
```

```
exit: OK
rf dot11a-radio-profile campus-table: OK
disable-arm-wids-functions off: OK
allowed-channels 44,48: OK
ch-bw-range 20MHz 20MHz: OK
exit: OK
rf dot11a-secondary-radio-profile campus-table: OK
disable-arm-wids-functions off: OK
allowed-channels 44,48: OK
ch-bw-range 20MHz 20MHz: OK
exit: OK
```

Question: Review the commands. Do you recognize them compared to the audit entry?

Answer: The commands on the AP should be the same as the commands that were added to the audit entry.

Question: Were the commands successfully processed by the AP?

Answer: For each line that is successfully processed by the AP, the AP will add the **OK** keyword at the end.

TIP: If you maximized the screen, use the X icon to revert to the standard view.

You have now verified that the configuration was pushed successfully.

Verify the Applied settings on the AP

In the previous steps you have verified that the commands have been pushed to the AP.
In the next steps you will verify that the radio profile settings are operational on the AP.

You will do this by reviewing:
- the active radio profile in use by the AP radios.
- the current radio channel in use by the AP 5 GHz radio.

34. In Aruba Central, you should still be at: Context: **Groups / campus-wifi-ui** > Navigation: **Tools** > Top: **Commands** with ap1 as selected device.
35. Click **Remove All** to clear the Selected Commands window.
36. Set the Commands filter to **radio**.
37. In the Commands list, select **show radio config** and click **Add**.
38. Set the **Commands** filter to **bss**.
39. In the Commands list, select **show ap bss-table** and click **Add**.

TIP: You can clear the previous command output by using the **CLEAR** button. This ensures you will only see the output of the currently requested command.

40. Click on **Run**.

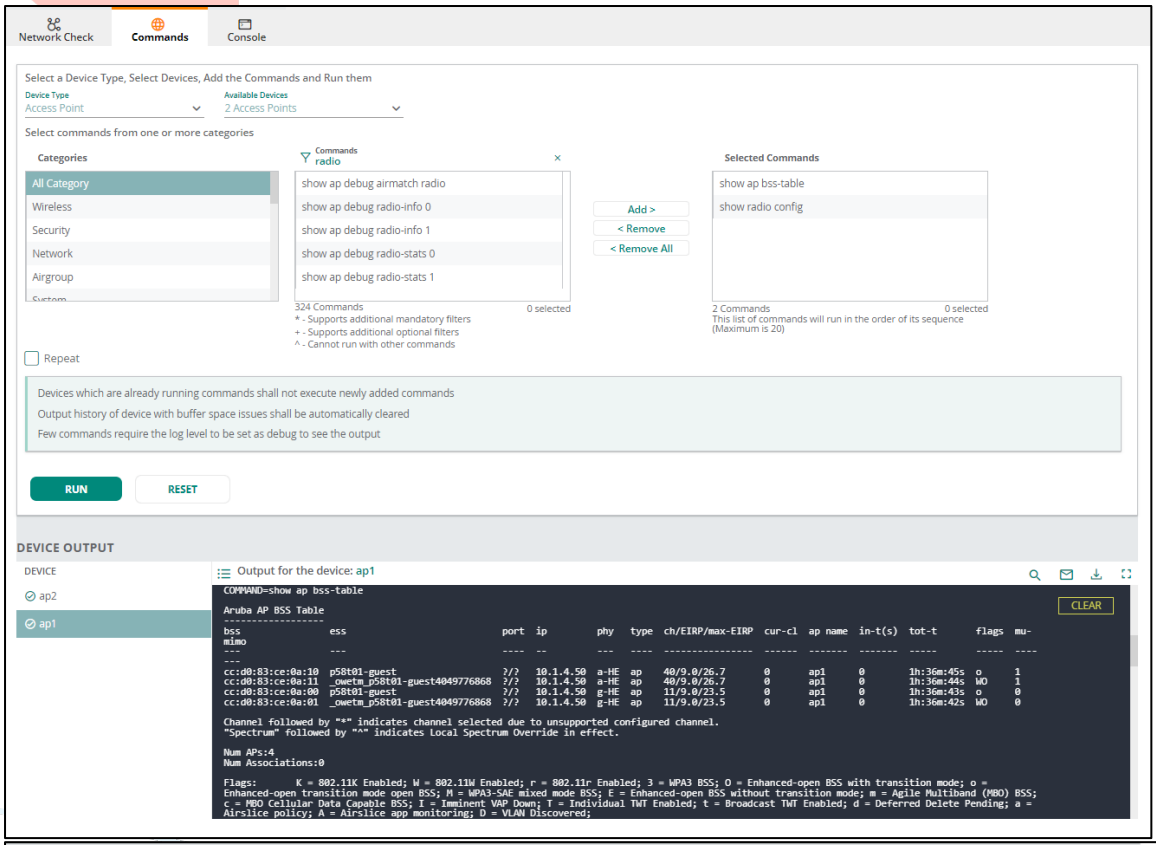

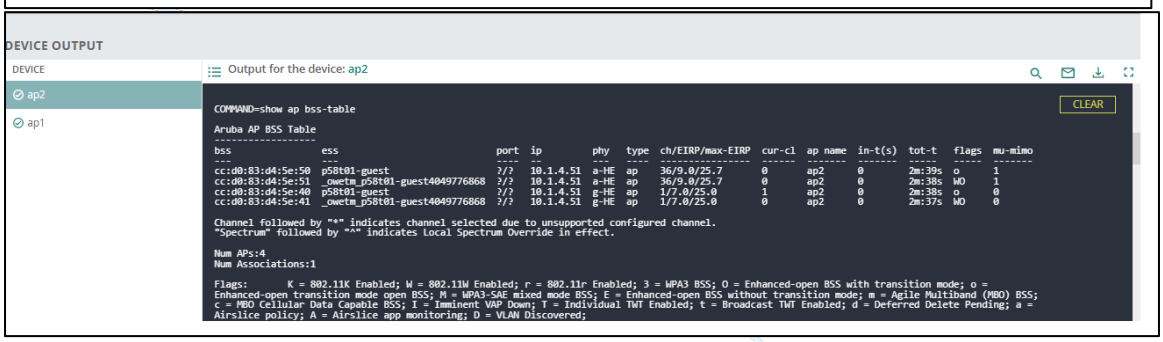

41. Review the deployed radio profile.

Example output (your channels may be different):

```
COMMAND=show radio config

------------------------------------------------------------
2.4 GHz:
Profile Name:campus-table
Zone Name:
Legacy Mode:disable
Single Chain Legacy:disable
Beacon Interval:100
```

```
802.11b-protection:enable
802.11d/802.11h:disable
Interference Immunity Level:2
Channel Switch Announcement Count:0
MAX Distance:0
Channel Reuse Type:disable
Channel Reuse Threshold:0
Background Spectrum Monitor:disable
Max Tx Power:12
Min Tx Power:6
Cell Size Reduction:0
Smart Antenna:disable
WIDS Override:off
40MHZ intolerance:disable
Honor 40MHZ intolerance:enable
Channel quality threshold:70
Channel quality wait time:120
High noise backoff time:720
Max channel bandwidth:20MHz
Min channel bandwidth:20MHz
Allowed channels:1,6,11
BSS Color:0
BSS Color Switch Countdown:10
Free Channel Index:25
Scanning-disable:No
Aggressive-scan:enable
Active-scan:disable
IoT Coexistence :enable
------------------------------------------------------------
5.0 GHz:
Profile Name:campus-table
Zone Name:
Very High Throughput:enable
Legacy Mode:disable
Single Chain Legacy:disable
Beacon Interval:100
802.11d/802.11h:disable
Interference Immunity Level:2
Channel Switch Announcement Count:0
MAX Distance:0
Channel Reuse Type:disable
Channel Reuse Threshold:0
Background Spectrum Monitor:disable
Max Tx Power:21
Min Tx Power:15
Cell Size Reduction:0
Smart Antenna:disable
WIDS Override:off
Channel quality threshold:70
Channel quality wait time:120
High noise backoff time:720
Max channel bandwidth:20MHz
Min channel bandwidth:20MHz
Radar backoff time:720
Allowed channels:44,48
BSS Color:0
```

```
BSS Color Switch Countdown:10
Free Channel Index:25
Scanning-disable:No
Aggressive-scan:enable
Active-scan:disable
Zero Wait DFS:disable
------------------------------------------------------------
5.0 GHz:
Profile Name:campus-table
Zone Name:
Very High Throughput:enable
Legacy Mode:disable
Single Chain Legacy:disable
Beacon Interval:100
802.11d/802.11h:disable
Interference Immunity Level:2
Channel Switch Announcement Count:0
MAX Distance:0
Channel Reuse Type:disable
Channel Reuse Threshold:0
Background Spectrum Monitor:disable
Max Tx Power:21
Min Tx Power:15
Cell Size Reduction:0
Smart Antenna:disable
WIDS Override:off
Channel quality threshold:70
Channel quality wait time:120
High noise backoff time:720
Max channel bandwidth:20MHz
Min channel bandwidth:20MHz
Radar backoff time:720
Allowed channels:44,48
BSS Color:0
BSS Color Switch Countdown:10
Free Channel Index:25
Scanning-disable:No
Aggressive-scan:enable
Active-scan:disable
Zero Wait DFS:disable
```

Question: What is the Max channel bandwidth for the 5GHz radio?

Answer: 20 MHz.

Question: What are the allowed channels?

Answer: This depends on your table number, but it should reflect your two assigned channels (for this lab setup).

42. Review the **BSS table** command output.

Question: What is the current active channel of the 5 GHz radio (the ‘phy’ with a-HE or a-VHT)?

Answer: The active channel should be one of the two channels you have allowed in the radio profile.

43. Click **Remove All** to clear the commands.
44. In the Command filter, enter '**arm**'.

Question: What does ARM represent?

Answer: Adaptive Radio Management (ARM) is part of the process that analyzes channel and power use, reports information to Aruba Central, and applies the changes as requested by Aruba Central.

45. Add the command **show ap arm history** and click **Run**.

Example output:

```
COMMAND=show ap arm history

Interface :wifi0
ARM History
-----------
Time of Change       Old Channel  New Channel  Old Power  New Power  Reason
Result
--------------       -----------  -----------  ---------  ---------  ------      ----
--
2022-09-08 04:03:31  44           40           18.0       18.0       AM RD       -
2022-09-08 03:35:01  44           44           9.0        18.0       AM MinEIRP  -
2022-09-08 03:28:03  0            44           0.0        9.0        AM Init     -

Interface :wifi1
ARM History
-----------
Time of Change       Old Channel  New Channel  Old Power  New Power  Reason   Result
--------------       -----------  -----------  ---------  ---------  ------   ------
2022-09-08 03:28:38  2+           11           0.0        9.0        AM RD    -
2022-09-08 03:28:03  0            11           9.0        9.0        AM Init  -

I: Interference, R: Radar detection, N: Noise exceeded, Q: Bad Channel Quality E:
Error threshold exceeded, INV: Invalid Channel, G: Rogue AP Containment, M: Empty
Channel, P+: Increase Power, P-: Decrease Power, 40INT: 40MHZ intol detected on 2.4G,
NO40INT: 40MHz intol cleared on 2.4G, OFF(R): Turn off Radio due to Radar,
OFF(CONFIG): Turn off Radio due to Wrong Config, ON: Turn on Radio, D: Dynamic
Bandwidth Switch, I*: CCA Interference, C: Radar cleared, DM: Dynamic Mode Change, O:
Opmode change, AIRMATCH: AirMatch Event, AM Solver: AirMatch(AM) service selected
channel/power, AM Init: Initialized channel/power from flash, AM N: Noise exceeded,
AM NC: Noise Cleared, AM RD: Reg-Domain Profile Change, AM Rogue: Rogue AP
Containment, AM DRT: DRT File Change, AM MinEIRP: Min EIRP Change, AM MaxEIRP: Max
EIRP Change, AM Freeze: set static channel/power, AM Unfreeze: unset static
channel/powerNC: Noise Cleared, Random: Random Channel, RMC: Radio Mode Change, RCP:
Radio Client Preference Change
```

Question: What are the two interfaces shown in the output?

Answer: wifi0 and wifi1.Wifi0 represents the 5 GHz radio. The Channel columns can be used as a hint (channels 1-11 for 2.4 GHz).

Question: This list shows the latest channel and/or power changes for each radio. What code is listed in the Reason column for the last change of the **wifi0** interface (5GHz)?

Answer: The Reason is AM RD.

Question: What does this mean? Use the legend at the bottom to find the explanation.

Answer: AM RD indicates Regulatory Domain Profile Change. This is used when the allowed channel list has changed.

Verify the Applied settings in Central

Now you review the same radio information from the Central UI.

46. In Aruba Central, navigate to Context: **Groups / campus-wifi-ui** > Navigation: **Devices** > Top: **Access Points** > Right-top: **Summary**
47. The top shows Access Points and Radios. Click on the **Radios** link.

Example screenshot. Your table channels may be different from the screenshot.

Question: Do you see radios active on the channels you configured?

Answer: Yes, after the radio profile configuration, the AP radio moved to one of the configured channels.

48. Under the channel distribution, the channel changes are reported.

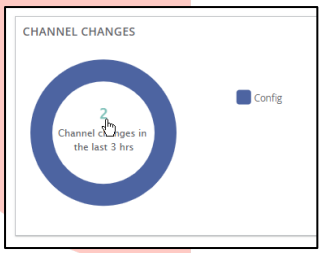

49. Click the number inside the circle of the “channel changes” pane to see the list of changes.

Example screenshot

Channel Changes (2)

Event Time	Reason	From Channel	To Channel	Band	Access Point
Sep 6, 2022, 11:17	Allowed Channels	36	48	5 GHz	ap2
Sep 6, 2022, 11:17	Allowed Channels	40	44	5 GHz	ap1

Question: For your ap1, what is the reason for the last channel change?

Answer: The ‘Allowed Channels’ was changed by the radio profile.

Question: What is the new channel for your ap1? Does it match the output from the ‘show arm history’ command?

Answer: Yes, the AP reports the new channel and ARM history to Aruba Central.

50. Close the Channel Changes window.

This concludes the radio profile channel configuration.

Changing Radio Power

In this section you will update the allowed transmit power in the radio profile.

51. In Aruba Central, navigate to Context: **Groups / campus-wifi-ui** > Navigation: **Devices** > Top: **Access Points**> Right top: **Config**
52. Click the **Radios** page.
53. On the **campus-table** record, click the **pencil** button to edit the profile.
54. Change the default power values to a lower value for this lab environment.

2.4GHz 6-9
5GHz 9-12

TIP: Using the mouse may be difficult to move the slider correctly. You can also click on the start or end marker and use the arrow keys (left and right) to change values.

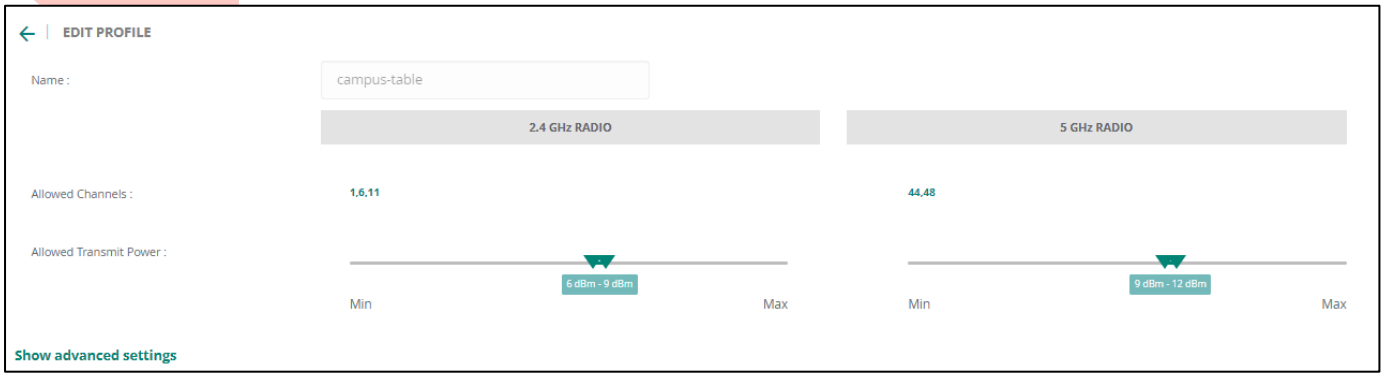

55. Click **Save** for the profile.
56. Use the **Audit Trail** to review the commands that affect the transmit power.

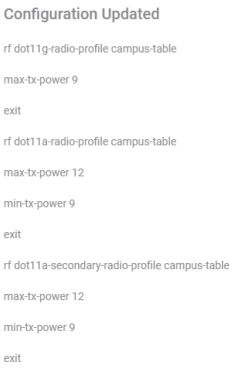

57. Navigate to **Tools > Commands** on **access point** "**ap1**" with the command **show radio config** to review the updated radio profile.

```
COMMAND=show radio config

------------------------------------------------------------
2.4 GHz:
Profile Name:campus-table
...
Max Tx Power:9
Min Tx Power:6
...
------------------------------------------------------------
5.0 GHz:
Profile Name:campus-table
...
Max Tx Power:12
Min Tx Power:9
```

```
...

=== Troubleshooting session completed ===
```

58. Use **Tools > Commands** on **access point ap1** with the command **show arm history** to review the radio change history. This should include a power change now for the wifi0 interface (5 GHz radio).

```
COMMAND=show ap arm history

Interface :wifi0
ARM History
-----------
Time of Change       Old Channel  New Channel  Old Power  New Power  Reason
Result
--------------       -----------  -----------  ---------  ---------  ------      ----
--
2022-09-08 04:26:12  40           40           18.0       12.0       AM MaxEIRP  -
2022-09-08 04:03:31  44           40           18.0       18.0       AM RD       -
2022-09-08 03:35:01  44           44           9.0        18.0       AM MinEIRP  -
2022-09-08 03:28:03  0            44           0.0        9.0        AM Init     -
...

I: Interference, R: Radar detection, N: Noise exceeded, Q: Bad Channel Quality E:
Error threshold exceeded, INV: Invalid Channel, G: Rogue AP Containment, M: Empty
Channel, P+: Increase Power, P-: Decrease Power, 40INT: 40MHZ intol detected on 2.4G,
NO40INT: 40MHz intol cleared on 2.4G, OFF(R): Turn off Radio due to Radar,
OFF(CONFIG): Turn off Radio due to Wrong Config, ON: Turn on Radio, D: Dynamic
Bandwidth Switch, I*: CCA Interference, C: Radar cleared, DM: Dynamic Mode Change, O:
Opmode change, AIRMATCH: AirMatch Event, AM Solver: AirMatch(AM) service selected
channel/power, AM Init: Initialized channel/power from flash, AM N: Noise exceeded,
AM NC: Noise Cleared, AM RD: Reg-Domain Profile Change, AM Rogue: Rogue AP
Containment, AM DRT: DRT File Change, AM MinEIRP: Min EIRP Change, AM MaxEIRP: Max
EIRP Change, AM Freeze: set static channel/power, AM Unfreeze: unset static
channel/powerNC: Noise Cleared, Random: Random Channel, RMC: Radio Mode Change, RCP:
Radio Client Preference Change
```

Question: What was the old and new power for the last change on the wifi0 interface?

Answer: The old power depends on your environment, but the new power should be in the range of 9-12dBm.

Question: What is the reason of the last change?

Answer: AM MaxEIRP: Max EIRP change.

Question: If your wifi1 (2.4GHz) interface does not report a new change, what could be the reason?

Answer: The wifi1 interface may already have been using a transmit power in the range of 6-9. Therefore, a power change was not required.

Review power change in Aruba Central

Now you will review the same radio power information from the Central UI.

59. In Aruba Central, navigate to Context: **Groups / campus-wifi-ui** > Navigation: **Devices**> Top: **Access Points** > Right-top: **Summary**
60. The top shows Access Points and Radios. Click the **Radios** link.
61. Change the selector Channel Distribution to **Power Distribution**.

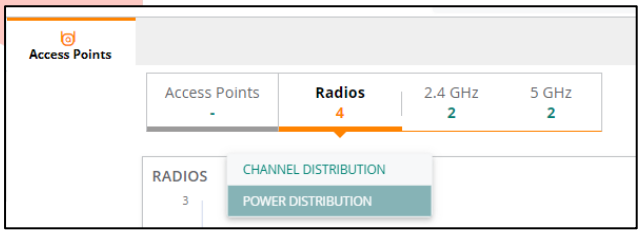

Example screenshot. Your table power distribution may be different from the screenshot.

62. Under the radio list, the number of power changes is shown. Click on the number in the circle to see the details.

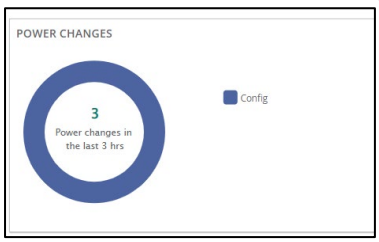

63. Review the details of the power changes.

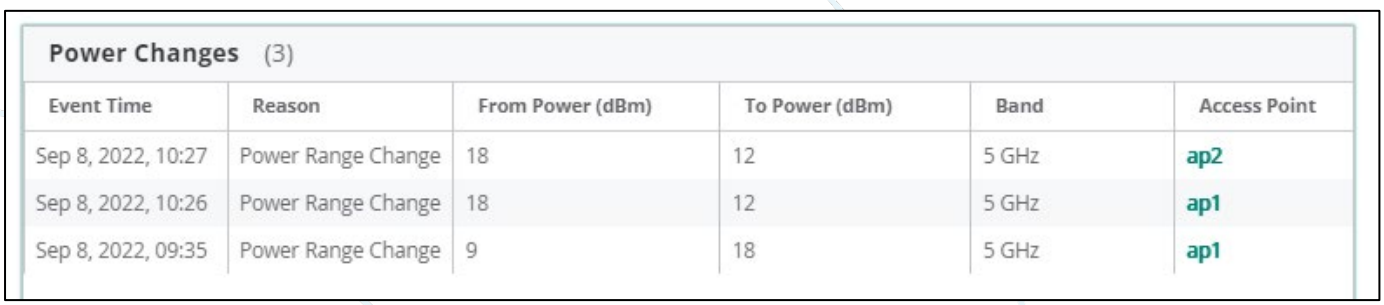

Power Changes (3)

Event Time	Reason	From Power (dBm)	To Power (dBm)	Band	Access Point
Sep 8, 2022, 10:27	Power Range Change	18	12	5 GHz	ap2
Sep 8, 2022, 10:26	Power Range Change	18	12	5 GHz	ap1
Sep 8, 2022, 09:35	Power Range Change	9	18	5 GHz	ap1

Question: What is the reason of the last power change on the 5GHz for ap1?

Answer: 'Power Range Change', since the radio profile applied a maximum of 12 for the 5GHz radio.

64. **Close** the power changes window.

> **NOTE:** Depending on the time of day that you do the lab, AirMatch (the cloud-based channel and power planner) may have also deployed a solution.
> In the example screenshot, two categories of power changes are listed: AirMatch and Config. In the current lab, you are looking for the 'config' power changes.
> In most lab environments, you will only see the Config category at this point, since AirMatch typically only runs once every day.

Example screenshots with AirMatch Algorithm Assigned changes:

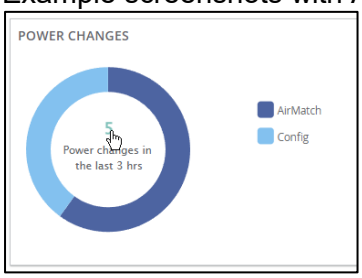

Power Changes (5)

Event Time	Reason	From Power (dBm)	To Power (dBm)	Band	Access Point
Sep 6, 2022, 13:43	Power Range Change	15	12	5 GHz	ap1
Sep 6, 2022, 13:42	Power Range Change	15	12	5 GHz	ap2
Sep 6, 2022, 12:02	Algorithm Assigned	18	15	5 GHz	ap1
Sep 6, 2022, 12:01	Algorithm Assigned	18	15	5 GHz	ap2
Sep 6, 2022, 12:01	Algorithm Assigned	9	7	2.4 GHz	ap1

Task 3: Configuring AirMatch

AirMatch is a service in Aruba Central that receives information from APs about activity on various channels. It processes this information into a channel plan for all the APs in your environment.

Objectives

- List AirMatch global schedule
- List AP timezone
- List ap debug command for timezone and airmatch/solver settings

Steps

The AirMatch deployment schedule is set at the global level. The application of the AirMatch solver settings will be applied based on the AP local time zone settings. Review the current time on your AP.

1. In Aruba Central, navigate to Context: **Global** > Navigation: **Tools** > Top: **Commands**
2. Select Device Type **Access Points** , select from available devices: **ap1**.
3. Select **All Category**, enter **clock** in the command filter.
4. **Add** the commands **show clock** and **show clock timezone** and **Run** the commands.

Example output:

```
COMMAND=show clock

Current Time     :2022-09-08 04:39:14


==================================
Output Time: 2022-09-08 08:39:16 UTC


COMMAND=show clock timezone

Current Timezone
----------------
Country       Timezone  DST Name  DST Recurring
-------       --------  --------  -------------
Eastern-Time  UTC-05    EDT       second sunday march 02:00 first sunday november
02:00
```

5. Take note of the current time.

Current Time: _____

Question: What is the timezone, how was this applied to the AP?

Answer: The timezone was configured at the AP group level > system settings. The APs automatically inherit the timezone from the group.

6. In Aruba Central, navigate to Context: **Global** > Navigation: **Devices** > Top: **Access Points** > Top right: **config**

> **NOTE:** You do this at the Global context level. The AirMatch deployment schedule is set once, as a global parameter in the system. The actual deployment to the devices will consider the device local time.

7. As an example in the lab, set the automatic deploy at the next top of the hour interval based on the AP time. For example, if the AP time showed 3:30, you may set the deploy time to 4:00.
8. Click **Save.**

Task 4: Client Association, Live Events, and Location information

Your customer would like to know how they can monitor, troubleshoot, and see events about clients that connect to the WLAN. You explain to the customer that Aruba Central provides client floorplan map location, historical events, and live events information. In this task you will explore the client information and client events.

Open Client Details context

You can find the various client events and monitoring options in the client details page. Let's start there.

1. In Aruba Central, navigate to Context: **Global** > Navigation: **Clients** > Top: **Clients**

Your wireless client should be in the list with a status of **connected**.

2. Click your **Client Name** in the list to see the client details > Overview page.

Client Location on Floorplan

You can see client locations on the floor plan from the client details context.

3. Use the navigation menu to click **Overview** > Top: **Location**.

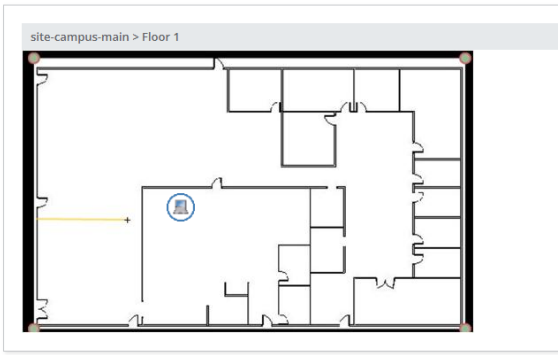

You see the site / floor of the client-connected AP. On the floorplan, only the current AP is displayed.

Association Events and Roaming latency

In these steps you will explore the historical association events for a client.

4. Use the navigation menu to click on **Overview** > Top: **Summary**.
5. Scroll down to the **Roaming Experience** tile.
6. On the top-right of this tile, click the **table** icon.

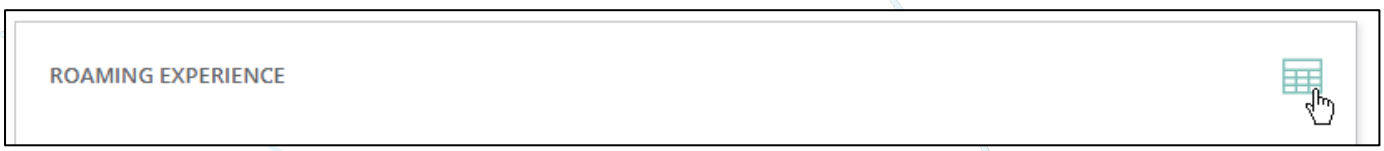

7. You see a historical list of association events and roaming latency.

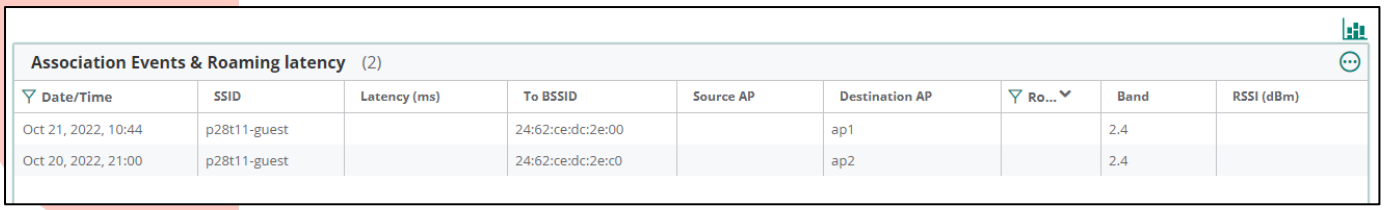

Association Events & Roaming latency (2)

Date/Time	SSID	Latency (ms)	To BSSID	Source AP	Destination AP	Ro...	Band	RSSI (dBm)
Oct 21, 2022, 10:44	p28t11-guest		24:62:ce:dc:2e:00		ap1		2.4	
Oct 20, 2022, 21:00	p28t11-guest		24:62:ce:dc:2e:c0		ap2		2.4	

Question: Must the client be connected to see this events table?

Answer: No, even when the client is offline, this history event table can be reviewed.

Clients Events

In the next steps you review client events recorded by Aruba Central.

8. You are still at the client device context.
9. Use the navigation menu to click on **Events** > Top: **Events**.
10. In this list you can explore various events for the client.
11. **Optional step**: You can change the displayed time period using the time selector at the top-right of the screen.

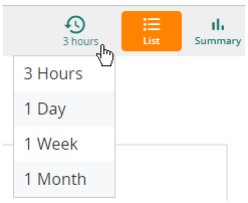

Client Live Events

Aruba Central supports Client Live Events for more advanced troubleshooting. This shows live client events as they connect, disconnect, or roam.

12. Use the navigation menu to click on **Live Events**.
13. The Live Events system automatically starts the troubleshooting session for the client MAC address.

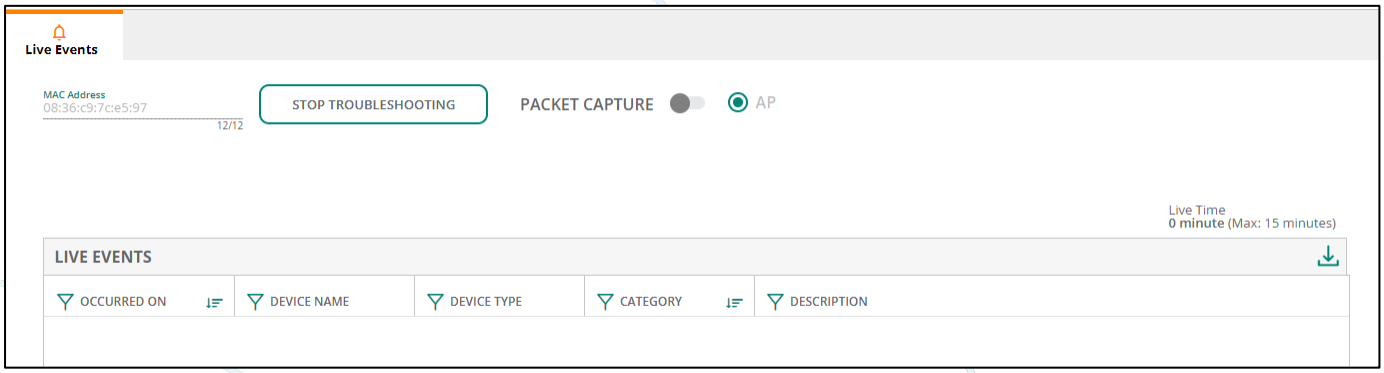

14. On PC1 disconnect from the guest WLAN.
15. In Aruba Central, the client Disassociation messages is displayed.

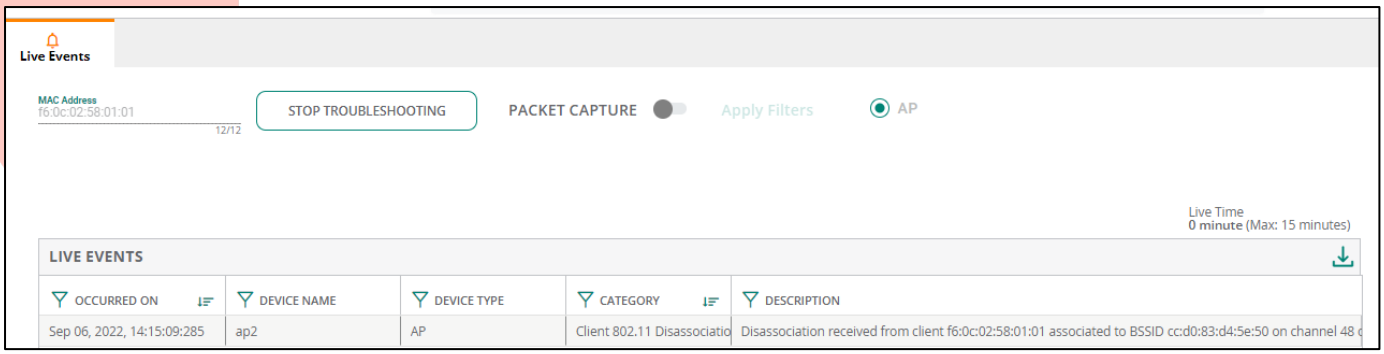

16. On PC1, connect to the p#tx-guest WLAN again.
17. In Aruba Central, review the events.

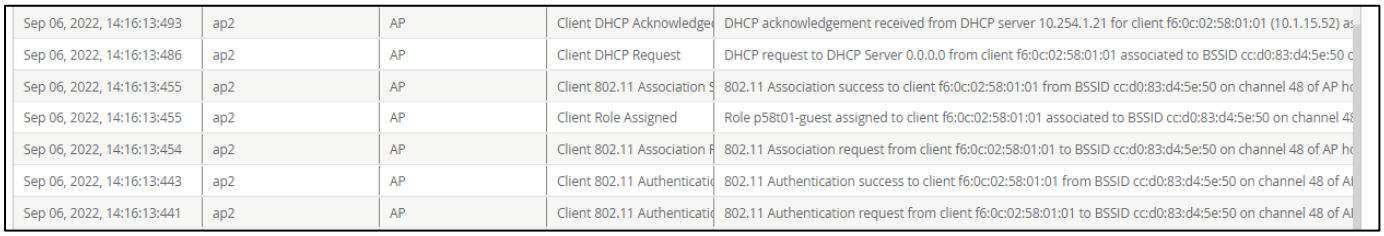

Question: What is the flow of the events?

Answer: Auth / Association / Role assignment / L3 address using DHCP

TIP:
You can export the events using the download CSV button, this can make it easier to share the data with a colleague.

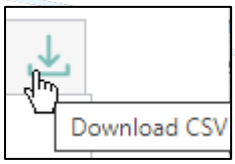

Example export of the events:

	A
1	OCCURRED ON,"DEVICE NAME","DEVICE TYPE","CATEGORY","DESCRIPTION"
2	Oct 21, 2022, 11:06:40:247,"ap1","AP","Client DNS Failure","DNS failure to _ldap._tcp.dc._msdcs.aruba-training.com detected for client 08:36:c9:7c:e5:97 associated to BSSID 24:62:ce:dc:2e
3	Oct 21, 2022, 11:06:40:245,"ap1","AP","Client DNS Failure","DNS failure to _ldap._tcp.dc._msdcs.aruba-training.com detected for client 08:36:c9:7c:e5:97 associated to BSSID 24:62:ce:dc:2e
4	Oct 21, 2022, 11:06:39:255,"ap1","AP","Client ARP Response","ARP response with sender IP 10.1.15.1 and sender MAC 12:01:00:00:01:00 received for target IP 10.1.15.50 and target MAC 08:
5	Oct 21, 2022, 11:06:39:253,"ap1","AP","Client ARP Response","ARP response with sender IP 10.1.15.1 and sender MAC 12:01:00:00:01:00 received for target IP 10.1.15.50 and target MAC 08:
6	Oct 21, 2022, 11:06:39:251,"ap1","AP","Client ARP Request","ARP request with sender IP 10.1.15.50 and sender MAC 08:36:c9:7c:e5:97 sent to gateway for target IP 10.1.15.1 in VLAN 15"
7	Oct 21, 2022, 11:06:39:248,"ap1","AP","Client ARP Response","ARP response with sender IP 10.1.15.1 and sender MAC 12:01:00:00:01:00 received for target IP 10.1.15.50 and target MAC 08:
8	Oct 21, 2022, 11:06:39:246,"ap1","AP","Client ARP Response","ARP response with sender IP 10.1.15.1 and sender MAC 12:01:00:00:01:00 received for target IP 10.1.15.50 and target MAC 08:
9	Oct 21, 2022, 11:06:39:243,"ap1","AP","Client ARP Request","ARP request with sender IP 10.1.15.50 and sender MAC 08:36:c9:7c:e5:97 sent to gateway for target IP 10.1.15.1 in VLAN 15"
10	Oct 21, 2022, 11:06:38:970,"ap2","AP","Client ARP Request","ARP request with sender IP 10.1.15.50 and sender MAC 08:36:c9:7c:e5:97 sent to gateway for target IP 10.1.15.1 in VLAN 15"
11	Oct 21, 2022, 11:06:38:967,"ap2","AP","Client ARP Request","ARP request with sender IP 10.1.15.50 and sender MAC 08:36:c9:7c:e5:97 sent to gateway for target IP 10.1.15.1 in VLAN 15"
12	Oct 21, 2022, 11:06:38:386,"ap1","AP","Client DHCP Acknowledged","DHCP acknowledgement received from DHCP server 10.254.1.21 for client 08:36:c9:7c:e5:97 (10.1.15.50) associated to
13	Oct 21, 2022, 11:06:38:378,"ap1","AP","Client DHCP Request","DHCP request to DHCP Server 0.0.0.0 from client 08:36:c9:7c:e5:97 associated to BSSID 24:62:ce:dc:2e:00 on channel 11 of AP
14	Oct 21, 2022, 11:06:38:294,"ap1","AP","Client Role Assigned","Role p28t11-guest assigned to client 08:36:c9:7c:e5:97 associated to BSSID 24:62:ce:dc:2e:00 on channel 11 of AP hostname a
15	Oct 21, 2022, 11:06:38:293,"ap1","AP","Client 802.11 Association Success","802.11 Association success to client 08:36:c9:7c:e5:97 from BSSID 24:62:ce:dc:2e:00 on channel 11 of AP hostnam
16	Oct 21, 2022, 11:06:38:292,"ap1","AP","Client 802.11 Association Request","802.11 Association request from client 08:36:c9:7c:e5:97 to BSSID 24:62:ce:dc:2e:00 on channel 11 of AP hostnam
17	Oct 21, 2022, 11:06:38:240,"ap1","AP","Client 802.11 Authentication Success","802.11 Authentication success to client 08:36:c9:7c:e5:97 from BSSID 24:62:ce:dc:2e:00 on channel 11 of AP h
18	Oct 21, 2022, 11:06:38:239,"ap1","AP","Client 802.11 Authentication Request","802.11 Authentication request from client 08:36:c9:7c:e5:97 to BSSID 24:62:ce:dc:2e:00 on channel 11 of AP h
19	Oct 21, 2022, 11:05:19:401,"ap2","AP","Client 802.11 Disassociation from Client","Disassociation received from client 08:36:c9:7c:e5:97 associated to BSSID 24:62:ce:dc:2e:c0 on channel 6 o
20	

An example export file can also be found in the ACAF Student Files on the MGMT PC desktop:

lab09.01-central-export_radio_list_1666291846829-example.csv

Task 5: Configuring WLAN Zones

The customer wants to ensure that the test guest WLAN is not advertised on the AP2.

Here are different strategies with pros and cons:

- **New Group**: Complete a new configuration. Complex to maintain the common settings with the other group.
- **Device Level Override**: Works for single device. Difficult to deploy to several devices.
- **AP Zones**: Easy to deploy. Only applies to the WLAN activation feature, no other configuration items can be pushed using this system.

In this task, the AP zones will be used to ensure that AP2 does not advertise the test guest WLAN anymore.

Objectives

- Understand AP Zones.
- Configure Zones.

Steps

1. In Aruba Central, navigate to Context: **Groups / campus-wifi-ui** > Navigation: **Devices** > Top: **Access Points**> Right top: **Config**
2. Click the **Access Points** page.

Question: What do you see under the WLANs column for your APs?

Answer: Both APs show 'All SSIDs selected'.

3. Select AP2 and use the **pencil** icon at the end of the record to edit the AP.
4. Click on the **WLANs** page to see the list of WLAN bindings.

Question: What WLANs are enabled?

Answer: The guest WLAN is enabled.

5. **Uncheck** the guest WLAN.
6. Click **Save Settings** to submit the configuration.
7. You will be returned to the AP list.

Question: What WLANs are listed for AP2 now?

Answer: There are no WLANs active for AP2.

Verify the BSS List on AP2

8. In Aruba Central, navigate to Context: **Groups / campus-wifi-ui** > Navigation: **Tools** > Top: **Commands**
9. In the **Device Type** dropdown select **Access Point**.
10. Click **Available Devices** and select **ap2**.
11. In the **Categories** list, select **All Category**. This will show all available commands.
12. In the **Commands** list, click the **filter** icon and type **bss**.
13. Select the **show ap bss-table** command, then click **Add**.
14. Click on **Run**.
15. After a few moments, review the commands.

Question: Are there any BSS entries active in the list?

Answer: No, AP2 does not have any WLANs active now.

Task Cleanup

16. Use the same procedure to enable all WLANs again on ap2.

Lab Cleanup: Disable guest SSID

To reduce the number of active SSIDs broadcasting in the remote lab, you may disable the guest SSID.

1. In Aruba Central, navigate to Context: **Groups / campus-wifi-ui** > Navigation: **Devices** > Top: **Access Points**> Top right: **config**
2. In the WLAN list, check the Network Enabled column.

Question: What is the current value?

Answer: Enabled.

3. Move the mouse over the SSID, icons will appear at the end of the line.

4. Click the **Wifi** icon to disable the WLAN.
5. Confirm the action with **Yes**.

Question: What is the value now for the Network Enabled column?

Answer: The value is “No” since you disabled the WLAN.

Wireless SSIDs

Name (Profile)	Security	Access Type	Traffic forwarding mode	Network Enabled
p58t01-guest	enhanced-open	Unrestricted	Bridge	No

NOTE: The same option can be used to enable the WLAN again.

You have completed this Lab !

Lab 10.01 Implementing a Corporate WLAN

Overview

Your customer would like to connect the internal employees and contractors to a secured WLAN that offers authentication and encryption. You learn that this can be achieved by configuring a WPA2 or WPA3 Enterprise WLAN. It uses an authentication (RADIUS) server for centralized authentication and authorization, and it offers encryption for the WLAN traffic.

In this lab you will configure and test the employee WLAN and test the connection using an employee and contractor user account.

Objectives

- Configure a WPA2/3-Enterprise WLAN.
- Configure a RADIUS server.
- Verify the secured WLAN connection.

Task 1: Create a Corporate WPA2/3 Enterprise WLAN

In this task you will create a WPA2/3 Enterprise secured WLAN, it will use 802.1X as the authentication protocol. You will explore the authentication and layer 2 VLAN access for the SSID. The AP should assign the traffic for the corporate wireless clients to VLAN 11.

Objectives

- Configure the WPA2/3 Enterprise WLAN
- Verify the deployment of the WLAN

Steps

1. In Aruba Central, navigate to Context: **Groups / campus-wifi-ui** > Navigation: **Devices**> Top: **Access Points**> Right top: **Config**

2. Under the WLAN page, click **Add SSID**.

General Page

3. On the General page, configure name **p#tx-employee**.

> **NOTE:** Make sure to replace the # value with your Pod number and x with your table number.
> For example, if you are using Table 07 in Pod 28, your WLAN name will be
>
> **p28t07-employee**
>
> Check with your instructor if you are not sure about the Pod and Table number.

Optional steps

The next steps show values as proposed in the Validated Solution Guide.
You may optionally apply them as an example to practice the settings.

> **TIP:** The latest version of the Aruba Validated Solution Guides can be accessed on this page:
>
> https://www.arubanetworks.com/techdocs/VSG/

4. Expand **Advanced Settings**.
5. Open the Broadcast/Multicast section.
6. For **Broadcast filtering**, set the value to **ALL**.
7. Set Dynamic Multicast Optimization (DMO) to **enabled**.
8. Set the DMO client threshold value to **40**.
9. Open the Transmit Rates (Legacy Only) section.

10. For **2.4 GHz** configure Minimum value as **5**.
11. For **5 GHz** configure Minimum value as **18**.

End of the optional steps.
Continue with the next steps!

5GHz Only Operation

In the remote lab environment, there is a lot of interference from neighboring APs, especially on the 2.4GHz band. You will limit this SSID to the 5GHz band.

1. Expand **Advanced Settings**.
2. Open the **Miscellaneous** section.
3. For the **Band** option select only **5 GHz**.
4. Click **Next**.

VLANs Page

5. On the **VLANs** page, for Client VLAN Assignment, select **Static**.
6. In the VLAN ID field, **remove VLAN 1**.
7. Enter VLAN ID **11**.
8. Click **Next**.

Security Page

9. On the **Security** page, move slider to **Enterprise**.
10. Leave key management to the default (**WPA3-Enterprise CCM128)**.

An enterprise WLAN requires a RADIUS authentication server. In the lab the ClearPass server with IP 10.254.1.23 will be used as the authentication server. It has been preconfigured to support 802.1X wireless authentication for your APs.

11. For **Primary Server**, click the '+' sign to add an authentication server.
12. In the New Server dialog box, enter these values:

Name	cppm1	
IP Address	10.254.1.23	
Shared key	Aruba123!	
Retype Key	Aruba123!	
Dyn Authorization	enabled	(You nay need to scroll down to see this option)

13. Click **OK** to add the server.
14. Verify that the **Primary Server** field now has **cppm1** selected.
15. Open the **Advanced Settings** section.

RADIUS Interim Accounting ensures that the AP updates the RADIUS server with status updates about the authenticated clients, such as the amount of traffic generated and whether the clients are still active (online). It is recommended to enable RADIUS Accounting.

16. Expand the **Accounting** section.

17. For the **Accounting** field, select **Use Authentication Servers**.
18. Set the **Accounting Interval** to **5 min**.
19. Open the **Fast Roaming** section.
20. Set **802.11K** to enabled.
21. Click **Next**.

Access Page

22. On the Access page. Leave the default to **unrestricted**.
23. Click **Next**.
24. To complete the configuration, click **Finish**.
25. Wait for the deployment wizard to complete and confirm by clicking **OK**.

You should now see the p#tx-employee WLAN in the Wireless SSIDs list.

Verify configuration deployment

26. Click the **Configuration Audit** link.

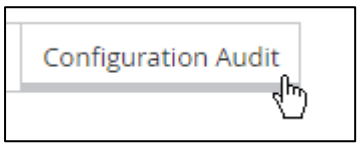

NOTE: If you don't see the Configuration Audit page, you should enable the **Advanced View**.

27. Under the LOCAL OVERRIDES & CONFIGURATION SYNC ISSUES, check the Configuration Status window. The **Not In Sync** value should be **0**.

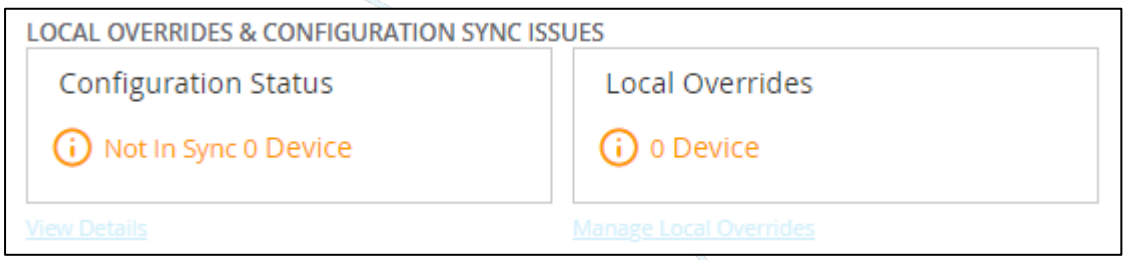

NOTE: It may take a minute for Central to push the configuration to the devices. In the meanwhile, **Not In Sync** will report **2** devices. Refresh this page until the Not In Sync shows 0 Devices, this indicates the synchronization was successful.

SSID configuration basics

Now you will review the commands that were generated by the UI Wizard.

28. In Aruba Central, navigate to Context: **Groups / campus-wifi-ui** > Navigation: **Audit Trail** > Top: **Audit Trail**
29. Look for the entry with description '**Created/Updated WLAN profile'**. Click on the three dots at the end of this entry to see the details.

There are some key sections in the configuration to note:

WLAN Access-rule

The WLAN Access-rule acts as an Aruba User role in the AP configuration. It is defined in the global configuration of the AP. For every WLAN that is created, there will automatically be a User role (access rule) based on the SSID name. This will be the default User Role that authenticated users get assigned when they connect to this SSID.

```
wlan access-rule p58t01-employee
 utf8
 rule any any match any any any permit
 exit
```

WLAN Authentication Server

The RADIUS authentication server is defined in the global configuration of the AP. This means that the same RADIUS server can be used in other SSID configurations as well.

```
wlan auth-server cppm1
 ip 10.254.1.23
 key ********
 port 1812
 acctport 1813
 rfc3576
 cppm-rfc3576-port 5999
 exit
```

WLAN SSID profile

The SSID profile is the WLAN object that defines the wireless properties and wireless authentication settings. It also includes the default Layer 2 forwarding settings (VLAN assignment) for clients that connect to this SSID.

```
wlan ssid-profile p58t01-employee
 essid p58t01-employee
 opmode wpa3-aes-ccm-128
 vlan 11
 rf-band 5.0
 type employee
 captive-portal disable
 dtim-period 1
 broadcast-filter all
 radius-accounting
```

```
radius-interim-accounting-interval 5
inactivity-timeout 1000
g-min-tx-rate 5
a-min-tx-rate 18
max-authentication-failures 0
blacklist
dynamic-multicast-optimization
dmo-channel-utilization-threshold 90
max-clients-threshold 64
enable
dot11r
utf8
okc
dot11k
openflow-enable
dmo-client-threshold 40
auth-server cppm1
exit
```

Question: Under the WLAN **ssid-profile**, what command controls how wireless clients will see this WLAN object when they scan for wireless networks?

Answer: The command **essid** controls the wireless ESSID name.

30. **Close** the audit details window.

Verify the received Cloud Commands on the AP

On the AP, you can view all commands that are pushed by Aruba Central. This can be useful in case a command fails to execute.

31. In Aruba Central, navigate to Context: **Groups / campus-wifi-ui** > Navigation: **Tools** > Top: **Commands**
32. In the **Device Type** dropdown list, select **Access Point.**
33. In the **Available Devices** dropdown list, select **ap1**.
34. In the **Categories** table, scroll down in the list and select **Central**.
35. In the **Commands** table, select **Show Central Last Configuration Received**.
36. Click **Add** to move it to the Selected Commands column.
37. Click **Run** to run the command on the selected devices. You will need to wait a few moments for the command to be executed.
38. The output will show all the commands that the AP received and the time of execution.

Question: How can you see if the command was executed successfully?

Answer: The output includes **OK** after a line that was accepted by the AP CLI.

NOTE: Renaming a WLAN SSID

The WLAN profile name cannot be changed in Aruba Central, however the ESSID name can be updated at a later point in time.
When this is done, please note that in Aruba Central, the WLAN profile name will not be the same as the actual ESSID name and it will be displayed behind the ESSID name.

Example configuration where the **profile name is the same as SSID name:**

Example configuration with the ESSID changed to p58t01-corporate on the profile name p58t01-employee, so the **profile name is different from the SSID name**:

There is no need to change the SSID name in the lab!

Verify the SSID is broadcasting on the AP

In this section, you will check the BSS table of the AP1 to verify that the SSID is being broadcasted by the AP using the command "**show ap bss-table**" on the AP.

39. Click **Remove** to clear the Selected Commands window.
40. In the **Categories** table, select **Wireless**.
41. In the Commands table, select AP BSSID Table.
42. Click **Add** to move it to the Selected Commands column.
43. Click **Run** to run the command on the selected devices. You will need to wait a few moments for the command to be executed.

Once the output is displayed, you can maximize the output screen.

44. Click the maximize button at the right-top of the command output.

Question: Do you see the employee ESS broadcasted on the ap1?

Answer: Yes, the ESSID is the name of the employee SSID as configured in the previous steps. Since it was only enabled on the 5GHz radio, it appears only on the **a** phy (radio).

45. Exit the maximized view by clicking on the  button at the right-top.

Test the AAA server

On the AP, you can use a **aaa test-server** command to verify RADIUS server communication. This will send a RADIUS access request to the configured RADIUS server. This test allows you to verify that the shared secret and RADIUS client settings are matching between the AP and the RADIUS server.
You must know an example username and password to run the test. In this lab, use username **employee** with password: **Aruba123!**

The **aaa test-server** command is not included in the standard command list, therefore you will open an SSH console session using Aruba Central.

46. In Aruba Central, navigate to Context: **Groups / campus-wifi-ui** > Navigation: **Tools** > Top: **Console**
47. In the **Device Type** dropdown list, select **Access Point.**
48. In the **Available Devices** dropdown list, select **ap1**.
49. For the **username** enter **admin**.
50. For the **password** enter **Aruba123!**
51. Click **Create New Session** to request the console session.

After a few seconds, the console connection will appear.

NOTE: If you have issues with this connection in your own production environment, make sure to check the Information icon  at the right-top:

The console connection is a reverse SSH connection over port 443 to an Aruba Central cloud host. Some new generation firewalls will detect that this is not HTTPS over port 443 and they may block this communication. You must make sure that this is explicitly allowed in the firewall.

Close the information box with the **X** button.

---

52. In the AP console connection, run the **aaa test-server** command. You can use ? and TAB for command completion. The server **cppm1** is the name you have used during the WLAN wizard.

```
aaa test-server cppm1 username employee password Aruba123! auth-type pap
```

Question: What is the result of the command?

Answer: The AP reports a success authentication.

> **NOTE:** If the test is not successful, you should run the WLAN wizard again and verify the CPPM1 IP address (10.254.1.23) and the shared secret (Aruba123!) on the security page of the wizard.

> **NOTE:** The PAP method (cleartext Password Authentication Protocol) is different from the EAP-PEAP-MSCHAPv2 authentication method that is used by the Wireless client. Your RADIUS server must be configured with PAP support to run this test command successfully. The lab ClearPass server has been configured with a test service for this purpose.

You will now verify on the ClearPass RADIUS server that this test request was processed.

53. On MGMT PC, open a browser to

```
https://10.254.1.23/tips
```

54. Login with username **admin** / password **Aruba123!**
55. In the left pane, expand the **Monitoring** section.
56. Under Monitoring, expand **Live Monitoring**.
57. Click on **Access Tracker**. Access Tracker shows the latest authentication requests that were processed by the server. This should include the test request you have made in the previous steps.
58. There will be a record with username **employee**.

Question: What is the Login Status for this entry?

Answer: The result for this request is ACCEPT - the server has accepted the login request. This resulted in the success message for the AP aaa test-server command.

59. This concludes the test procedure. You may leave the ClearPass session open.

Task 2: Verify Corporate Access with Wireless Client

In this task you will connect to the corporate wireless network using PC1. The authentication method will be EAP-PEAP-MSCHAPv2.

The wireless client in the lab supports only WPA2, therefore only WPA2 can be tested.

Objectives

- Connect with a Windows client to a WPA2 or WPA3-Enterprise network.
- Verify the WiFi client connection in Aruba Central.
- Verify the WiFi client connection in the wired network.
- Verify the WiFi client connection on the AP.

Steps

Use the PC1 to make a Wireless Connection

1. Open a connection to your PC1 system.
2. Wireless LAN list, select **p#tx-employee**, connect with username **employee** password **Aruba123!**

Example screenshot, your pod and table may be different.

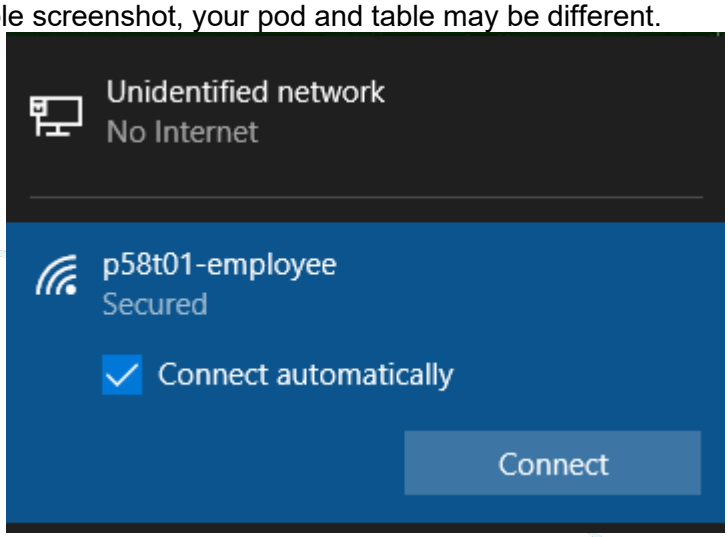

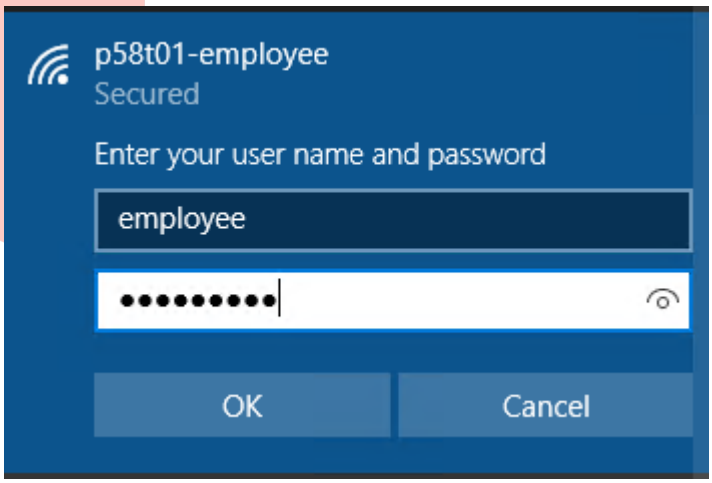

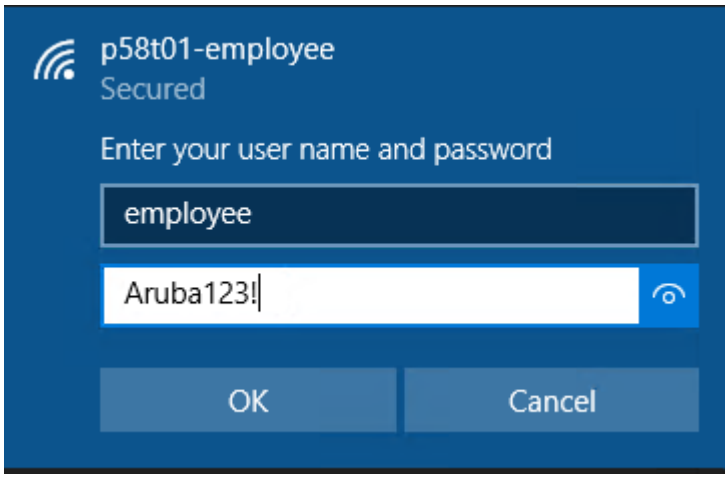

3. Click **OK** to submit your credentials.
4. Click **Connect** to accept the certificate warning.

NOTE: The lab test PC is not part of a managed environment. In a managed, corporate environment the WLAN profile can be provisioned using Group Policies or a Mobile Device Manager.

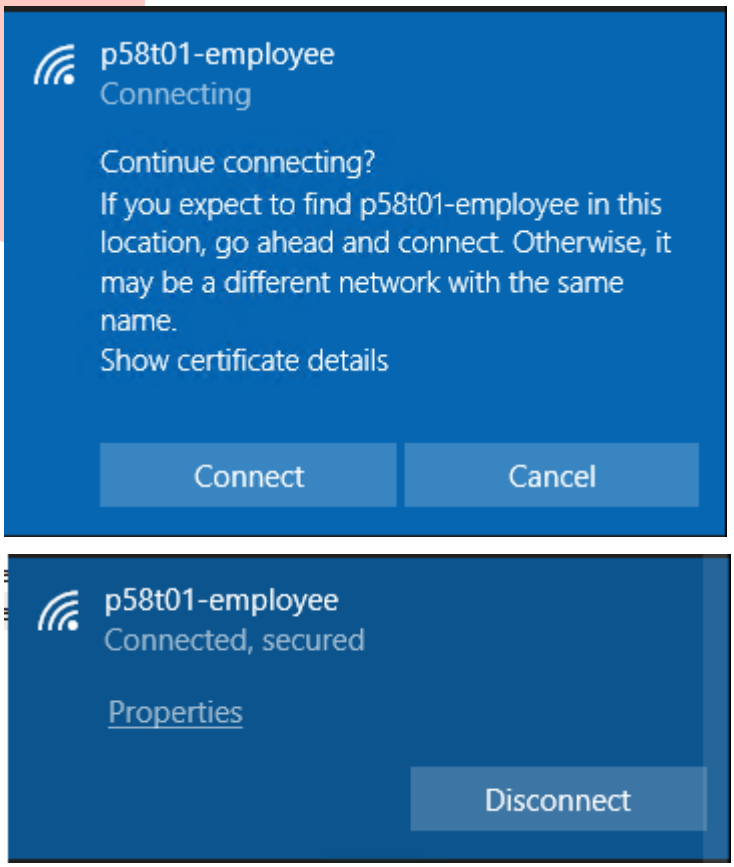

Check on the client side the assigned IP address.

5. Right-click the wireless icon in the status bar.
6. Click **Open Network & Internet Settings**.

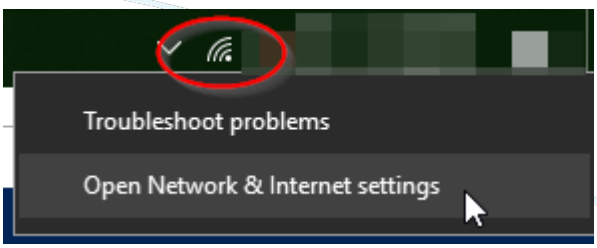

7. Click Change Adapter Options.
8. Right-click **Wi-Fi** connection > **Status.**
9. In the Status dialog > **Details**

Example Screenshot

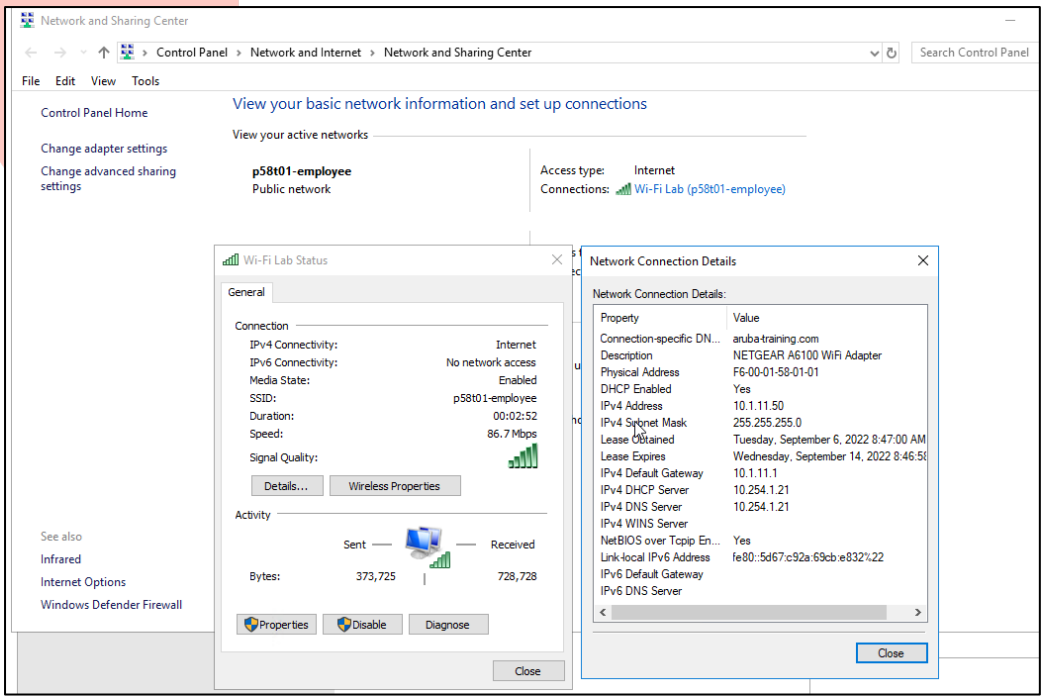

Question: Take note of the physical address for the client?

Answer: _____

Question: What is the IPv4 address the client received? Does this match with the configured VLAN?

Answer: The IP address should be in the 10.1.11.0/24 subnet, this is indeed VLAN 11.

10. Close the **Details** and **Status** windows.

Verify the client connection in Aruba Central

11. In Aruba Central, navigate to

Context: **Groups / campus-wifi-ui** > Navigation: **Clients**> Top: **Clients**

Question: Do you see the client employee in the list?

Answer: Yes, the client employee is listed.

Question: What is the status for the client?

Answer: The status is Connected. This indicates that the user is now online.

> **NOTE:** Aruba Central can also show the historical clients when they are no longer connected, this is convenient for troubleshooting.

You can use the top-right time period button to control how far back the historical clients should be displayed.

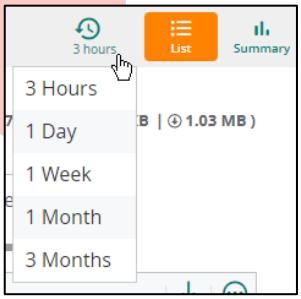

Question: What is the VLAN and IP address that was assigned to the client?

Answer: The VLAN id is 11, the IP address should match the address you found on the PC1 client.

Client Details

The previous screen shows a list of all the clients connected at the group level. You will now investigate client details for an individual client.

12. Click on client name '**employee'** to access the client details.

TIP: You can also use the AI search at the top. Simply type employee<ENTER> and the client will be listed in the output. Click on the name employee to access the client details page.

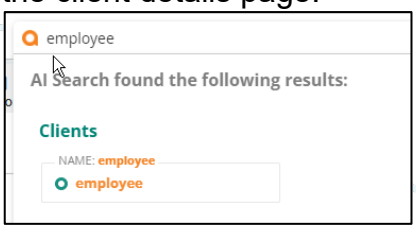

.

The context will change to the client name.

13. You are now on the client employee details > Overview > Summary page.

Question: What are the elements you see in the DATA PATH?

Answer: The **client** is connected via the **SSID** employee.to the **AP**1(or AP2).

14. Take note of the access point the client is connected to:

_____

Question: What are the other information tiles you can see in the Client details Overview > Summary page? Make sure you scroll down to see all the tiles.

Answer: The other tiles show:
Client information
Network
Connection
Throughput
Health
Signal Quality
Retry Frames
Tx/Rx Rate
Roaming Experience

Example of the **Client, Network** and **Connection** tiles:

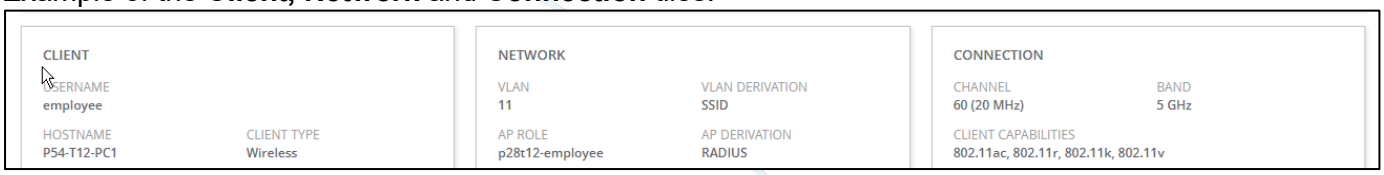

Verify the client connection in the wired network .

You will now look at the access switch details to verify you can see the client MAC address. The wireless client MAC can be seen on the access switch since the wireless client is connected to a bridged WLAN.

15. In the AI search bar, enter the name of the access switch. Use sw-edge1 or sw-edge2, depending on the AP your client was connected to.
16. In the AI search result, click on the name of your switch.

17. While on the edge switch details page, navigate to

Navigation: **Tools**> Top: **Console**

18. For username enter **admin**, password **Aruba123!**
19. Click **Create New Session**.
20. When the console session is connected, check the VLANs on the AP port.

```
show vlan port 1/1/2
```

Example for sw-edge1, in your setup you may be connected to sw-edge2.

```
sw-edge1# show vlan port 1/1/2

--------------------------------------------------------------------------------
VLAN  Name                                Mode            Mapping
--------------------------------------------------------------------------------
4     VLAN4                               native-untagged port-access
11    VLAN11                              trunk           port-access
12    VLAN12                              trunk           port-access
13    VLAN13                              trunk           port-access
14    VLAN14                              trunk           port-access
15    VLAN15                              trunk           port-access

Overridden VLAN list: 1
```

Question: Is VLAN 11 listed, what is the mode for VLAN 11?

Answer: Yes, VLAN 11 is listed, it is in trunk mode, this is VLAN tagged traffic.

21. Check the MAC-address table on the AP port.

```
show mac-address-table port 1/1/2
```

```
sw-edge1# show mac-address-table port 1/1/2
MAC age-time            : 300 seconds
Number of MAC addresses : 3

MAC Address          VLAN     Type                        Port
-------------------------------------------------------------
20:4c:03:c6:09:50    4        port-access-security        1/1/2
20:4c:03:c6:09:50    11       port-access-security        1/1/2
08:36:c9:7c:e5:97    11       port-access-security        1/1/2
```

Question: Do you see the PC wireless MAC address in the list and in the correct VLAN?

Answer: Yes, the PC wireless MAC is listed in VLAN 11.

22. Exit the console connection.

```
exit
```

Verify the client connection on the AP.

In this section you will verify on the AP the connected client status.

1. Use the **AI search** box to navigate to the **employee** client details page.
2. In the **DATAPATH** tile, click on the **AP**. Central will take you to the AP device details page.

> **NOTE:** It takes some time for the device links and topology to be established. Therefore, it is possible that you don't see the AP connected to the edge switch, or you may even see the AP as connected to the Aggregation switch.
> You may ignore this and continue with the next steps. This will be corrected by Central over time.

3. While on the AP details page, navigate to Navigation: **Tools**> Top: **Console**
4. For username enter **admin**, password **Aruba123!**
5. Click Create New Session.
6. When the console session is connected, check the clients.

```
show clients
```

```
ap1# show clients

Client List
-----------
Name       IP Address  MAC Address        OS      ESSID            Access Point
Channel  Type  Role             IPv6 Address                Signal    Speed (mbps)
----       ----------  -----------        --      -----            ------------  -----
--  ----  ----             ------------                    ------    ------------
employee  10.1.11.52  08:36:c9:7c:e5:97  Win 10  p28t11-employee  ap1           56
AC    p28t11-employee  fe80::753c:42f5:cdf5:6f7a  69(good)  173(good)
Number of Clients   :1
Info timestamp      :69628
```

Question: Do you see the client's name and MAC address?

Answer: Yes

7. Check the detailed output of the clients using show clients debug.

```
show clients debug
```

```
ap1# show clients debug

Client List
-----------
Name      Host Name    IP Address  MAC Address        OS      ESSID            Access
Point  Channel  Type  Role             IPv6 Address              Signal    Speed
(mbps)  *Reauth Age  *Reauth Interval  *Reauth ESSID    Auth Type    *Authenticated
DEL  Age  Vlan      *ESSID                    *Private role info  Accouting Session
Name  BSSID             Idle Timeout  csum      *mcast groups  *Acct Interval
*Class Attribute
*Dhcp-Opt Vlan  *Dhcp-Opt role  Intercept  *Offline  *FB Token  *FB RxBytes  *FB
TxBytes  *SLAAC IP Address  *Link Local IP Address     *DHCP Status  *DHCP v6 Status
----      ---------    ----------  ----------         --      -----             ------
------  -------  ----  ----             ------------              ------    --------
----  -----------  ----------------  -------------    ---------    --------------  --
-  ---  ----      ------                    ------------------  --------------------
--  -----             ------------  ----      -------------  --------------  -------
---------
--------------  --------------  ---------  --------  ---------  -----------  --------
---  -----------------  ----------------------     ------------  ---------------
employee  P54-T11-PC1  10.1.11.52  08:36:c9:7c:e5:97  Win 10  p28t11-employee  ap1
56       AC    p28t11-employee  fe80::753c:42f5:cdf5:6f7a  69(good)  173(good)
432          0                 p28t11-employee  802.1x/PEAP  yes             no   13
11(SSID)  p28t11-employee(EAP-PEAP)  148(RADIUS-7fff)    employee
24:62:ce:dc:2e:10  1000          2ed05adc  (0)            300
5e95b071935f47e09e15952e26b2116ab90b0000000000052303030303030312d30312d36333532376
263350000000000000000000000  0,NONE          ,,0-0           no         no
null       null         null         0.0.0.0           fe80::753c:42f5:cdf5:6f7a
Completed     Soliciting
```

Hint: You may want to maximize the screen.

Question: What is the VLAN this client is connected to?

Answer: The VLAN shows 11.

8. Check the mac-address-table on the AP using the show datapath bridge command.

```
show datapath bridge
```

9. You may also apply an filter to include lines with the text '11' to reduce the output.

```
show datapath bridge | include 11
```

Example output:

```
ap1# show datapath bridge | include 11
8    bond0                    3      1       0/0/0   106  1500   117049     79215
FTQBW
25   aruba000                 1      11    148/0/0     0  1500      915       809  B
12:01:00:00:01:00  11    11            dev8                G         0
B8:D4:E7:D9:E5:00  11    11            dev8                          0
B8:D4:E7:DA:80:00  11    11            dev8                          0
08:36:C9:7C:E5:97  11    11            dev25                         0
ap1#
```

Question: Do you see the client MAC address in the list?

Answer: Yes, it is listed in VLAN 11.

Question: Are there any other MAC addresses in the VLAN 11?

Answer: Yes, the aggregation switch's Active Gateway MAC address 12:01:00:00:01:00 is also listed. This is the default gateway for the clients in VLAN 11.

This concludes the AP client review.

10. Exit the console connection.

```
exit
```

Task 3: Connect to Corporate Network with Contractor User

In this task you will connect to the corporate wireless network using PC4 as a contractor user.

Objectives

- Practice the client connection.

Steps

Use PC4 to make a Wireless Connection

1. Open a connection to your PC4 system.
2. Verify the LAN interface is disabled. Disable it if needed.
3. Verify the Wifi interface is enabled. Enable it if needed.

Note: Use the steps from the previous task if you need help to find the network adapters in Windows.

4. Wireless LAN list, select **p#tx-employee**, connect with username **contractor** password **Aruba123!**
5. In Aruba Central, navigate to Context: **Groups / campus-wifi-ui** > Navigation: **Clients**> Top: **Clients**

6. Verify that the contractor user is connected to the network. You can use the Refresh button to refresh the client list.

Example client list:

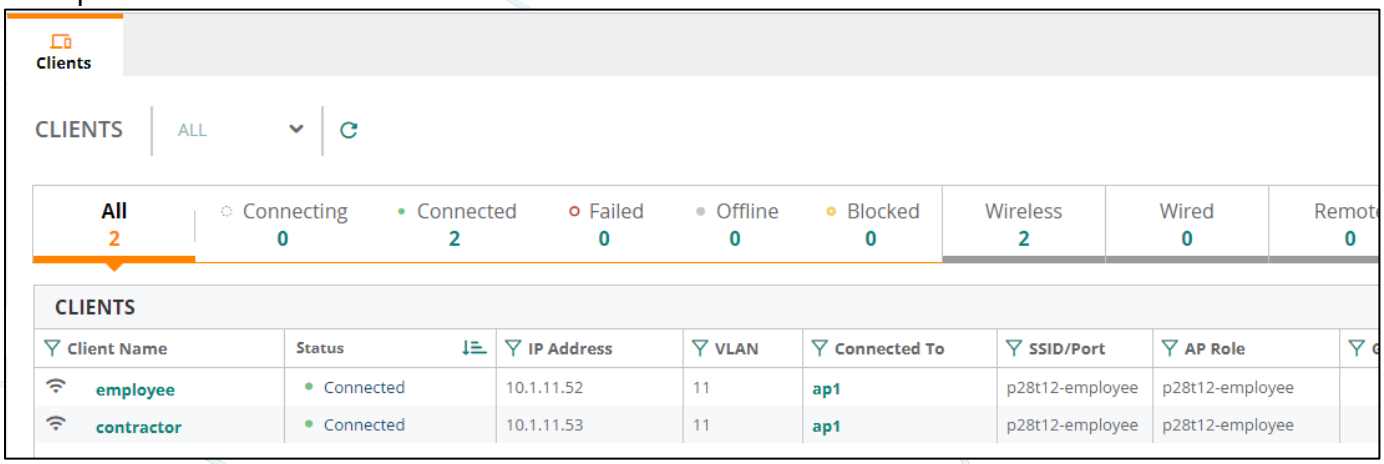

You have completed the lab!

Lab 10.02 Implementing Access Control

Overview

Your senior colleague is pleased with how you configured and validated the 802.1X secured WLAN. Now the customer would like to enforce some access control policies to the wireless clients. In this lab you will explore different access control options for AOS10 bridged WLANs. This will include the default role configuration, creating new user roles for the employee and contractor roles, applying access control using roles and assigning VLANs.

Objectives

- Explore the default SSID user role access
- Leverage authentication-based user roles
- Use Aruba user role VSAs
- Perform user role VLAN assignment
- Explore AppRF and Application Statistics

Task 1: Explore the default SSID user role access

In this task, you explore the default SSID-based user role access. When you create a WLAN, a default user role is automatically created for that WLAN, based on the WLAN(SSID) name. The rules defined in this default SSID role are applied to any device that connects to this SSID by default. The default role can be overruled by RADIUS (ClearPass) assigned roles or custom role derivation rules. You setup a RADIUS example in upcoming tasks. In this task, you will also explore the default SSID role.

Objectives

- Understand the default SSID user role.
- Apply network access control using the default role.

Steps

Explore the default SSID role assignment.

1. In Aruba Central, navigate to Context: **Groups / campus-wifi-ui** > Navigation: **Clients**

Both PC1 (employee) and PC4 (contractor) wireless clients should be connected.

> **NOTE:** If either PC1 or PC4 lacks a wireless connection, set that up now.

2. The client list includes an AP role column. You may need to scroll to the right to see it.

> **TIP:** You can change the column display order by grabbing the top line of the column header and moving it left or right.
>
> 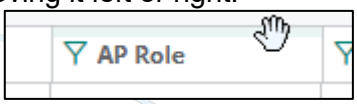
>
>
> Example clients tables:
>
> CLIENTS
>
Client Name	MAC Address	Status	IP Address	VLAN	Connected To	AP Role
employee	f6:00:01:58:01:01	Connected	10.1.11.50	11	ap1	p58t01-employee
contractor	f6:00:01:58:01:04	Connected	10.1.11.53	11	ap1	p58t01-employee
>
> CLIENTS
>
Client Name	Status	IP Address	VLAN	Connected To	SSID/Port	AP Role
contractor	Connected	10.1.11.53	11	ap2	p28t12-employee	p28t12-employee
employee	Connected	10.1.11.52	11	ap1	p28t12-employee	p28t12-employee

Question: What is the AP Role that is currently assigned to the employee and contractor users?

Answer: Both users currently have the role that is named p#tx-employee.

Question: What does the role name seems to match?

Answer: The role name matches the WLAN object name. When you create a new WLAN, a new 'default access role' is automatically created, based on the WLAN name.

Review the default SSID role

3. In Aruba Central, navigate to

Context: **Groups: campus-wifi-ui** > Navigation: **Devices**> Top: **Access Points**> Right top: **Config**

4. Click on the **Security** page.
5. Expand the **Roles** section.

Question: What roles currently exist in the configuration?

Answer: There are currently four roles:
- Default_wired_port_profile
- p#tx-employee
- p#tx-guest
- wired-SetMeUp

Two of these are default. The employee and guest roles were created when you added a new WLAN.

Review the p#tx-employee role.

6. Click the role name **p#tx-employee**.

Question: What are the access rules for this role?

Answer: There is one access rule: "Allow any to all destinations". This means the role does not have any restrictions and allows unrestricted network access.

Configure Network Based Access control

The customer wants to start with a simple access control measure. They want to ensure that whoever connects to the corporate wireless network WLAN cannot access IP subnet 10.1.0.0/24. This is the subnet with aggregation switch loopback IP addresses. You have been asked to configure this network access control on the employee WLAN.

Setup test traffic for Access Control

Configure both clients with a continuous ping to the 10.1.0.2 IP address (the loopback IP of sw-agg1).

7. On PC1 (the employee PC), start a continuous ping to 10.1.0.2. This should be successful. Leave the ping running.

```
ping 10.1.0.2 -t
```

8. On PC4 (the contractor PC), start a continuous ping to 10.1.0.2. This should be successful. Leave the ping running.

```
ping 10.1.0.2 -t
```

Configure the Access Control

Now you will implement access control to block access to the “off-limits” subnet - 10.1.0.0/24.

9. In Aruba Central, navigate to Context: **Groups / campus-wifi-ui** > Navigation: **Devices**> Top: **Access Points**> Right top: **Config**
10. On the **WLANs** page, select the **employee** SSID and **edit** it using the pencil icon.
11. Click the **Access** page.

Question: What is the current Access rules slider status?

Answer: Unrestricted. This reflects the access rule: **allow any to all destinations** in the default user role.

12. Move the slider to Network Based.

Question: What are the listed access rules?

Answer: There is 1 access rule: Allow any to all destinations.

13. Click on **Add Rule** to create a new entry.
14. In the **Rule Type** dropdown list, select **Access Control.**
15. For **Service**, leave the default Network – **Any**.
16. For **Action**, select **Deny**.
17. For Destination, select To a network: IP 10.1.0.0 and netmask 255.255.255.0.
18. Click **OK** to add the rule.
19. Verify in the Access Rules table that the Deny rule is listed **before** the Allow any rule.

NOTE: You can use the  icon at the start of the line to change the order of entries by dragging them up or down.

Question: Why is it important that the Deny rule is listed first?

Answer: Rulesets are processed from top to bottom. When there is a match on the condition, the action is applied. If the '**Allow any to all destinations**' were the first entry, all traffic would match this rule and be allowed. No traffic would ever hit the next rule, and so no traffic would be denied.

20. Click **Save Settings** to save the configuration.
21. Wait for the configuration wizard to complete and confirm with **OK**.

Verify the Access Control.

22. On PC1, verify that the ping no longer responds.
23. On PC4, verify that the ping no longer responds.
24. In Aruba Central, navigate to Context: **Groups / campus-wifi-ui** > Navigation: **Clients**

Question: Did the user role for the employee and contractor change?

Answer: No, they still have the p#tx-employee AP Role assigned to them.

25. In Aruba Central, navigate to Context: **Groups / campus-wifi-ui** > Navigation: **Devices** > Top: **Access Points** > Right top: **Config**
26. Click on the **Security** page, open the **Roles** section.
27. Select the role **p#tx-employee**.

Question: What do you notice for the Access Rules for this role?

Answer: The deny access rule was added to the rule set of this role.

Question: What do you know now about the WLAN wizard **Network Based** access and **Unrestricted** options?

Answer: The unrestricted option configures the Allow any to all destinations in the default SSID role, while the Network Based access configures user-defined rules in the default SSID role.

> **NOTE:** You could also have configured the custom rule directly under the Security > Roles. The WLAN Wizard would have reflected this custom rule in the UI as Network Based access.

Similarly, if you removed the custom deny rule in the role, the WLAN Wizard UI automatically shows Unrestricted access again.

Task 2: Authentication based user roles

The customer is happy that you managed to implement this access control, but they realized that this access control is not specific enough for their requirements. They want the employee users to have full access to the network, while the contractors should be denied access to subnets 10.1.0.0/24 and 10.1.3.0/24. The customer wants an easy way to block additional contractor destinations, should the need arise.

Your senior colleague explains that each client can be assigned a different user role, based on RADIUS authentication. To do this, you leverage an Aruba Vendor-Specific RADIUS attribute, named Aruba-User-Role. In fact, your colleague has already configured the ClearPass Policy Manager to return the user role *employee* for the user employee and the Aruba-User-Role *contractor* for the user contractor.

Review the Assigned RADIUS attributes

In these steps you will login to the ClearPass RADIUS server and review the RADIUS attribute that is returned when employee and contractor users login.

1. Using the MGMT PC, open a browser and connect to ClearPass Policy Manager

https://10.254.1.23/tips

2. Login using admin / Aruba123!
3. In the left pane, navigate to Monitoring > Live Monitoring > Access Tracker
4. Click the latest **employee** record in the list to open its **Summary** page.
5. Click the **Output** page and expand the RADIUS Response.

3.	P58-T01-CPPM	RADIUS	10.1.4.50	0	F6-00-01-58-01-01	employee	acaa-wireless-dot1x	ACCEPT	2022/08/02 12:40:11	aruba-role-employee

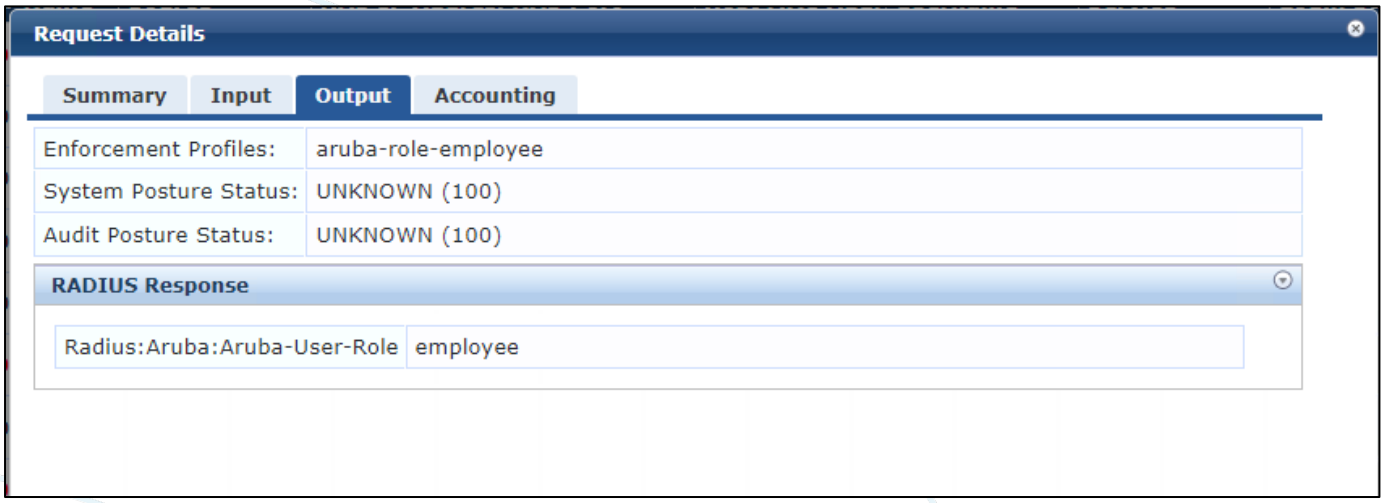

Q: What is the Aruba-User-Role RADIUS attribute that ClearPass returns (output) to the AP when the employee connects?

Answer: The value is **employee**

6. Click **Close** for the Request details.
7. Repeat this for the **contractor** authentication.

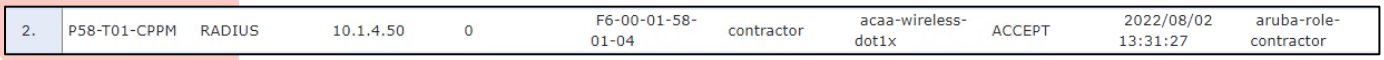

2.	P58-T01-CPPM	RADIUS	10.1.4.50	0	F6-00-01-58-01-04	contractor	acaa-wireless-dot1x	ACCEPT	2022/08/02 13:31:27	aruba-role-contractor

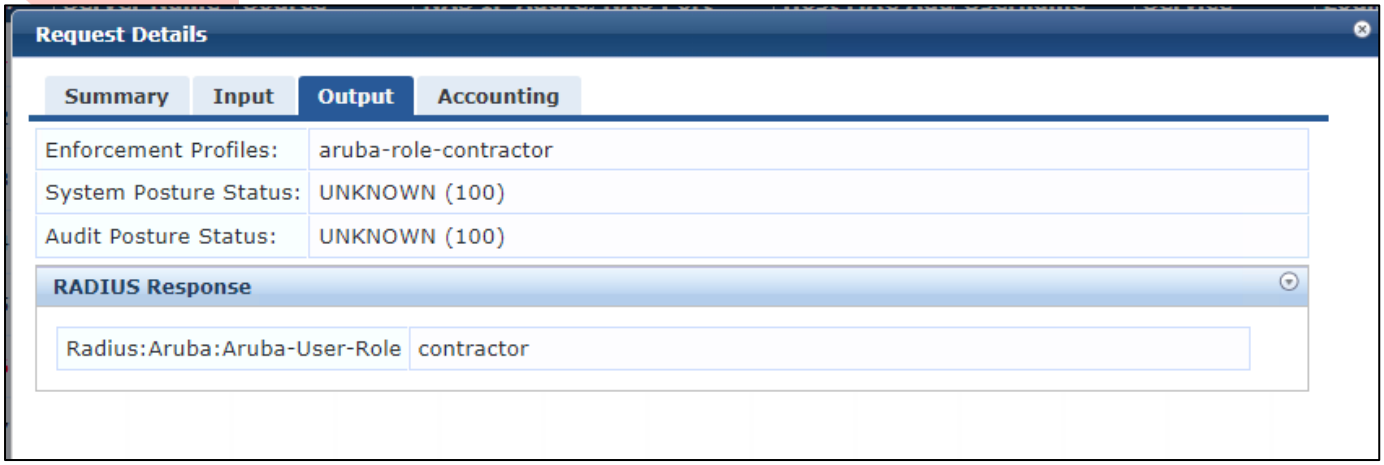

Question: What is the Aruba-User-Role RADIUS attribute that ClearPass returns (output) to the AP when the contractor connects?

Answer: The value is **contractor**

Question: Since the RADIUS server was already assigning a specific user role for both employee and contractor, why didn't these connections receive the contractor and employee user roles assigned to them?

Answer: The AP tried to process the requested role, but since the roles employee and contractor don't exist in the AP configuration, the users were assigned the SSID default role.

Question: For what scenario would you make the default SSID role very restrictive?

Answer: If a mistake is made on the CPPM (RADIUS) server, user access would be very limited, instead of full access.

8. **Close** the request details.

Configure user roles on the AP

In this section you will create additional user roles for employee and contractor.

9. In Aruba Central, navigate to Context: **Groups / campus-wifi-ui** > Navigation: **Devices**> Top: **Access Points** > Right top: **Config**
10. On the **WLANs** page, select the **employee** SSID and **edit** it using the pencil.
11. Click the **Access** page.
12. Move the slider to **Role Based**.

Question: What roles do you see in the Role list?

Answer: The same roles as under Security > Roles.

Question: What rules do you see under Role Assignment Rules?

Answer: There is only one rule: the Default role: p#tx-employee.

Question: Do you need to configure a custom role assignment rule for the Aruba-User-Role assignment?

Answer: No, the RADIUS Aruba-User-Role is automatically processed by the AP. The only reason you would create a new “Role Assignment” rule is if you had a third-party RADIUS server that returns some other attribute, such as the Filter-id. To map some RADIUS attribute to an Aruba User Role, use the Role Assignment rules. This is not required in a typical deployment: typically the RADIUS server will assign the Aruba-User-Role attribute during authentication.

13. In the Role list, click the + sign to **add a Role**.
14. For the Role name, enter **employee** and click **OK**.
15. Add another role, with name **contractor** and click **OK**.

Verify in the Role list that the 2 new roles are listed.

> **IMPORTANT:** Although you can also create the roles under Security > Roles, it is recommended to create new roles under the WLAN wizard. In deployments where a gateway is used for tunneled WLANs, the wizard ensures that roles are created in both the AP and the gateway configuration.

16. Click the **employee** role.

Question: What is the default access rule in a new role?

Answer: There is one default rule: **Allow any to all destinations**.

17. Click Save Settings.
18. Wait a few seconds for the wizard to complete, then confirm with **OK**.

Review the WLAN Configuration Access UI slider

> **IMPORTANT:** The Access type slider in the UI for the WLAN configuration will only change when a custom rule is added for the role mapping.

This does **not** mean that Aruba User roles are not processed by the WLAN!

19. In the **WLAN list**, click on the p#tx-employee WLAN.
20. Use the **pencil** icon to edit the WLAN.
21. Click on the **Access** page.

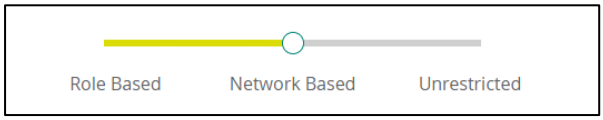

Question: What does the Access rule slider show?

Answer: The slider is set to Network Based.

Question: Does this mean that RADIUS-assigned Aruba User Roles are not applied?

Answer: No, the RADIUS VSA Aruba-User-Role is automatically accepted by the AP. There is no need to make any custom role assignment rule to use the Aruba-User-Role VSA.

Re-authenticate Clients and Verify the New Role Assignments

Simply adding the roles to the AP configuration will not reauthenticate the clients that are already connected.

You will now disconnect the two clients and then verify that they have received the new roles.

22. In Aruba Central, navigate to

Context: **Groups / campus-wifi-ui** > Navigation: **Clients** > Top: **Clients** > Right top: **List**

23. Move the mouse over the employee record. At the right end of the line the option Disconnect from AP should be shown.

24. Click on the link **Disconnect from AP** link and confirm the disconnect with **Yes**.
25. Repeat the disconnect for the **contractor** client.
26. Wait a few moments, then refresh the Clients list using the refresh button next to the CLIENTS title.

NOTE: You may need to repeat this refresh a few times until you see that the AP Role for both employee and contractor user has changed.

IMPORTANT: Some of the remote lab wireless NICs do not automatically reconnect to the WLAN after a disconnect (even when the Connect Automatically checkbox is enabled).

Please check the wireless connection status on your PC1 and PC4, you may need to manually reconnect to the employee WLAN!

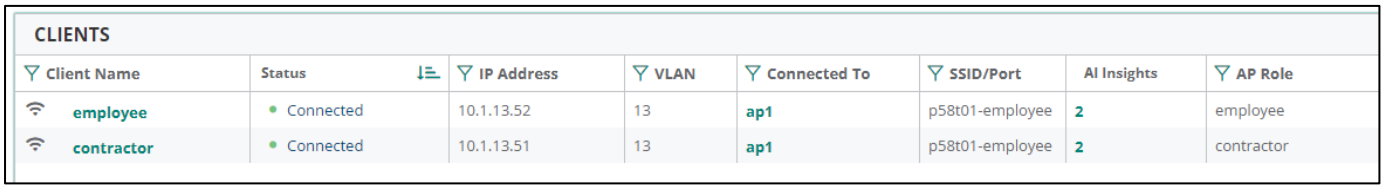

CLIENTS

Client Name	Status	IP Address	VLAN	Connected To	SSID/Port	AI Insights	AP Role
employee	Connected	10.1.13.52	13	ap1	p58t01-employee	2	employee
contractor	Connected	10.1.13.51	13	ap1	p58t01-employee	2	contractor

27. On PC1, verify the ping to 10.1.0.2 works.
28. On PC4 (contractor), verify that the ping also works, for now. This was not the desired outcome, but you will adjust this configuration in the next section.

Configuring an Alias for ACL use

The customer requested that you block several IP subnets for contractors, and they may want to block additional subnets in the future. They want a simple administration method to maintain this list of blocked IP ranges. After discussing this request with your senior colleague, you learn that this is possible using an alias.

In the AP global configuration, you will create an alias for the **contractor-blocked** subnets and use this alias for the contractor's deny rule. If the customer must block additional subnets later, they simply add them to the alias.

Define Network Alias

In this section you define an alias with blocked subnets.

29. In Aruba Central, navigate to Context: **Groups** / **campus-wifi-ui** > Navigation: **Devices** > Top: **Access Points** > Right top: **Config**
30. Click the **Security** page to open the **Aliases** section.
31. Click the + sign to add a new alias.
32. For the Name, use **contractor-blocked**.
33. Click the + sign to add an item.
34. In the **Type** dropdown list, select **Network**.
35. For Network Prefix, enter **10.1.0.0**, Network Mask **255.255.255.0**.
36. Click **OK** to add the item.
37. Click the + sign to add another item.

38. In the **Type** dropdown list, select **Network**.
39. For Network Prefix, enter **10.1.3.0**, Network Mask **255.255.255.0**.
40. Click **OK** to add the item.
41. Click **OK** to add the alias.
42. Click **Save Settings** to submit the configuration.

Modify the contractor role

You can change rules for a role in the WLAN Wizard, but also directly on the role under Security.

43. On the **Security** page, expand the **Roles** section.
44. Select the role **contractor**.
45. On the **Access Rules** table, Click the + sign to add a new rule.
46. In the **Rule Type** dropdown list, leave the default **Access Control.**
47. For **Service**, leave the default **Network – Any**.
48. For **Action**, select **Deny**.
49. For Destination, select To a network alias.
50. For Network Alias, select contractor-blocked from the list.
51. Click **Save** to add the rule.

Verify that your new rule is displayed in the Access Rules list.

Question: What is the current position for the rule?

Answer: When adding a rule directly in the role, the rule is added to the end of the ruleset.

Question: What would be the result of this order?

Answer: Rules are inspected top to bottom, and the action of the first matching rule is applied. Traffic is not blocked since all traffic matches the ‘any to all destinations’ and is allowed.

52. At the end of the **Deny** rule, use the arrow up button to move the line up.
53. Now verify in the Access Rules table that the Deny rule is now listed **before** the Allow any rule.

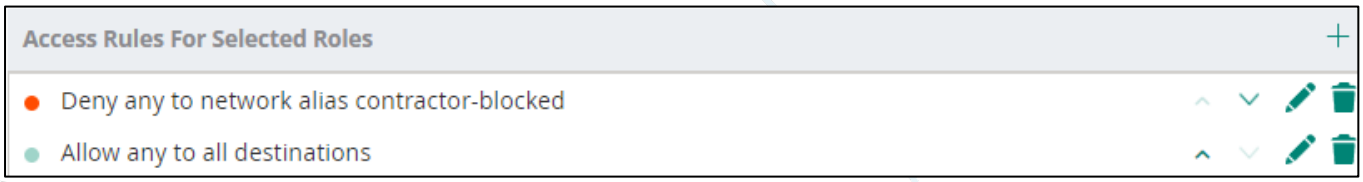

54. Click **Save Settings** to save the role.

Verify the updated contractor Role Access Control

55. In Aruba Central, in the **AI search** bar, enter **contractor** and click on the name contractor. This takes you to the contractor user detail page.
56. At the top, click on **Sessions** to see the client's current firewall sessions.

Question: Do you see any ICMP entries?

Answer: Yes, one or more ICMP sessions can be seen.

Question: What is the Action, based on the Action column?

Answer: The Deny action was applied.

Question: What is the Last refreshed time on this page?

Answer: That will be close to your current time. You should note that Central does not store firewall session information. The information on this page is retrieved on demand from the device. Click on the refresh icon next to the **Last Refreshed** time to get an updated list of firewall sessions.

57. On the PC4, verify that the ping to 10.1.0.2 no longer works.
58. Attempt to ping to 10.1.3.2. This should also fail.

Optional: Alias update

Make sure to continue the lab after these optional steps.

As you are completing and testing the configuration, the customer already has a new subnet to block for the contractors: 10.1.4.0/24. Since your configuration is now ready for this request, you can simply change the alias and add the new subnet.

This section does not include the detailed steps. Use the steps from the previous section if you need assistance. The remainder of the labs do not depend on the change in this section.

59. Using PC4, start a continuous ping to the IP address 10.1.4.2.

```
ping 10.1.4.2 -t
```

60. Add the **10.1.4.0/24** subnet to the contractors-blocked alias.
61. Using PC4, verify that you can no longer access the IP address 10.1.4.2.

This confirms the change was successfully applied using the alias.

62. Stop the ping on PC4.

End of the optional steps.

Task 3: User Role based VLAN assignment

Your customer is pleased with the contractor access control that you configured. Now they want to ensure that contractors are separated on VLAN 12. The default VLAN for the corporate WLAN should still be VLAN 11.

Objectives

- Learn about VLAN WLAN assignment.
- Configure a VLAN in a user role.
- Verify the VLAN assignment in Central.

Steps

Check the default VLAN Derivation Method

In these steps you review how the default client VLAN derivation is reported in Central.

1. In Aruba Central, enter **contractor** in the AI search bar.
2. Click the **contractor** name in the search output.
3. On the **contractor** user details page, review the NETWORK tile.

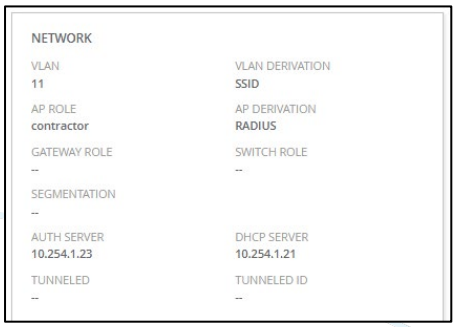

Question: What is the VLAN ID and VLAN Derivation method for the contractor?

Answer: VLAN 11, the method is SSID. This means that the client is assigned to VLAN 11 based on the WLAN's default VLAN configuration.

Configure the contractor User Role with VLAN 12

Now you will configure the contractor user role with VLAN 12. Thus, any client assigned to the contractor user role is placed in VLAN 12.

4. In Aruba Central, navigate to Context: **Groups / campus-wifi-ui** > Navigation: **Devices**> Top: **Access Points**> Right top: **Config**
5. Click the **Security** page, open the **Roles** section.
6. Select the **contractor** role.
7. Click the **+** sign to add a new Access Rule.
8. In the **Rule Type** dropdown list, select **VLAN Assignment**.

9. In the **VLAN ID** field, enter **12**.
10. Click **Save** to add the Rule.
11. Verify the new rule is in the list.

12. Click **Save Settings** to save the role.

Reauthenticate the Client

Since the AP caches the client VLAN information, a regular disconnect will not assign the client to the new VLAN. You must do a hard-disconnect on the connected AP's console - using the 'disconnect-user' command:

13. In Aruba, enter in the **AI Search bar**: **contractor** and press ENTER.
14. In the result list, click the name **contractor**.
15. In the DATAPATH tile, click the AP. This will take you to the AP device details page.
16. Click Navigation: **Tools** > Top: **Console**
17. Enter username **admin** / password **Aruba123!**
18. Click Create New Session.
19. Once the session is active, review the client list.

```
show client
```

20. Take note of the IP address for the contractor user?

_____

21. Use the disconnect-user command with the IP Address of the contractor user.

```
disconnect-user 10.1.11.x
```

22. Check the client list again until you see the contractor with an IP address in VLAN 12 (10.1.12.0/24).

```
show client
```

NOTE: If you don't see the client reconnect, switch to the PC4 windows system. The PC4 wireless client should automatically reconnect but may need a manual connect in the lab environment. If the PC4 is not connected to the wireless network anymore, make the connection manually.

23. In Aruba, enter in the **AI Search bar**: **contractor** and press ENTER.
24. In the result list, click the **name contractor** to see the client details page.
25. On the CLIENT DETAILS page, check the NETWORK tile.

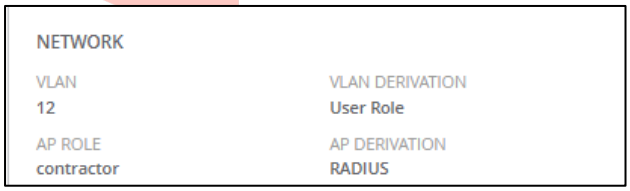

Question: What is the VLAN ID and the VLAN Derivation method for the client?

Answer: The contractor is now in VLAN 12, the VLAN derivation is now User Role. Previously, the DERIVATION was SSID based, since the user role did not have a VLAN assigned to it. The User role VLAN assignment overrules the default SSID based VLAN assignment.

> **NOTE:** It may take a few minutes before the information in Central reflects the changed access. Check the status again after a few minutes.

Task 4: AppRF and Application Statistics

Your customer has heard that the Aruba WLAN solution includes Deep Packet Inspection (DPI), an Application Recognition system that can filter and report on application use.

In this task you review DPI activation and the options to review statistics.

Objectives

- Enable AppRF on the APs.
- Use Central to monitor application usage.
- Use application-based access control.

Steps

Enable Deep Packet Inspection on the AP

1. In Aruba Central, navigate to Context: **Groups / campus-wifi-ui** > Navigation: **Devices**> Top: **Access Points**> Right top: **Config**
2. Click the **Services** page, open the **AppRF** section.
3. Set the Deep Packet Inspection value to All.

Question: What are the options under Deep Packet Inspection?

Answer: None, App, WebCC and All.
App indicates Application inspection and classification.
WebCC indicates Web Content Classification. This can be used for URL filtering and URL reputation filters.

4. Click Save Settings.

Reboot APs

5. Reboot AP1 and AP2.

Optional: Test access control

In these optional steps you will configure the contractor role to block access to the Streaming Application Category.
You will then use the PC4 (contractor) to attempt to open a streaming site.

6. In Aruba Central, navigate to Context: **Groups / campus-wifi-ui** > Navigation: **Devices**> Top: **Access Points**> Right top: **Config**
7. On the **Security** page, expand **Roles**.
8. Select the **contractor** role.
9. Add a new Access Rule

Type	Access Control
Service	Application Category: enable checkbox for Streaming
Action	Deny
Destination	To all destinations

10. Click Save.
11. Make sure the new Access rule is at the top of the list.

Test with client

12. On PC4, attempt to navigate to a streaming site for example to www.netflix.com
13. This connection should be blocked by the AP.

End of the optional steps.
Continue with the next steps.

Review Applications monitoring

The application monitoring report pages need some time to aggregate the information.

14. In Aruba Central, navigate to Context: **Groups / campus-wifi-ui** > Navigation: **Applications**> Top: **Visibility**
15. Click the **Applications** page under Visibility.

To see historical statistics, you can change how far back in time the data should be presented.

16. At the right-top, change the history period to, say 3 months, for example.

17. This reveals the applications that have been used in the requested period.

> **NOTE:** Please note that in the lab environment, the collection has only recently been started. You will probably not see the same number of applications as the example below, or the output may even be completely blank. You may continue the lab and check these screens again at a later moment.

Example:

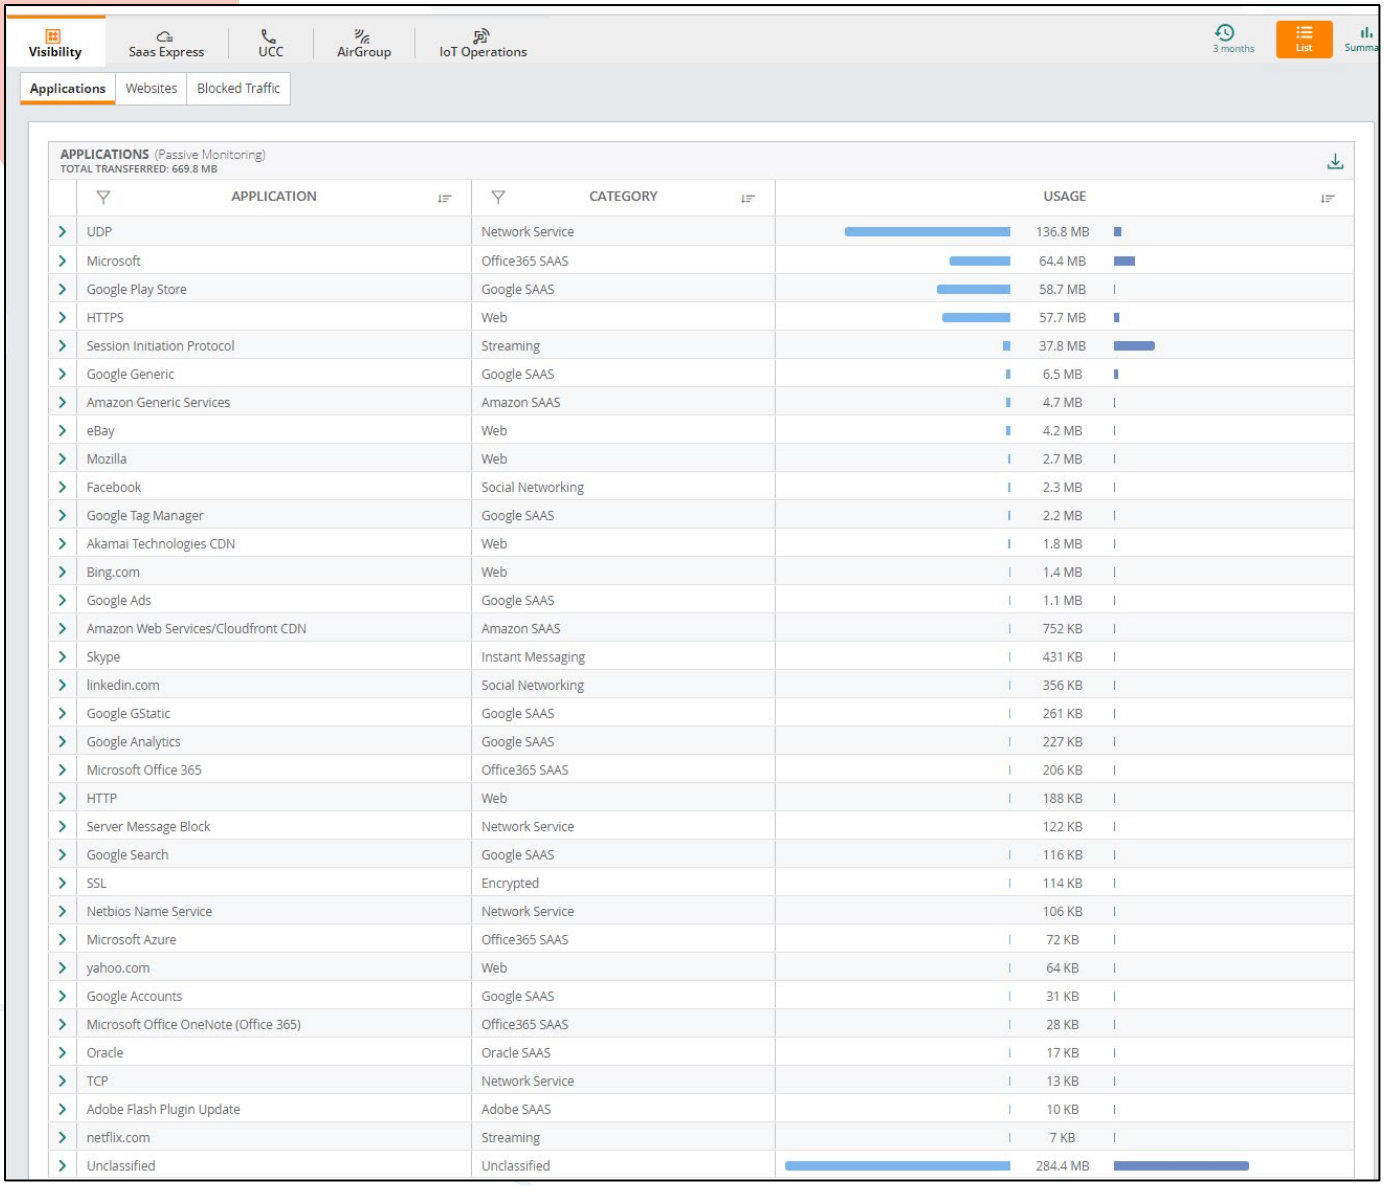

18. You can expand any application in the list to see application details.

Example:

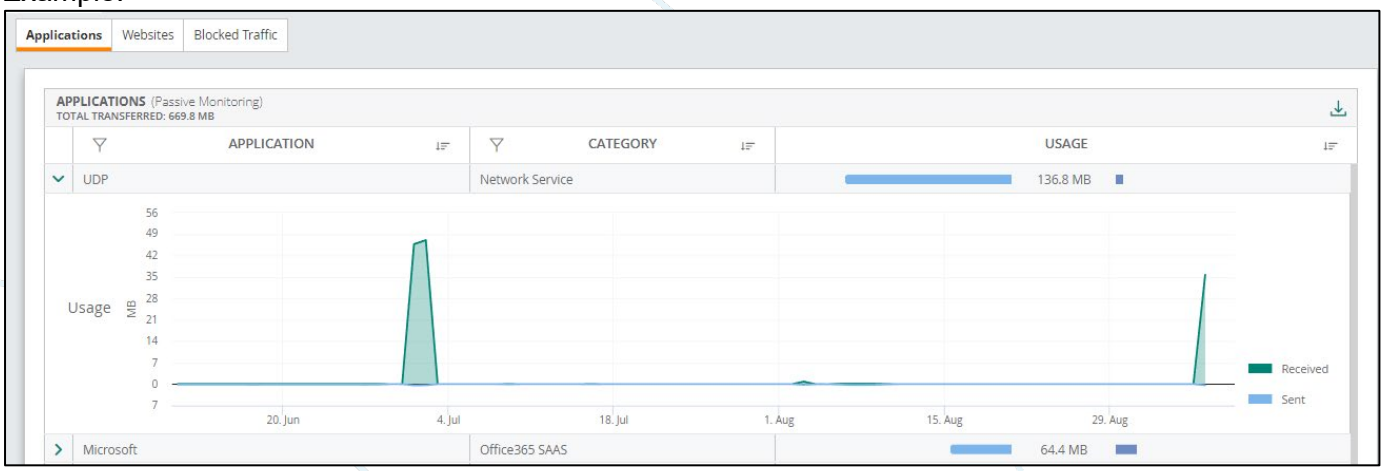

19. Click on the **Website** page under Visibility to see Web Reputation statistics for the current group and history period.

Example:

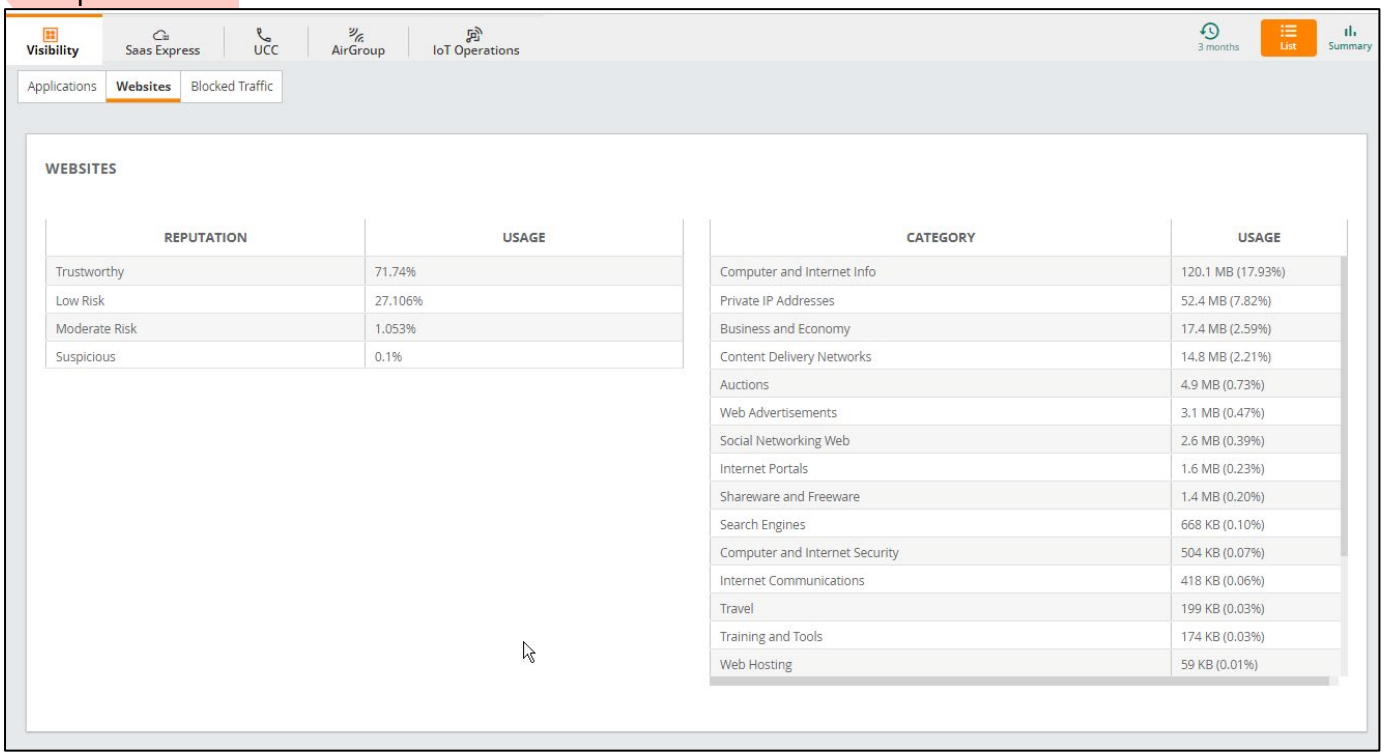

Review blocked sessions

You can collect a report about the sessions that were blocked by the firewall.

20. In Aruba Central, navigate to Context: **Groups / campus-wifi-ui** > Navigation: **Applications**> Top: **Visibility**
21. Click on Blocked Traffic.
22. If blocked traffic has been registered, you can download the list.
23. **When available**, click **Download CSV**, you can open the downloaded CSV file with Excel or a text editor.

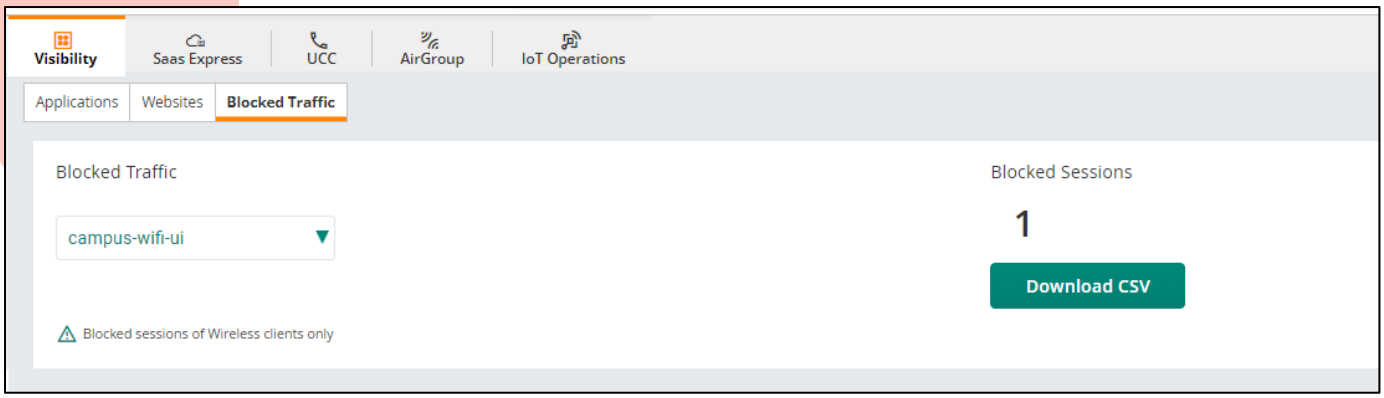

	A	B	C	D	E	F	G	H	I	J	K	L	M	N
1	Timestamp	Client Mac	UserName	Cluster	SSID	UserRole	ApSerial	DeviceType	AppName	AppCat	WebCat	WebRep	TotalBytes	
2	34:01.9	f6:00:01:58:01:04	-	87	p58t01-employee	contractor	CNKWKSM0JN	Windows	netflix.com	Streaming	Streaming Media	Trustworthy	5	
3														
4														
5														
6														

Cleanup

Disconnect the client from the wireless network and disable the employee WLAN to reduce the number of active SSIDs in the remote lab.

24. On PC1 and PC4, disconnect from the employee WLAN.
25. In Aruba Central, navigate to Context: **Groups / campus-wifi-ui** > Navigation: **Devices**> Top: **Access Points**> Right top: **Config**
26. In the WLAN list, click on the **employee** WLAN.
27. On the right end of the record, use the WLAN icon to disable the WLAN.

28. Confirm the Action with **Yes**.

    Verify in the WLAN list that the **Network Enabled** column reports **No** for the employee WLAN.

You have completed this Lab!

Lab 11.01 Implementing Guest Access

Overview

Your customer has heard that Aruba Central includes a Cloud Guest captive portal solution. They have asked if you can configure an example Guest WLAN with Cloud Guest. They want to have a basic captive portal page where they can show some legal information about the use of the guest network and guests should agree with those terms. They don't require formal username and password registration for this example.

Objectives

- Configure a splash page in Aruba Cloud Guest.
- Configure a guest WLAN using Aruba Cloud Guest.
- Verify the operation of the guest solution.

Task 1: Configure a Cloud Guest Splash Page

In this task you will prepare the webpage that guest users see when they connect to the guest WLAN. This is also known as a splash page. Aruba Central provides integrated, customizable splash pages. The Aruba Central solution also includes guest user account management and a cloud-based RADIUS system to support guest logins.

A Cloud Guest splash page could be used by AP WLANs of multiple AP groups (configurations), but by default it is only available to the AP group where it is created.

Objectives

- Configure a Cloud Guest splash page.

Steps

Configure the Cloud Guest Splash page

1. In Aruba Central, navigate to Context: **Groups / campus-wifi-ui** > Navigation: **Guests** > Top: **Guest Access** > **Splash Pages**

The list of configured splash pages is displayed.

Question: How many splash pages do you see?

Answer: Only one, the default page.

Question: What is the Type of this page?

Answer: The default page is configured as Type **Anonymous**. This means that guest users need not enter credentials to access the Guest network. The system automatically creates a user account and they automatically login with these generated credentials.

Question: What is the State of the page?

Answer: The default page is **Shared**. This means the page can be use by all groups in Central, it is not exclusive for this group. This allows you to configure and customize a splash page and use that same splash page on multiple AP groups.
If you don't want the splash page to be available on other Aruba Central configuration groups, it should not be marked as 'Shared'.

2. Click on the **+** sign to create a new Splash Page. A wizard with 3 main steps will appear:

Configuration – Customization - Localization

Configuration Page

3. On the **Configuration** page, configure these settings:
4. Name p**#**t**x**-guest-page

> **NOTE:** Make sure to replace the # value with your Pod number and x with your table number.
> For example, if you are using Table 07 in Pod 28, your page name will be
>
> p**28**t**07**-guest-page
>
> Check with your instructor if you are not sure about the Pod and Table number.

5. Type Leave the default (Anonymous)
6. Authentication Success Redirect URL
7. Redirect URL https://www.arubanetworks.com

> **NOTE:** This forces all authenticated guest to this website. By default the guest is redirected to the original requested web page**.**

8. Session Timeout for lab purposes, set it to **5 minutes**.

> **NOTE:** The session timeout means that 5 minutes after guest users have successfully connected and authenticated to the network, they will be logged out and presented the login page again. This is only set for lab testing purposes.

9. Share This Profile leave the default (enabled). This ensures the same Page and page customization can be used over multiple AP Groups.
10. Leave the other settings default.
11. Click **Next**.

Customization Page

12. On the **Customization** page, configure these settings:
13. Optional step: Feel free to change the colors and logo on this page.
14. Expand the Terms and Conditions Settings section.
15. Enter some terms. For example, enter: **Use at your own risk**.

NOTE: In the real worlds, you should involve your corporate legal department in drafting this text.

16. Set Display 'I Accept' to **Yes, display checkbox**.
17. Click **Next**.

Localization Page

18. On the Localization page, the local language translation can be configured.
19. For the **Login page title** enter Aruba Guest Access P#Tx

NOTE: Feel free to complete some of the other fields as well.

20. At the bottom of the page, click on **Preview**.

NOTE: Make sure that popups are enabled in your browser.

21. Pay attention to the URL. This is the same base URL the guest users are shown when they connect to the guest WLAN.

NOTE: If you need to make any changes, you can still go back and enter different values in the splash page configuration.

22. When you have completed the review, close the preview browser tab.
23. In the splash page wizard, click **Finish** to complete the creation of the page.
24. Verify the new splash page is listed.

Task 2: Configure WLAN profile with Cloud Guest Splash page

In this task you will configure your guest WLAN profile with the splash page that was just created.

Objectives

- Enable the guest captive portal on a WLAN.

Steps

1. In Aruba Central, navigate to Context **Groups: campus-wifi-ui** > Navigation **Devices**> Top **Access Points**> Right top: **Config**
2. In the WLAN list, move the mouse over the **p#tx-guest** and use the WLAN icon to enable it. (You created this open WLAN in Lab 09.01: WLAN Fundamentals).

> **NOTE:** You have previously disabled this WLAN to reduce the number of broadcasting WLANs in the remote lab environment.

3. Confirm with **Yes**.
4. In the WLAN list, use the pencil icon to edit the **p#tx-guest** WLAN profile.

5. You are now on the **General** page of the WLAN configuration.
6. Expand Advanced Settings.

> **NOTE:** In a production deployment you should apply the recommended best practice settings based on the latest Aruba Validated Solution Guide, such as the Broadcast/Multicast optimization.
> In this lab task you will skip this best practice configuration.

In the remote lab environment, there is much interference from neighboring APs. Therefore, you will limit this SSID to the 5Ghz band.

7. Open the **Miscellaneous** section.
8. For the **Band** option select only **5 GHz**.
9. Click Save Settings.

Guest WLAN Security and Access configuration

10. In the WLAN list, use the pencil icon to edit the **p#tx-guest** WLAN profile.
11. On the **Security** page, move the **Security Level** slider to **Visitors**.
12. Set the Type to Cloud Guest.
13. Set the **Guest Captive Portal Profile** to **p#tx-guest-page**. This is the name of the splash page you created in the previous task.

Example screenshot

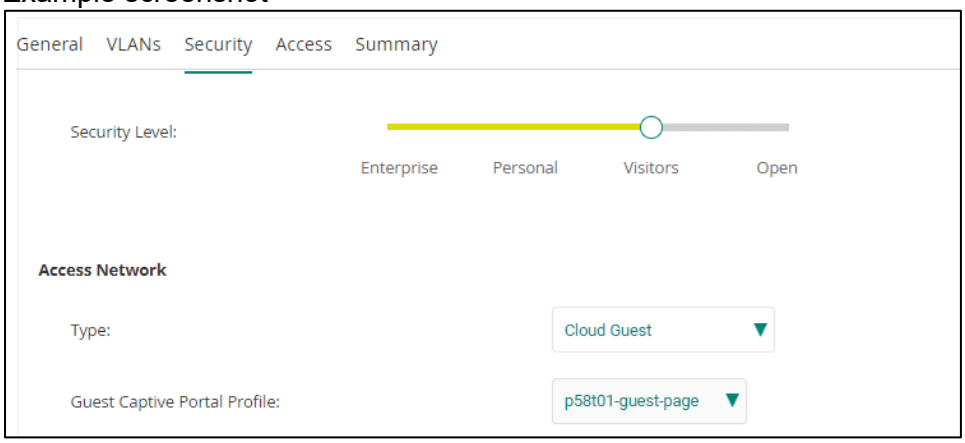

14. Click the **Access** page.
15. Move the slider to **Network Based**.

Question: What would be the effect of leaving the access rule to Unrestricted?

Answer: The guest users would have full access to internal, corporate network resources.

Question: How can you easily block access to internal resources and allow access to the rest of the internet?

Answer: Most businesses use an internal IP range from some of the subnets of RFC1918:
10.0.0.0/8
172.16.0.0/12
192.168.0.0/16
For guests, you will block access to these subnets and allow access to everything else.

16. Click **Add Rule**, deny access to the network **10.0.0.0** mask **255.0.0.0**

> **NOTE:** This is the only range you need for the lab environment. If you are unsure in your production environment, you can add the other RFC1918 ranges.

While this rule works fine, it also blocks access to the lab DNS server. In this lab environment, you should add an exception for traffic to 10.254.1.21. Also, you use this host for connectivity testing, so you should allow any traffic to this host (do not limit the rule to DNS).

17. Click **OK** to add the rule.
18. Click **Add Rule**, allow Network service **any** to a particular server **10.254.1.21**.

19. Click **OK** to add the rule.
20. Review the order of the rules.

> **NOTE:** Remember rules are processed top to bottom.
> If the rule 'Allow any to all destinations' is at the top, the guest users will have full access to the internal resources.
> Similarly, when the deny 10.0.0.0/8 is on top of the Allow any to server 10.254.1.21, all DNS traffic will be blocked.

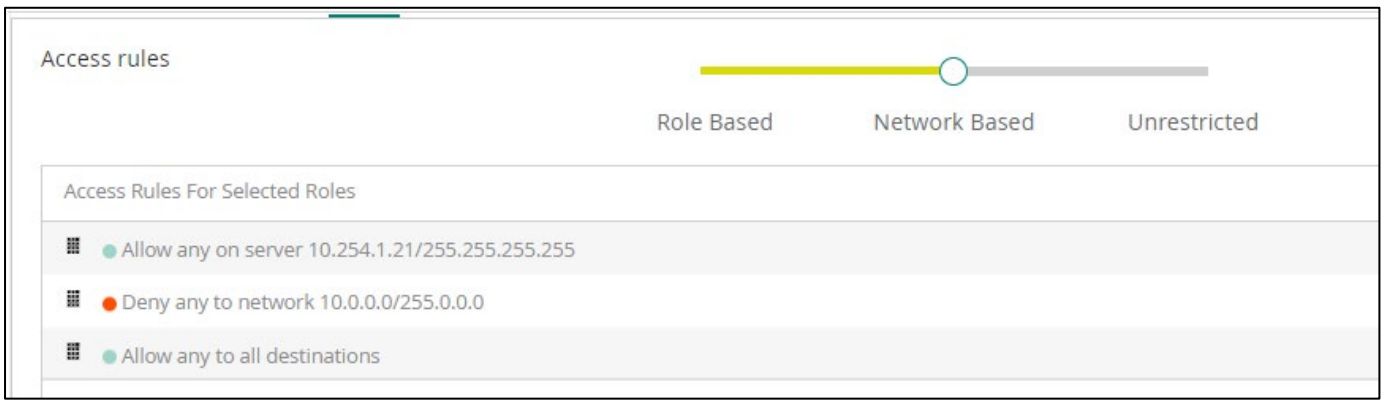

21. Click **Save Settings** to save the WLAN security and access changes.
22. After a few seconds, click **OK** to confirm the wizard has completed.

Task 3: Test Cloud Guest access

In this task you will use PC1 to make a connection to the guest WLAN that has now been configured with the Cloud Guest captive portal.

Objectives

- Test the Cloud Guest connection.

Steps

1. Open a Connection to PC1.

> **NOTE:** After performing the **next** step, a browser page will popup. Do not log in yet at this point. You will first explore the current status of the client known as the pre-auth (before the client is authenticated) stage.

2. Make a wireless connection to p#tx-guest. **Do not login** yet.
3. In Aruba Central, navigate to Context: **Groups / campus-wifi-ui** > Navigation: **Clients**

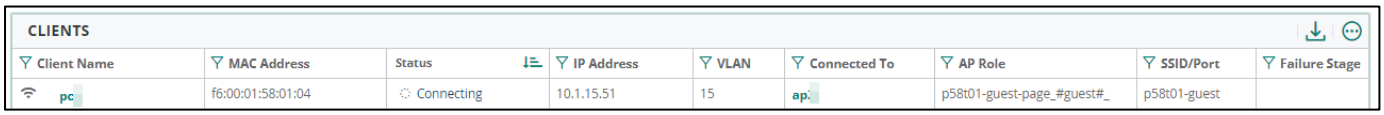

Question: What is the status of the PC1 client?

Answer: The client state is connecting. Note that the client WLAN state is already connected, but since the client has not completed the captive portal login, the status will be reported as connecting.

Question: What is currently the assigned AP Role?

Answer: The role is p#tx-guest-page_#guest#_.

Question: Did you create this role?

Answer: No, this role was automatically created when you enabled the captive portal on the WLAN profile. This role contains the captive portal redirect instructions to the splash page named p#tx-guest-page.

4. On the PC1 client, open a new tab on the browser and navigate to a secure (HTTPS) website.

```
https://www.arubanetworks.com
```

Question: What do you notice?

Answer: A certificate warning is displayed.

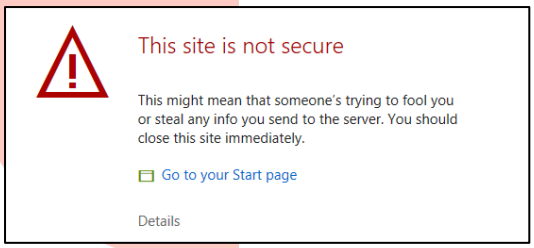

Question: Why would there be something wrong with the certificate of www.arubanetworks.com?

Answer: There is nothing wrong with the actual webserver. However, at this point, the guest user does not have full access, so the HTTPS session is captured by the AP. Since the AP does not have the www.arubanetworks.com certificate, the guest browser will show the certificate error.

Question: Can you avoid or solve the certificate warning when guest users make an initial HTTPS connection?

Answer: No, the AP can never have the original certificate of the initial HTTPS connection destination.

5. Attempt to open an HTTP connection (**not** HTTPS) to

```
http://1.1.1.1
```

or

```
http://neverssl.com
```

Question: What do you notice?

Answer: The session is now redirected to the Cloud Guest page without certificate warning. This is because the HTTP connection can be hi-jacked by the AP and redirected to the correct HTTPS website that Cloud Guest is using.

Question: Does this mean that every guest user needs to open a browser to an HTTP site manually?

Answer: Theoretically yes, however that would be challenging to explain to every guest user. This is why most OS and browser vendors have solved this by making a test HTTP connection automatically for you. In the lab example, a Windows 10 client, a test connection is automatically made by your Windows client to

```
http://www.msftconnecttest.com/connecttest.txt
```

When the test fails to download the correct test file, your windows client will report that it is not connected to the internet:

After successfully connecting to the internet, the test connection is successful, and the client reports a normal connection:

NOTE: The remote lab blocks access to microsoft.com domains, so the connectiontest.txt download file is hosted on the 10.254.1.21 host.

Review the captive portal user role

In these steps you will review the configuration that was applied to this captive portal user role. You will first try to use the UI:

6. In Aruba Central, navigate to Context: **Groups / campus-wifi-ui** > Navigation: **Devices** > Top: **Access Points** > Right top: **Config**
7. Click **Security**, expand **Roles**.

Question: Do you see the p#tx-guest-page_#guest#_ user role in the roles list?

A: No, this role is automatically created for you and it is hidden in the UI.

In the next steps you will explore this system _#guest#_ role in the AP configuration.

1. In Aruba Central, navigate to

Context: **Groups / campus-wifi-ui** > Navigation: **Tools** > Top: **Commands**

2. Select Device Type: **Access Point**, Available devices: **ap1**.
3. Select Category: **Security**, Commands: **AP Access Rule Table**.
4. **Add** the command and click **Run**.
5. Use the **maximize** button to increase the output window size.
6. In the AP1 output, check the access rule for the role p#tx-guest-page_#guest#_, this access rule will be at the end of the output.

```
...
Access Rule Name :p58t01-guest-page_#guest#_
In Use           :Yes
ACL Vlan     :0
ACL Captive Portal:external
ACL ECP Profile    :p58t01-guest-page_#guest#_
CALEA            :disable
Redirect Blocked HTTPS Traffic  :disable
DPI error page URL:
Bandwidth Limit   :downstream disable upstream disable
airslice-application-list       :
monitoring-application-list     :
```

```
Access Rules
------------
Dest IP  Dest Mask              Eth Type  Dest Match  Protocol (id:sport:eport)
Application  Action  Log  TOS  802.1P  Blacklist  App Throttle (Up:Down)  Mirror
DisScan  ClassifyMedia  TimeRange
-------  ---------              --------  ----------  ------------------------  ----
-------  ------  ---  ---  ------  ---------  --------------------  ------  -------
-------------  ---------
alias    licdn.com              IPv4/6    match       https
permit
ClassifyMedia
alias    twimg.com              IPv4/6    match       https
permit
ClassifyMedia
alias    bam.nr-data.net        IPv4/6    match       https
permit
ClassifyMedia
alias    js-agent.newrelic.com  IPv4/6    match       https
permit
ClassifyMedia
alias    symcb.com              IPv4/6    match       http
permit
ClassifyMedia
alias    symcd.com              IPv4/6    match       http
permit
ClassifyMedia
alias    digicert.com           IPv4/6    match       http
permit
ClassifyMedia
alias    any                    IPv4/6    match       http
permit
ClassifyMedia
```

Log in to the Cloud Captive Portal

Now you will login with the guest user and verify the updated access to the network.

7. On PC1, click on the '**I accept'** checkbox and **Log in** on the captive portal webpage.

After the successful authentication, the browser will redirect you to the configured welcome page (https://www.arubanetworks.com).

8. In Aruba Central, navigate to Context: **Groups / campus-wifi-ui** > Navigation: **Clients**

9. Click the Refresh button to get the latest client information.

Question: What is the Status for the client now?

Answer: After the successful captive portal authentication, the status changes to **Connected**.

Question: What is the guest client name now?

Answer: The client's name changed from the original hostname to the username. Cloud Guest with anonymous access is using username **anonymous**, this is now shown as the Client name.

Question: What is the AP Role for this user?

Answer: The AP User Role is **p#tx-guest**. After Cloud Guest returns a RADIUS accept to the AP. The AP assigns the user the SSID default role.
The pre-authentication roles were based on the splash page name. The post-authentication default role is the default SSID user role.

Phase	Role Name
pre-authentication	p#tx-guest-page_#guest#_ Splash Page name + _#guest#_
post-authentication	p#tx-guest (SSID default role)

Verify Guest Access Control

During configuration of the guest WLAN, you configured the default SSID role (p#tx-guest) with some restrictions. The goal was to ensure that the authenticated guest users cannot access the internal network resources.
You blocked access to the entire 10.0.0.0/8 network, except host 10.254.1.21 - a DNS/DHCP server that also hosts the msfconnecttest.txt.

10. On PC1, attempt to ping to 10.254.1.21. This should be successful.
11. Attempt to ping to 10.254.1.23 (the ClearPass server), this connection is blocked, along with all else in the 10.0.0.0/8 range.
12. Based on your configuration, 5 minutes after the guest user logged in on the network, the session will expire.

Optional: Review the guest session details - remaining time

Make sure to continue the lab after these optional steps!

13. In Aruba Central, navigate to Context: **Groups / campus-wifi-ui** > Navigation: **Clients**
14. Click your guest client to open the client details page.
15. On the client details page, navigate to **Tools > Commands**. Central automatically selects Device Type : Access Point and Available Devices: ap1 or ap2 for you, since the client is connected to this AP.
16. In Categories, select All Categories.

17. In the Commands window, use the filter **clients**. Thus, you only see commands that contain the string “clients”.

18. Add **show clients debug** to the Selected Commands window, then click **Run** to execute the command.

After a few moments, the command output is shown.

Tip! The output may be easier to analyze if you use the maximize button.

Question: What is the Reauth Interval for your client?

A: This should be 300 (seconds). The Reauth Interval is based on the Cloud Guest RADIUS session timeout value. You have set the Cloud Guest session timeout to 5 minutes which equals 300 seconds.

Question: What is the Reauth Age for your client?

Answer: This value is different for each client. This is the number of seconds since the client’s last authentication, but it should be lower than the Reauth interval.

NOTE: If 5 minutes has passed since you logged in as a guest, the client are logged out and it reverts to the pre-authentication role. You see this as the Authenticated value = No, and the Role is p#tx-guest-page_#guest#_ (the pre-authentication role).

In that case, make a new connection on PC1 to your guest SSID and login again. Then repeat the **show clients debug** command and check the output based on the instructions above.

End of the optional steps, continue with the next steps!

Cleanup

Disconnect the client from the wireless network and disable the Guest WLAN.

19. On PC1, disconnect from the guest WLAN.
20. In Aruba Central, navigate to

Context: **Groups / campus-wifi-ui** > Navigation: **Devices**> Top: **Access Points**> Right top: **Config**

21. In the WLAN list, click on the guest WLAN.
22. On the right end of the record, use the WLAN icon to disable the WLAN.

23. Confirm the Action with **Yes**.

Verify in the WLAN list that the **Network Enabled** column reports **No** for the guest WLAN.

You have completed the Lab!

Lab 11.02 Guest Access with ClearPass Guest

Overview

The customer is pleased with your demonstration of Cloud Guest. They are planning to complete the migration to AOS10 over the next year, at which point they will also migrate to Cloud Guest. During the migration however, they would like to have the same Guest experience as they have today on their sites with AOS8 deployments, using a ClearPass Guest solution.

You have been asked to setup an AOS10 Guest WLAN that integrates with the existing ClearPass Guest solution, using ClearPass as the external captive portal server.

Objectives

- Review the ClearPass Guest page configuration
- Configure a Guest WLAN using an external Captive Portal server.
- Configure ClearPass as an external captive portal.
- Verify access to the Guest WLAN.

Task 1: Verify a ClearPass Guest page

In this task you will review the existing guest page configuration on the ClearPass Guest server.

Objectives

- Understand the URL of a ClearPass Guest Page.
- Understand the NAS Vendor Settings on the ClearPass Guest page for a default AOS10 AP.

Steps

Review the ClearPass Guest page

1. On the MGMT PC, open a command prompt and ping to cppm.aruba-training.com

```
ping cppm.aruba-training.com
```

Question: What is the IP address for this host?

Answer: The name resolves to IP address 10.254.1.23. This is the IP address of the ClearPass server.

2. On the MGMT PC, open a browser to ClearPass to:

https://cppm.aruba-training.com/guest

> **NOTE:** The FQDN cppm.aruba-training.com is only available inside the remote lab environment.

3. Login with username **admin** / password **Aruba123!**
4. On the left side, open the section **Configuration** > **Pages** > **Self-Registrations**.

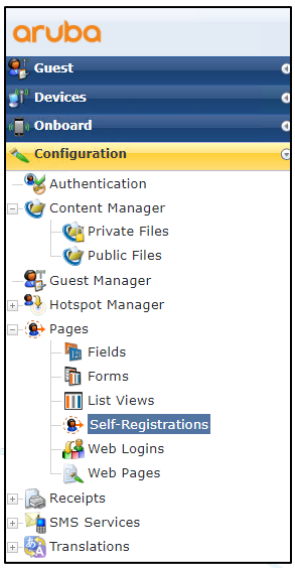

Question: What are the page names you see in the list?

Answer: There is a **Guest Self-Registration** and **acaa-wifi** page. The **acaa-wifi** was created for this training lab.

5. Click on the **acaa-wifi** entry. Some action buttons will appear under the line.
6. Click on **Launch** to see an example of the self-registration page. Pay attention to click on Launch, **not** on “Launch Self-Service” or “Launch Login”

> **NOTE:** You should not enter any credentials in the form at this point, since you are connected using the MGMT PC; you are not using a guest client PC now!

Question: What is the full URL of the page?

Answer: The full URL is https://cppm.aruba-training.com/guest/acaa-wifi.php. Take note of this URL, you will need this URL for the Guest WLAN redirect page configuration.

> **NOTE:** ClearPass is using a server certificate that was signed by a private Lab CA in this lab environment.

Question: What is the FQDN used to reach the ClearPass Guest URL?

Answer: The FQDN is cppm.aruba-training.com.

7. After the preview, you may **Close** the preview web page.

NAS Vendor Settings

On the external captive portal, the administrator needs to provide the hostname to which the guest browser will submit the guest credentials.

The AOS10 APs are by default configured with a public certificate with subject name of **securelogin.hpe.com**.

This is automatically installed and configured by Aruba Central on the AOS10 APs.
The Captive Portal administrator only needs to refer to this name in the captive portal configuration.

In the next steps you will review that this information was set on the guest captive portal page.

8. In the Self-Registrations list, click **Edit** for the **acaa-wifi** page.

9. On the customize self-registration page, click on the **NAS Vendor Settings**.

10. Review the value that has been set in the Address field.

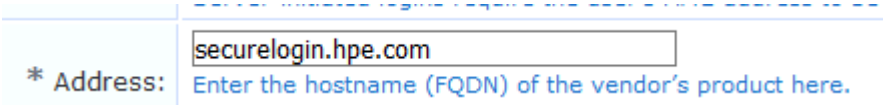

This value must match the captive portal certificate name that has been installed on the APs. In this lab setup, it matches the AOS10 default - the public certificate with subject name of securelogin.hpe.com.

This concludes the ClearPass page review. You may close the web browser on MGMT PC.

Task 2: Configure WLAN profile with ClearPass Guest Splash page

In this task you will configure a new guest WLAN profile with the ClearPass splash page that you have just reviewed.

Objectives

- Enable Guest captive portal on a WLAN.
- Configure the external splash page with a ClearPass guest page.

Steps

Create the Guest CPPM WLAN

1. In Aruba Central, navigate to Context: **Groups / campus-wifi-ui** > Navigation: **Devices**> Top: **Access Points** > Right top: **Config**
2. On the WLAN page, click **Add SSID**.
3. On the **General** page, in the **Name** field, enter **p#tx-guest-cppm**.

> **NOTE:** Make sure to replace the # value with your Pod number and x with your table number.
> For example, if you are using Table 07 in Pod 28, your WLAN name will be
>
> **p28t07-guest-cppm**
>
> Check with your instructor if you are not sure about the Pod and Table number.

4. Expand the **Advanced Settings** section. Under **Miscellaneous** set band to **5GHz**.

> **NOTE:** In a real deployment you should also apply the best practices. These steps are omitted in this task. Refer to the previous labs to see the WLAN general best practices.

5. Click **Next.**
6. On the **VLANs** page, set **Static** and VLAN id **15**.
7. Click **Next**.

Configure the External Captive Portal

In the next steps you will define the ClearPass Guest page as an external captive portal.

8. On the **Security** page, move the **Security Level** slider to **Visitors**.
9. Set the Type to External Captive Portal.
10. For the **Captive Portal Profile**, use the **+** button to add a new profile.

11. In the **External Captive Portal New** window enter these settings:
    - Name cppm-guest
    - IP or hostname cppm.aruba-training.com
    - URL /guest/acaa-wifi.php
    - Port leave default to 443
12. Click **OK** to save the profile.

Question: Why did you configure cppm.aruba-training.com as the hostname?

Answer: This is the DNS name that was registered in the lab environment for the ClearPass guest server.

Question: Why are you using /guest/acaa-wifi.php as the URL name?

Answer: You have reviewed on ClearPass guest the Splash Page. Based on the launch example, you have seen that the page name was /guest/acaa-wifi.php. You need to update this field to match the page name of the ClearPass system.

13. In the **Captive Portal Profile** field, verify **cppm-guest** is now selected.
14. For Primary server, select cppm1.

Question: When was this RADIUS server created?

Answer: You created the cppm1 RADIUS server during the Employee WLAN lab activity.

15. Open Advanced Settings > Accounting.
16. Set the Accounting to Use Authentication Servers.
17. Set the Accounting Interval to 5 min.
18. Click **Next**.

Authenticated Guest Access Control

Using the Access Control, you can control the level of network access that the guest users will have on the network. In this example, some basic restrictions will be applied.

19. On the **Access** page, move the slider to **network based**. Make sure authenticated guests can only reach the host 10.254.1.21 (for any service) and block traffic to all other 10.0.0.0/8 IP addresses.
20. Click on **Add rule** and create a rule with these settings:
    - Type Access Control
    - Services Any (default)
    - Action Deny
    - Destination network IP: 10.0.0.0 netmask: 255.0.0.0

21. Click **OK** to save the rule.
22. Click **Add rule** and create a rule with these settings:
    - Type Access Control
    - Services Any (default)
    - Action Allow (default)
    - Destination particular server 10.254.1.21.
23. Click **OK** to save the rule.
24. Click **Next** to move to the Summary page.
25. On the summary page, click **Finish**.
26. After a few seconds, click **OK** to confirm success configuration message.

Optional: Check audit trail

Make sure to continue the steps after this optional section!
Check the Audit Trail configuration changes that have been pushed to the AP.

27. In Aruba Central, navigate to

Context: **Groups / campus-wifi-ui** > Navigation: **Audit-Trail**

28. After a few moments, the new configuration will be shown. This may take a minute to show up, refresh the page until you see the new audit entry.
29. Click on the three dots to see the details for the configuration.

The WLAN access-rule is the network access role that controls the IP access for the guest users.

```
wlan access-rule p58t01-guest-cppm
 utf8
 rule 10.254.1.21 255.255.255.255 match any any any permit
 rule 10.0.0.0 255.0.0.0 match any any any deny
 rule any any match any any any permit
 exit
```

The SSID-profile controls the WLAN object, the SSID name and the VLAN assignment. It also references the external captive portal profile to ClearPass.

```
wlan ssid-profile p58t01-guest-cppm
 essid p58t01-guest-cppm
 opmode enhanced-open
 vlan 15
 rf-band 5.0
 type guest
 captive-portal external profile cppm-guest
 dtim-period 1
 broadcast-filter none
 radius-accounting
 radius-interim-accounting-interval 5
 inactivity-timeout 1000
```

```
 ax-authentication-failures 0
 blacklist
 dmo-channel-utilization-threshold 90
 max-clients-threshold 64
 enable
 utf8
 openflow-enable
 auth-server cppm1
 exit
```

The external captive portal refers to the ClearPass FQDN and the URL for the guest page.

```
wlan external-captive-portal cppm-guest
 server cppm.aruba-training.com
 url "/guest/acaa-wifi.php"
 port 443
 auto-whitelist-disable
 https
 exit
```

End of the optional steps, continue with the next steps!

Task 3: Test ClearPass Guest access

In this task you will use a wireless client to test the guest access.

Objectives

- Test the access to a Guest WLAN with external captive portal.

Steps

Make the WLAN connection

In these steps you will connect to the guest-cppm WLAN, but you will **not login** yet.

1. Open a connection to PC1.
2. Make a wireless connection to your own guest CPPM WLAN p**#**t**x**-guest-cppm.

> **NOTE:** A browser page will pop up at this point, **do not log in yet** at this point. You will first explore the status of the client during the pre-authentication phase.

3. In Aruba Central, navigate to Context: **Groups / campus-wifi-ui** > Navigation: **Clients**

Question: What is the status of the PC1 client?

Answer: The client state is connecting. Note that the client WLAN state is already connected, but since the client has not completed the captive portal login, the status in Aruba Central will be reported as **connecting**.
You may need to wait a few moments for the client status to be updated in Aruba Central.

Question: What is currently the assigned AP Role?

Answer: The role is **External CP**.

Question: Did you create this role?

Answer: No, this is an internal role that is used when an External Captive Portal (External CP) is enabled on a WLAN profile.

Log in to the ClearPass Captive Portal

In these steps you will login to the captive portal page.

4. On the PC1, complete the guest self-registration form

Full name guest
Email guest@aruba.local

5. Click on the '**I accept'** checkbox and **Register** to create a new account.

A new guest account is now created, and the details will be displayed.

6. Click on **Log In** to connect to the network.
7. You should now have access to the internet.

After the successful authentication, the browser will redirect you to the configured welcome page.

Verify

In these steps you will check the updated client status in Aruba Central after the guest login has completed.

8. In Aruba Central, navigate to

Context: **Groups / campus-wifi-ui** > Navigation: **Clients**

> **NOTE:** It may take a minute before you see the updated client information and username in Aruba Central. Use the refresh button to get the latest status.

Question: What is the guest Client Name now?

Answer: The client's name changed from the original hostname to the username that was entered during the self-registration. If you used guest@aruba.local, this will be shown as the Client name.

Question: What is the AP Role for this user?

Answer: The AP User Role is p#tx-guest-cppm. After ClearPass returns a RADIUS accept to the AP, the AP assigns the user the SSID default role.

guest@aruba.local	Connected	10.1.15.51	15	ap2	p28t11-guest-cp...	p28t11-guest-cppm

This completes the guest WLAN with ClearPass Guest configuration.

Cleanup

Disconnect the client from the wireless network and disable the Guest CPPM WLAN to reduce the number of active SSIDs in the remote lab.

9. On PC1, disconnect from the guest WLAN.
10. In Aruba Central, navigate to

Context: **Groups / campus-wifi-ui** > Navigation: **Devices** > Top: **Access Points** > Right top: **Config**

11. In the WLAN list, click on the **p#tx-guest-cppm** WLAN.
12. On the right end of the record, use the WLAN icon to disable the WLAN.

13. Confirm the Action with **Yes**.

    Verify in the WLAN list that the **Network Enabled** column reports **No** for the guest-cppm WLAN.

You have completed the Lab!

Lab 12.01 WLAN Security Features

Overview

Your customer has purchased several IoT devices that must be connected to the wireless network. The network configuration of the IoT devices is limited - the devices do not support 802.1X or WPA2/WPA3-Enterprise authentication. They only support an open or Pre-Shared Key (PSK) configuration. The customer would like to connect IoT devices to a PSK network and they have asked you to build this WLAN.

Objectives

- Configure a PSK WLAN.
- Configure access control on a PSK WLAN.
- Configure multiple keys on the PSK WLAN.
- Apply access control to different keys in an MPSK setup.

Task 1: Create PSK WLAN

In this task you will create a PSK WLAN that can be used by the IoT devices. The initial configuration in this task will use a single key for the PSK WLAN.

Objectives

- Create a PSK WLAN.
- Verify access control on the PSK WLAN.

Steps

Create the PSK WLAN

In the next steps you will create the PSK WLAN.

1. In Aruba Central, navigate to Context: **Groups / campus-wifi-ui** > Navigation: **Devices**> Top: **Access Points**> Right top: **Config**
2. On the WLAN page, click **Add SSID** to add a new WLAN profile.
3. On the **General** page, complete the **Name(SSID) p#tx-psk**

> **NOTE:** Make sure to replace the # value with your Pod number and x with your table number.
> For example, if you are using Table 07 in Pod 28, your WLAN name will be
>
> p28t07-employee
>
> Check with your instructor if you are not sure about the Pod and Table number.

4. Open the **Advanced Settings** section.

In the remote lab environment, there is much interference from neighboring APs. Therefore, you will limit this SSID to the 5Ghz band.

> **NOTE:** Many IoT devices only support the 2.4GHz band, you should not disable this in a production environment. Since you will use the lab PCs to simulate IoT devices, your lab devices can work on 5GHz only.

5. Open the **Miscellaneous** section.
6. For the **Band** option select only **5 GHz**.
7. Click **Next**.
8. On the **VLANs** page, configure **Static** VLAN **14.**
9. Make sure to **remove VLAN 1**.
10. Click **Next**.

PSK WLAN Security configuration

In the next steps you will configure the PSK Security and apply the PSK key.

11. On the **Security** page, verify the **Security Level** slider is set to **Personal**.
12. Leave the Key Management to the default: **WPA3-Personal**.
13. For the Passphrase, enter **Aruba123!**
14. Open the **Advanced Settings** section.
15. Verify that the **WPA3 Transition** is **enabled**. This is the default setting.

> **NOTE:** Many current and future IoT devices may not support WPA3 yet. To support WPA2 clients on a WPA3 WLAN, the WPA3 Transition setting must be enabled.

16. Click **Next**.

Access Control

In the next steps you will set the network access control.

17. On the **Access** page you can the IoT devices level of access.

The customer has indicated that IoT devices can only access the internet and the server subnet 10.254.1.0/24. Access to the rest of the internal network must be blocked.

18. Move the slider to **Network based**.
19. Click **Add rule** and create a rule with these settings:

Type	Access Control		
Services	Network: Any (default)		
Action	Deny		
Destination	network	IP: 10.0.0.0	netmask: **255.0.0.0**

> **NOTE:** Pay attention to configure the correct subnet mask!

20. Click **OK** to save the rule.
21. Click **Add rule** and create a rule with these settings:

Type	Access Control	
Services	Any (default)	
Action	Allow (default)	
Destination	network	IP: 10.254.1.0 netmask 255.255.255.0.

22. Click **OK** to save the rule.
23. Verify the rules have been added in the correct order.

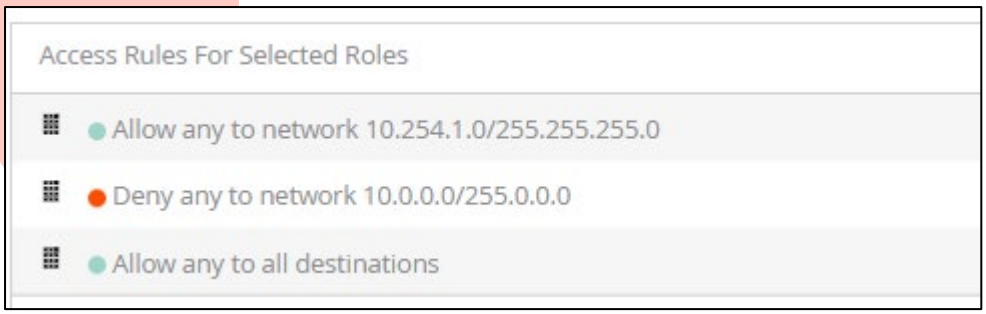

24. Click **Next** to move to the Summary page.
25. On the summary page, click **Finish**.
26. After a few seconds, click **OK** to confirm the configuration message.

Task 2: Test PSK Access with IoT air sensor

The customer is planning to deploy some IoT air sensors, and they want you to test access to the PSK network. To simulate the air sensor, you will apply a different WLAN MAC to the PC1.

Objectives

- Test access to a PSK WLAN.
- Verify PSK WLAN access control.

Steps

Verify WLAN access

In the next steps, you will change the WLAN MAC.

1. On PC1's desktop, open the ACAF student folder.
2. Double-click the file **host-wifi-pc1-iot-sensor-air.cmd**. This will set your WLAN MAC to F60**C02***P#Tx*01. (**Co2** is the MAC hint for the air sensor). If you are prompted for administrator access, click Yes.

> **NOTE:** Setting a different client MAC is not strictly required for this lab to function. It only ensures that Aruba Central sees the client as a new device (new MAC address). This ensures there will be no client history or IP address information from previous lab activities.
>
> P# will be replaced with your Pod number, for example 28.
> Tx will be replaced with your Table number, for example 01 or 12.

3. Connect to the WLAN **p#tx-psk** using the key **Aruba123!**
4. In Aruba Central, navigate to Context: **Groups / campus-wifi-ui** > Navigation: **Clients** > Top: **Clients** > Right top: **List**

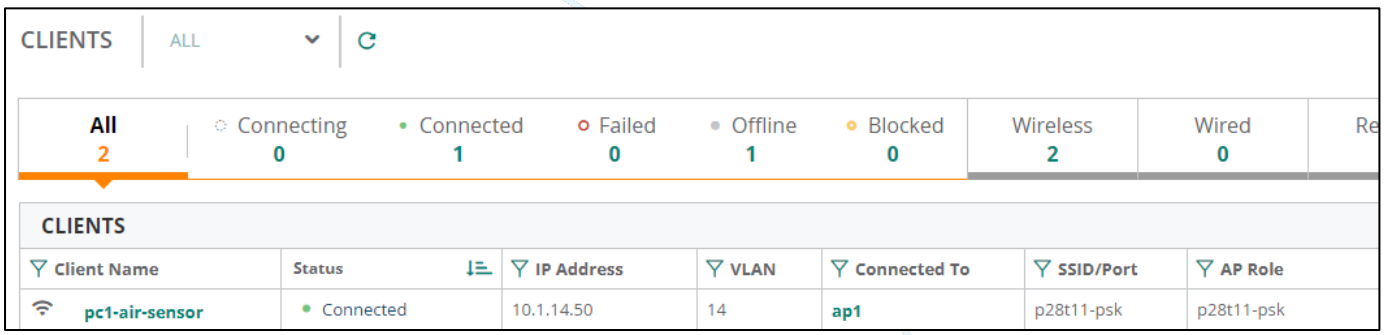

Question: Do you have a connected client on the psk SSID? (use the SSID/Port column to see the SSID)

Answer: Yes, the client (simulated air sensor) is now connected.

Question: To which VLAN is the client connected?

Answer: The client is on VLAN 14, based on the WLAN default VLAN configuration.

Verify the Access Control

The customer requested access control for the IoT devices, so you validate access control in the following steps.
IoT devices should be able to access the internet (8.8.8.8 for example), the server subnet (10.254.1.21), but no other devices on the internal network (10.1.14.1 for example).

5. On PC1, test the configured access control.
6. Open a command prompt and ping to

8.8.8.8	should be successful
10.254.1.21	should be successful
10.1.14.1	should fail

Task 3: MPSK with Differentiated Access Control

The customer was happy with the easy access control configuration for IoT air sensors. Now they want a new set of HVAC air-conditioning (AC) devices that must connect to the WLAN.

They want the AC units to connect with a PSK (**Airco123!**) that is different from the air sensors. The AC units also have different access control needs. They should only be able to access the AC unit management software on the server 10.254.1.21. They also want to make sure that anyone using the default PSK key Aruba123! can only access the internet.

After consulting with your senior colleague, you learn about the MPSK local feature on the APs. This feature allows you to configure multiple Pre-Shared Keys (MPSK) for the same WLAN. The feature also enables different user roles (access control) for the different keys.

The MPSK configuration will consist of the following steps:

1. Configure User Roles for *iot-air-sensor* and *iot-ac* with unique access control.
2. Configure an MPSK (multiple Pre Shared Key) rule set to link unique PSK keys to the user roles.
3. Change the default SSID role Access Control to internet only.

Objectives

- Configure MPSK local on a PSK WLAN.
- Configure unique access control with MPSK.

Steps

Update the default Access

In these steps you will limit access for the default user role on the PSK WLAN.

1. In Aruba Central, navigate to Context: **Groups / campus-wifi-ui** > Navigation: **Devices**> Top: **Access Points**> Right top: **Config**
2. Select the **p#tx-psk** WLAN, use the pencil icon to **edit**.
3. Click the **Access** page to see the Network Based access.
4. Move the slider to **Role Based**.

Question: What role is automatically selected when you moved the slider?

A: This is the default SSID role, based on the WLAN profile name: **p#tx-psk**.

Question: What is the current rule set for this role?

Answer: Based on the requirements in the previous task, the default role allows:
access to the 10.254.1.0/24 network
blocks all other access to the 10.0.0.0/8 network.

Update the rule set so only internet access is allowed. Remember that access to the host 10.254.1.21 is always required in the lab for several network services such as DNS and DHCP.

5. Select the rule **Allow any to network 10.254.1.0/255.255.255.0**, use the pencil **edit** the rule.
6. Set the Destination to **particular server 10.254.1.21**.

   Click **OK** to save the rule.

Create new user roles for iot-sensor-air

The air sensor will now get its own user role. You will apply the same access control as that previously applied to the default SSID:

- Allow Access to the server subnet 10.254.1.0/24
- Deny Access to internal network 10.0.0.0/8
- Allow Access to Internet

7. Click the **Add Role** button to create a new role.
8. For the name, enter **iot-sensor-air** and click **OK**.

Verify that the new user role iot-sensor-air is now selected.

9. Click **Add rule** and create a rule with these settings:

Type	Access Control		
Services	Any (default)		
Action	Deny		
Destination	network	IP: 10.0.0.0	netmask: 255.0.0.0

10. Click **OK** to save the rule.
11. Click **Add rule** and create a rule with these settings:

Type	Access Control	
Services	Any (default)	
Action	Allow (default)	
Destination	network	IP: 10.254.1.0 netmask 255.255.255.0.

12. Click **OK** to save the rule.

Verify the rules for iot-sensor-air are listed in the correct order. If necessary, you can adjust the order by dragging the rules up or down in the list using the dots at the start of the line:

Drag: 

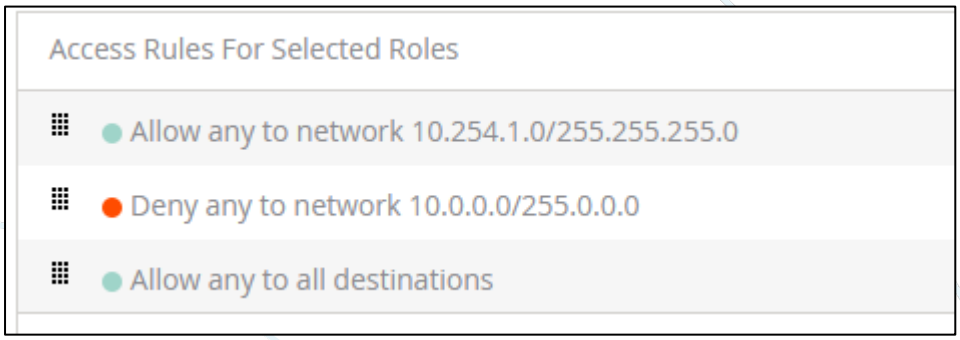

Create new user role for iot-ac Units

The HVAC AC units will get their own user role. You will apply these access controls based on the customer requirements:

- Allow Access to server 10.254.1.21.
- Deny all other access.

13. Click the **Add Role** button to create a new role.
14. For the name, enter **iot-ac** and click **OK**.

Verify that the new user role iot-sensor-air is now selected.

15. Click **Add rule** and create a rule with these settings:

Type	Access Control
Services	Any (default)
Action	Allow (default)
Destination	particular server 10.254.1.21

16. Click **OK** to save the rule.
17. Click the rule **Allow any to all destinations**.
18. Use the bin icon to **delete** the entry.

19. Confirm the action with **yes**.
20. Verify the rule of the role.

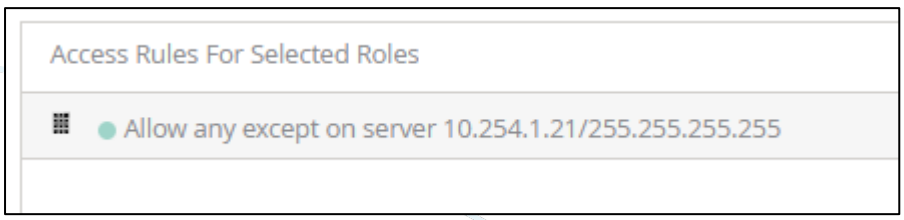

21. Click **Save Settings** to save the Access Control changes.

> **NOTE:** You have created the new roles after moving the slider to 'Role based', but you did not create any custom role-derivation rules. Therefore, if you would edit the WLAN again, the access control slider would still show 'Network Based' instead of Role Based, this is expected.

Create MPSK Local rule set

In the previous steps you have prepared dedicated roles for the AC and the air-sensors. Now you need to prepare an MPSK Local rule set where unique PSK keys can be added and linked to unique user roles.

22. In Aruba Central, navigate to Context: **Groups / campus-wifi-ui** > Navigation: **Devices**> Top: **Access Points**> Right top: **Config**
23. Click on the **Security** page, open the **MPSK Local** section.
24. Click on the **+** button to add a new rule set.
25. For Name enter **psk**.
26. In the key list, use the **+** button to add a key with these settings:

Name	**iot-ac**
Passphrase	**Airco123!** (then retype the Passphrase) IMPORTANT: Make sure you **don't configure Aruba123!**
Role	**iot-ac**

27. Click **OK** to add the key.
28. In the key list, use the **+** button to add another key with these settings:

Name	**iot-sensor-air**
Passphrase	**Sensorair123!**
Retype the Passphrase.	
Role	**iot-sensor-air**

29. Click **OK** to add the key.

Now you have added two keys that are specifically linked to a user role. The customer also wants to support the existing PSK, which can be linked to the SSID default role. Thus, you need not specify a specific role.

30. In the key list, use the **+** button to add a key that is not linked to a user role:

Name	**default**
Passphrase	**Aruba123!**
Role	<do not select a role>

31. Click **OK** to add the key.

> **NOTE:** You could also have selected the **p#tx-psk** user role since that is the default user role for the PSK WLAN.

32. Click **OK** to add the MPSK local rule set.
33. Click **Save Settings** to submit the changes.

You have now created a MPSK local rule set in which different keys will be linked to different user roles.

Update PSK WLAN Profile with the MPSK Local rule set

In the previous steps you have created a rule set. The rule set doesn't do anything by itself, it must be enabled on a WLAN profile.
You will now enable it on your PSK WLAN Profile.

34. In Aruba Central, navigate to

Context: **Groups / campus-wifi-ui** > Navigation: **Devices**> Top: **Access Points**> Right top: **Config**

35. On the WLANs page, select the **p#tx-psk** WLAN, use the pencil to **edit** the WLAN.
36. Click on the **Security** page.
37. For **Key Management** select **MPSK Local** from the list.

Question: The list also contains MPSK AES, what is the difference with MPSK Local?

Answer: MPSK-AES requires the use of ClearPass RADIUS server to store and manage the PSKs. It provides per user or per device keys to be generated. In this lab you will only configure MPSK Local setup.

38. In the **MPSK Local** field, select **psk**. This is the list you have prepared in the previous steps.
39. Click **Save Settings** to submit the configuration.
40. After a few seconds, click **OK** to confirm the success message.

Task 4: Test MPSK Access

In this task you test the PSKs defined in the MPSK Local rule set. You verify that the correct role was assigned for each key.

Objectives

- Test access to an MPSK WLAN.
- Verify MPSK WLAN role assignment.

Steps

Verify WLAN access for AC Unit

1. On PC1, forget the existing WLAN connection

Click on the Wireless icon in the status bar.
Then right-click on the active connection to see the popup menu.
In the menu, select **Forget**.

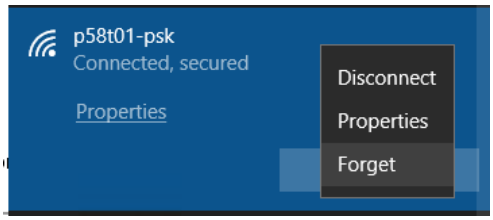

2. Connect again to your psk WLAN, this time use the key **Airco123!**
3. In Aruba Central, navigate to Context: **Groups / campus-wifi-ui** > Navigation: **Clients** > Top: **Clients** > Right top: **List**

Question: What AP role is assigned to the client?

Answer: Based on the PSK, the client has received the AP Role **iot-ac**

CLIENTS

Client Name	Status	IP Address	VLAN	Connected To	SSID/Port	AP Role
pc1-air-sensor	Connected	10.1.14.50	14	ap1	p28t11-psk	iot-ac

NOTE: It may take a moment for Central to reflect the role update.

Verify Access control for the AC Units

The AC units are only allowed to access the host 10.254.1.21.

4. On the PC1, attempt to ping to these destinations:

10.254.1.21 should be successful
10.254.1.23 should fail

8.8.8.8 should fail

Verify WLAN access for IoT Sensor Air

In these steps you will test the MPSK access with the IoT Air Sensor PSK.

5. On PC1, forget the existing PSK WLAN connection.
6. Connect to your psk WLAN with the key **Sensorair123!**
7. In Aruba Central, navigate to Context: **Groups / campus-wifi-ui** > Navigation: **Clients** > Top: **Clients** > Right top: **List**

Question: What AP role is assigned to the client?

Answer: Based on the PSK, the client has received the AP Role **iot-sensor-air**.

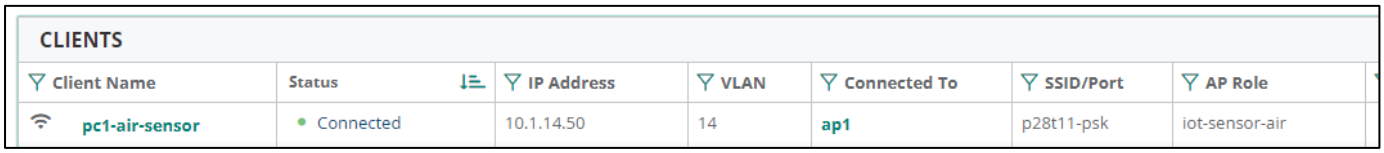

NOTE: It may take a minute before the status of the client is updated in the Central list.

Verify Access control for the IoT Sensor Air

The air sensors are allowed to access the server subnet 10.254.1.0/24 and the internet. Other internal networks should be blocked.

8. On the PC1, attempt to ping to these destinations:

10.254.1.21 should be successful
10.254.1.23 should be successful
8.8.8.8 should be successful
10.1.14.1 should fail

Verify WLAN access with the default key

You have configured two additional keys for the MPSK WLAN, but clients that have been configured with the original Aruba123! should still be able to connect to the PSK. In the next steps you will verify that clients can still use the original PSK to connect to the WLAN.

9. On PC1, forget the existing WLAN connection.
10. Connect to your psk WLAN with the key **Aruba123!**
11. In Aruba Central, navigate to

Context groups: campus-wifi-ui > Navigation Clients > Top Clients > Right top: List

Question: What is the AP role that is assigned to the client?

Answer: Based on the PSK, no role was configured for this key. Therefore, the client has received the SSID default Role **p#tx-psk**.

CLIENTS

Client Name	Status	IP Address	VLAN	Connected To	SSID/Port	AP Role
pc1-air-sensor	Connected	10.1.14.50	14	ap1	p28t11-psk	p28t11-psk

Optional: Verify Access control for the default key

Anyone connecting with the default key Aruba123! should only be allowed internet access.

12. On the PC1, attempt to ping to these destinations:

10.254.1.21 should be successful based on the allowed host entry.
10.254.1.23 should fail
10.1.14.1 should fail
8.8.8.8 should be successful

Task 5: Troubleshoot clients connecting to PSK

In this task you will enter an invalid PSK and see how to troubleshoot this using Aruba Central.

Objectives

- Troubleshoot a PSK Key mismatch.

Steps

Troubleshoot PSK Key mismatch

In these steps you will attempt to make a WLAN connection with an invalid PSK.
Using the client details and Live Events in Aruba Central you will see the issues.

1. On PC1, forget the existing WLAN connection.
2. Attempt to connect to your psk WLAN with an invalid key, for example **test123!**
3. In Aruba Central, navigate to Context: **Groups / campus-wifi-ui** > Navigation: **Clients** > Top: **Clients** > Right top: **List**

Example:

Client Name	MAC Address	Status	IP Address	VLAN	Connected To	AP Role	SSID/Port	Failure Stage
pc1	f6:0c:02:58:01:01	Failed			ap1		p58t01-psk	KEY EXCHANGE

4. Click on PC1 to access the client details page.

Question: The **DATA PATH** window is shown at the top of the page. What is the **Failure Stage** for the client?

Answer: The Failure Stage is KEYEXCHANGE.

> **NOTE:** If the client is not shown as failed, you can repeat entering the invalid key on the client.

5. In the Navigation pane, click on **Live Events**. The troubleshooting will automatically start for a maximum duration of 15 minutes.
6. On PC1, attempt to connect again with an invalid key, for example **test123!**
7. In Aruba Central, check the event descriptions.

TIP: If the events are difficult to read, you can download the events as a CSV file and open them on your local system with a text viewer or spreadsheet application such as Excel.

The download button at the top of the event list: -

Question: What is the result message for the Client exchanged key events?

Answer: The reported result is MIC failure. This indicates an invalid PSK.

Cleanup

Disconnect the client from the wireless network and disable the **p#tx-psk** WLAN.

8. On PC1, disconnect from the psk WLAN.
9. In Aruba Central, navigate to

Context: **Groups / campus-wifi-ui** > Navigation: **Devices**> Top: **Access Points**> Right top: **Config**

10. In the **WLAN** list, click on the **psk** WLAN.
11. On the right end of the record, use the WLAN icon to **disable** the WLAN.

12. Confirm the Action with **Yes**.

    Verify in the WLAN list that the **Network Enabled** column reports **No** for the psk WLAN.

You have completed the Lab!

Lab 13.01 Monitoring and Maintenance

Overview

Your customer has asked you how they can monitor the network and perform maintenance on the APs and switches of their Aruba solution. In this lab you will see how you can use Aruba Central to configure alerts and generate reports. You will also explore how firmware compliance can be configured for the different device types in your setup (APs and switches).

Objectives

- Configuring alerts
- Configuring reports
- Configuring firmware compliance

Task 1: Configuring Alerts

Your customer wants to use Aruba Central alerts to help with network monitoring, such as when an AP goes offline. After contact with your senior colleague, you know that you can use Aruba Central to alert on several common network and wireless deployment issues.

In this task you will configure several alerts:
- When AP is offline for more than 5 minutes
- When a configuration change is made in Aruba Central

The customer would like to see the alert in Aruba Central, as well as receive an email about the alert.

Objectives

- Configure alerts in Aruba Central.
- Configure default recipients for alerts.
- Review the generated alerts.

Steps

Configure the Alerts

In these steps you will be configuring the following alert definitions in Aruba Central:

AP

AP disconnected:	major	over 5 min

Switch

Switch disconnected:	major	over 5 min

Audit

Config change detected	minor

1. In Aruba Central, navigate to Context: **Global** > Navigation: **Alerts & Events** > Top: **Alerts** > Right top: **Config**

Set Default Recipients

You can set a default email address. This makes it easy to set the email destination when creating alert definitions.

2. At the top-right, click **Default Recipients**.
3. Use the **+** button to add a new recipient.
4. **Optional**: You may enter your own email address here to test the alerts. If you don't want to use your own email address, you can leave the field blank.

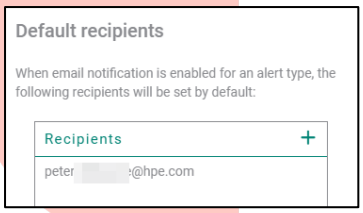

5. Click **Save**.

Set Access Point Alerts

6. Several categories are used to group the alert definitions.

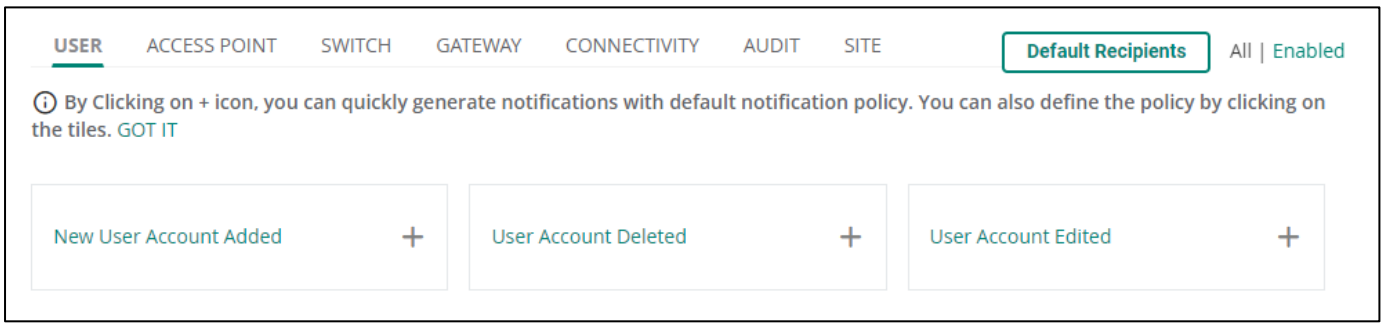

Question: What are the categories for the Alerts?

A: User / Access Point /Switch / Gateway / Connectivity / Audit / Site

7. Click the **Access Point** category to see the AP Alert options.

Access Point Disconnected

8. Click **AP Disconnected** text.

> **NOTE:** Do not click the + sign next to the text, this will only enable the alert with the default settings.

The severity level controls how alert severity is set for this alert.

9. Set Severity to Major.
10. Set **Duration** to **5** mins. This is the minimum value.
11. There is no need to set the device filter options.
12. **Optional**: If you have set your email, **Enable** the **Email** checkbox. The Default Recipients are automatically enabled. You could optionally add another email address for this specific alert.
13. Click **Save**, wait a few seconds, and confirm the message.
14. Verify that the alert is enabled, there should be a check mark next to the alert.

15. Repeat the previous steps to enable the **Virtual Controller Disconnected** alert with the same settings.

Review configured Alerts

You can adjust the view so you only see alerts that are enabled.

16. At the top-right, click the **Enabled** link to filter the Alerts.

17. Now only the alerts that you have enabled are displayed.

Set Switch Alert

18. Click the **Switch** category to see switch alert options.
19. **Enable** the Switch Disconnected alert.

Severity	Major
Duration	5 min
Email	enabled (optional, if you have configured your email)

Set Audit Alert

To receive an alert when a device configuration changes, you can enable the audit alert.

20. Click the **Audit** category to see the switch alert options.
21. Enable the Config Change Detected alert.
22. Severity **Critical** / email enabled
23. **Save** and confirm the message.

Generate an action that triggers an alert

You will now make a very basic configuration change to verify the Audit Alert operation.

24. In Aruba Central, navigate to Context: **Groups / campus-wifi-ui**> Navigation: **Devices** > Top: **Access Points** > Right top: **Config**
25. In the WLAN list, enable or disable the guest WLAN by clicking the Wi-Fi icon on the right-end side of the record.

Review the Alerts

In the next steps you review where you can find alerts.

26. In Aruba Central, navigate to Context: **Global** > Navigation: **Alerts & Events** > Top: **Alerts**
27. Review the list of **Open Alerts**.

Question: What is the category for the last alert?

Answer: Config change detected.

28. Hover your mouse over the description field to see alert details.

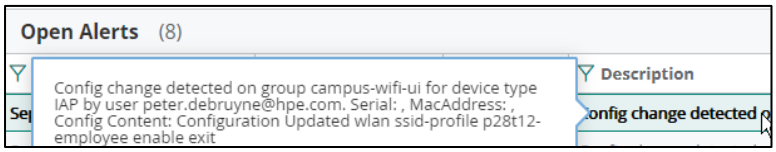

The **Open Alerts** list reflects a list of 'active' alerts. With an alert acknowledge action, the alert will be removed from the Open list.

29. Select the alert (if there are any other alerts, you may select them as well).
30. At the bottom of the page, click **Acknowledge**.

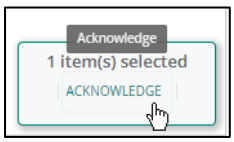

31. Verify the Open Alert list is now empty.

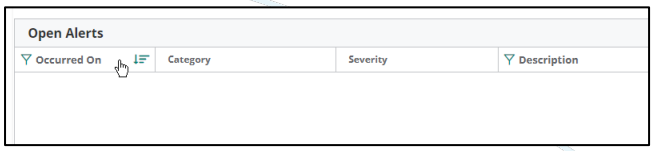

Reviewing historical Alerts

If you need to review an Acknowledged Alert, you can still show all alerts.

32. Click Show Acknowledged Alerts.

You can enable some extra columns to see who acknowledged the alert and the alert's elapsed time. The elapsed time is the time between the alert start time and the acknowledge time.

33. Click the **3 dots** to see columns display options.
34. Enable the Acknowledged By and Elapsed time columns.

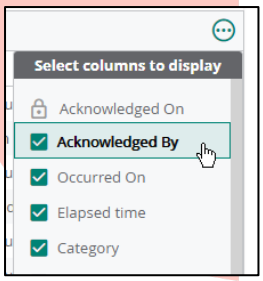

35. Review the last alert about the config change.

Example screenshot:

Acknowledged On	Acknowledged By	Occurred On	Elapsed time	Category	Severity	Description
Sep 30, 2022, 13:50:03	peter.debruyne@hpe.com	Sep 30, 2022, 13:44:47	5 minutes 16 seconds	Config change detec...	critical	Config change detected on group campus-wifi-ui for device type IAP by us

Question: Who acknowledged the alert?

Answer: Your login name should be listed.

Question: What is the elapsed time?

Answer: Depending on your lab progress, this will be time between the alert detection and the Acknowledge action.

36. **Optional**: If you have enabled your email in the default recipients, check your email for the new alert.

Task 2: Generating Reports in Aruba Central

In this task you will use Aruba Central to generate a report about the network RF Health.

Objectives

- Use Aruba Central to generate a report.

Steps

Generate a Report

1. In Aruba Central, navigate to Context: **Global** > Navigation: **Reports** > Right-top: **Summary**
2. Click **Create**.
3. Review report categories.
4. Under Infrastructure, click on **RF Health**.
5. Click **Next**.
6. Select the context and site.

Context	Site
Site	site-campus-main

7. Click **Next**.
8. Report Period: **Last 7 days** and click **Next**.
9. Leave recurrence to the default of **One Time (Now)**.
10. Set report title to **Campus-RF-Usage**.
11. **Optional**: Enter your email address, email format: enable PDF.
12. Click **Generate**, confirm the message.

Review the report

13. In Aruba Central, navigate to Context: **Global** > Navigation: **Reports** > Right-top: **List**

> **NOTE:** It may take a few minutes to complete the report generation, repeat checking the list until you see your report.

This will show the list of generated reports.

14. Click the entry with title **Campus-RF-Usage**.
15. Review the report. It contains details for the different bands, channel changes, channel use, errors etc.
16. At the top of the page, you can download the report as PDF.

17. Review the RF Health in the downloaded PDF document.

Task 3: Firmware Compliance

Your customer is thinking about the required firmware on the devices during their deployment. They will be deploying factory default devices with Aruba Central, but when the devices are unboxed, they may be running a different firmware version compared to the validated production version.

The customer knows they can manually update the firmware to the validated version, but they have asked you to simplify this operation. After some research, you learn that Aruba Central has a firmware compliance feature. This allows you to set a desired version on a global or group level. When a new device is added to that context, Aruba Central automatically updates the device to the desired version. This may be an upgrade or downgrade of the firmware version.

The compliance can be set for all groups or for specific groups, it is not possible to set the compliance at a site level. The logic for this is that one group represents one configuration, and a configuration may have features that depend on a specific firmware version. Within a site, you may have multiple device configurations, represented by multiple groups.

Objectives

- Set firmware compliance.
- Verify firmware compliance configuration.

Steps

Configure compliance on a new group

> **IMPORTANT:** You will **not** change the actual firmware version on the remote lab devices in the lab. In this task you will create a new group and set the compliance on this new group as an example of the steps. However, you will not move the devices to this new group.

The customer would like to test a new configuration and firmware version on a subset of their devices.
Since they don't want to change the existing production configuration, you should not make changes to the existing group. Therefore, you will first create a copy of the existing group. On this new group you can then set the new firmware compliance version.

> **NOTE:** Copying a group is fine for switches, gateways and APs operating only in bridge mode, but must be handled with care for APs with tunneled WLANs.

1. In Aruba Central, navigate to Context: **Global**> Navigation: **Organization**> Top: **Network Structure**> Tile: **Groups**
2. In the list of groups, look for the group **campus-sw-edge-tpl**.
3. Move your mouse over the group to see the action icons.

4. Click the **Clone Group** icon to make a copy of the group.

5. In the Clone Group dialog, for the name field, enter **campus-sw-edge-tpl-v2**.
6. Click **Clone** to copy the group.
7. Verify the new group is now displayed in the group list.

Verify the Cloned Configuration

8. In Aruba Central, navigate to Context: **Groups / campus-sw-edge-tpl-v2** > Navigation: **Devices**> Top: **Switches**> Right top: **Config**

Question: Do you see a template in the Templates list?

Answer: Yes, the sw-edge template is available on this new group. It was copied from the original group.

9. At the top, click the **Variables** page.

Question: Are there any variables listed here?

Answer: No, since there are no devices assigned to the group yet, no variables will be displayed.

Question: Will this cause any issues if you move a device from another group? (Hypothetical question since you will not move your devices!).

Answer: No, Aruba Central stores the device variables at the device level. The variables will not change when the device is moved to another group.

Set the firmware compliance at group level

With the new group defined, you are now ready to set the desired firmware version.
Compliance can be set at the global and group level.
You will now configure the compliance at the group level.

10. In Aruba Central, navigate to Context: **Groups / campus-sw-edge-tpl-v2** > Navigation: **Firmware**> Top: **Switches**

> **NOTE:** Make sure to access the correct device type (Access Points/Switches/Gateways) at the top of the screen!

11. At the top-right, click **Set Compliance**. The Manage Firmware compliance window will be displayed.

NOTE: The Compliance window reflects the context in which you are currently working. Since you navigated to this screen via the Group context, this screen shows that no compliance is currently set.

12. Move the slider for **Set firmware compliance** to enabled.

NOTE: Compliance may have been enabled already at the MSP level in Aruba Central, you can continue the steps even when it was enabled already.

13. Leave the field **AOS-S Version** default.

NOTE: There are no AOS-S devices in your lab, so this setting has no effect in this lab.

14. Set **CX Firmware Version** to 10.09.1050.

NOTE: This is just an example version, any version is fine since you will not actually update the devices.

15. Leave the field **When** default (**Now**).

NOTE: This can be useful to control when the new compliance policy shall take effect. This should typically be set to a time in a maintenance window.

16. Leave Install on set to primary partition.
17. Enable the **Automatically reboo**t checkbox.

NOTE: Make sure to enable this checkbox, otherwise Aruba Central will only update the firmware in the flash, but you must still decide when to reboot the devices. This may be convenient if you want to control the order of rebooting devices.

18. Click **Save** to submit the configuration.

You have now created a new group and configured firmware compliance for that group. Any AOS-CX switch you might add to this group is automatically updated to the configured compliance version. There is no need to do this in this lab environment.

Access Point compliance

You configure Access Point firmware compliance like the previous steps.

19. For example, in Aruba Central, navigate to Context: **Groups / campus-wifi-ui** > Navigation: **Firmware**> Top: **Access Points**

20. At the top-right, click **Set Compliance**. You see the Manage Firmware compliance window with the available AP firmware version.
21. Click **Cancel**.

> **IMPORTANT:** You should **not** set the compliance on the APs in the lab.

You have completed the Lab!

Lab 14.01 Troubleshooting Overview

Overview

Your senior colleague has asked you to become more familiar with some troubleshooting tasks and scenarios. In this lab you will collect Tech Support log file information from the Aruba APs and Switches. These files can be used by your senior colleague or Aruba TAC to analyze the system.

You will troubleshoot the connection between a switch and Aruba Central and explore Aruba Central's site-level status windows.

Objectives

- Collect log information for senior support or TAC.
- Troubleshoot the connection between a switch and Aruba Central.
- Review site-level troubleshooting information.

Task 1: Collection Log Information for Senior Support or TAC

In this task you will learn to collect tech support information from the Aruba APs and switches.

Objectives

- Collect tech support information for an AP
- Collect tech support information for a switch

Steps

Access Point Tech Support Collection

In the next steps you will collect the support information for an AP.

1. In Aruba Central, navigate to Context: **Global** > Navigation: **Devices** > Top: **Access Points**
2. Click **ap1** to view its device page.

> **NOTE:** You can also use the AI Search bar:
> Enter **ap1**, press <ENTER> and click on the ap name in the search result window.

You are now at this location in Aruba Central:
Context: **Device / ap1** > Navigation: **Overview** > Top: **Summary**

3. At the top-right, click the Actions dropdown list and select **Tech Support**.

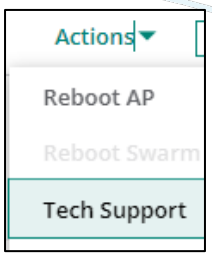

This action will take you to the Tools > Commands page.

Question: What is the currently selected Command?

Answer: The command **AP Tech Support Dump** was automatically selected.

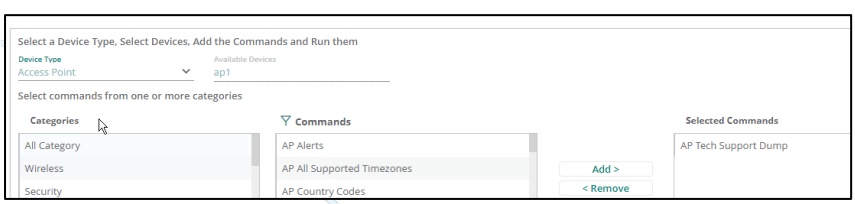

4. Click **Run**.
5. It may take some time to complete this command. Upon completion, download the command output using the **download** icon.

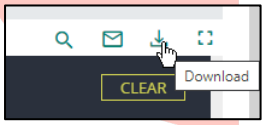

This downloads a zip file to your PC that contains the tech support files.
Example on a Windows system:

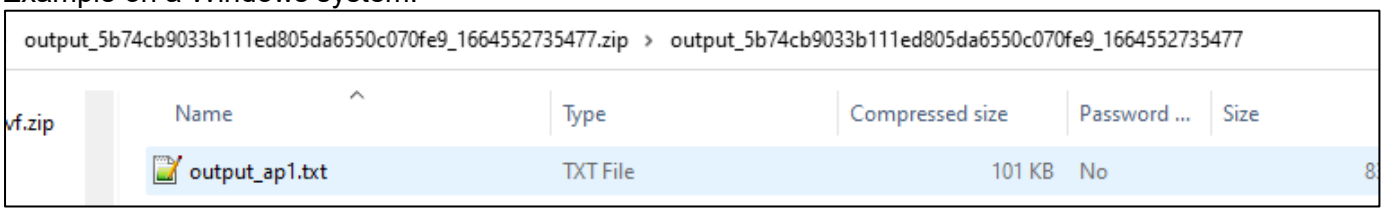

You can share this file with Aruba TAC or your senior engineer for further analysis.

Switch Tech Support Collection

In the next steps you will collect switch support information.

6. In Aruba Central, navigate to Context: **Global** > Navigation: **Devices** > Top: **Switches**
7. Click sw-edge1 to see its device page.

> **NOTE:** Just like with the AP, you can also use the AI Search bar:
> Enter **sw-edge1**, press <ENTER> and click the switch name in the search result window.

You are now at this location in Aruba Central:
Context: **Device / sw-edge1** > Navigation: **Overview** > Top: **Summary**

8. Take note of the product code of the switch, this typically starts with Jxxx. You will need this information later in this lab.

Example output showing JL666A:

Product code:_____

9. At the top-right, click the **Actions** dropdown list and select **Tech Support**.

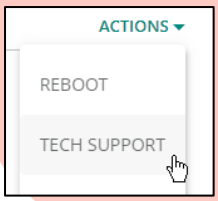

This action takes you to the Tools > Commands page.

Question: What is the currently selected *Command*?

Answer: The command **show tech** was automatically selected.

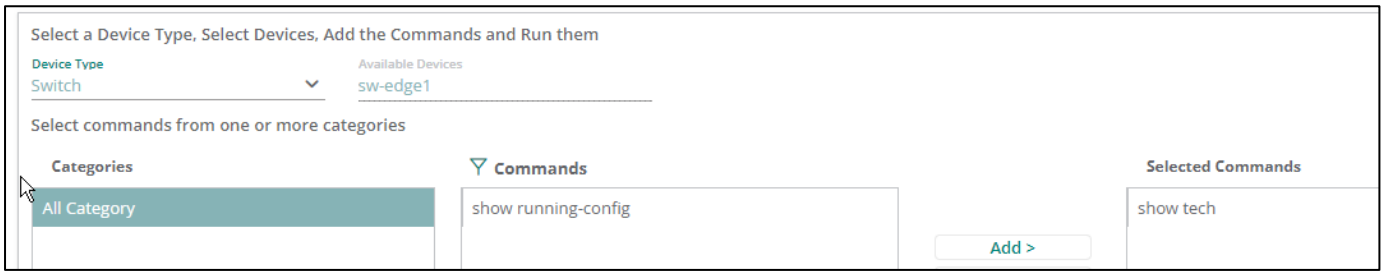

10. Click **Run**.

> **NOTE:** Aruba Central reserves a limited buffer for the command output. If the buffer is full, the system will not display the output, but provide you with a download or email option.

11. It may take some time to complete this command. Then download the command output using the **download** icon.

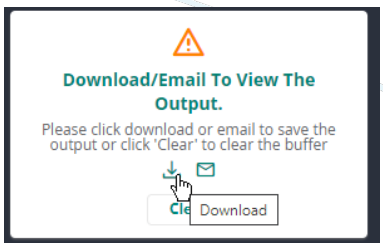

This downloads a zip file to your PC that contains the tech support files.
Example on a Windows system:

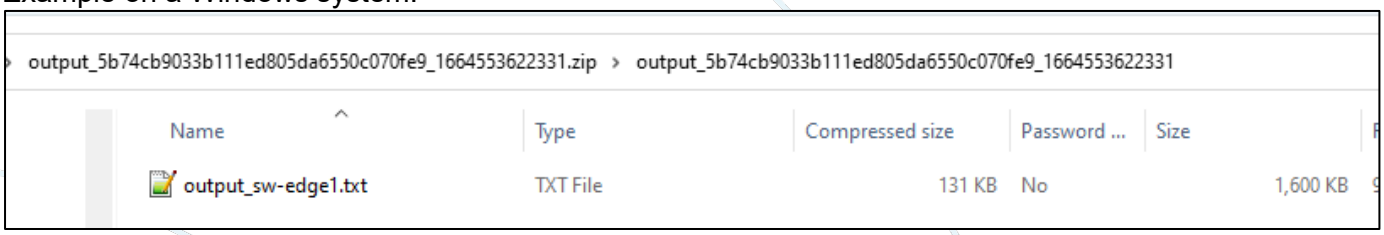

Again, you can share this file with Aruba TAC or your senior engineer for further analysis.

Task 2: Troubleshoot Switch to Aruba Central connection

In this task you will learn to troubleshoot the connection between an AOS-CX switch and Aruba Central. In the scenario, the customer has uploaded a template to Aruba Central that contains a configuration error. As a result of the configuration error, the switch can no longer connect to Aruba Central. The customer wants you to resolve this issue: to ensure that the switch connects to Aruba Central again.

Objectives

- Verify the Aruba Central connection.
- Understand the function of the Aruba support command.

Steps

Initially, the lab environment is still functional. You will explore the admin access level you have on the switches.

Explore default configuration status

When a device is managed by Aruba Central, access to configuration mode is limited. This prevents changing the configuration locally on the switches. All configuration changes must be done via Aruba Central.

1. Use MGMT PC to open an SSH connection to sw-edge1.
2. Login using **admin** / **Aruba123!**
3. Review the current Aruba Central output.

```
show aruba-central
```

```
sw-edge1# show aruba-central
Central admin state                       : enabled

Central location                          : device-uswest4-d2.central.arubanetworks.com
VRF for connection                        : default
Shared Token                              : N/A
Central connection status                 : connected

Central source                            : activate
Central source connection status          : connected
Central source last connected on          : Thu Oct 20 14:44:59 UTC 2022
System time synchronized from Activate    : True

Activate Server URL                       : devices-v2.arubanetworks.com
CLI location                              : N/A
CLI VRF                                   : N/A

Source IP                                 : 10.1.3.4
Source IP Overridden                      : False

Central support mode                      : disabled
```

Question: What is the Aruba Central connection status?

Answer: Connected.

Aruba Central Configuration Lockout

Check the status of the Configuration Lockout

4. Review the running configuration, look for the configuration lockout command.

```
show running-config | include config
```

```
sw-edge1# show running-config | include config
Current configuration:
configuration-lockout central managed
```

Question: What does it mean?

Answer: When the device is connected to Aruba Central, it is not possible to make configuration changes.

5. Attempt to access the configuration mode and review the available options.

```
config
?
```

```
sw-edge1# configure
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
```

6. Attempt to create a new VLAN.

```
vlan 1000
```

```
sw-edge1(config)# vlan 1000
Invalid input: vlan
```

Question: What do you observe?

Answer: The command is not available. The device is currently managed by Aruba Central, to which it has an active connection. You are now locked out of making configuration changes.

> **NOTE:** Please note that it is still possible to disable the Aruba Central feature on the switch in this mode. If you do this, the switch reverts to local configuration mode.

Aruba Support Mode

In the next steps you will review the default Aruba Support mode status.

7. Review the previous show aruba-central output.

```
show aruba-central
```

```
sw-edge1(config)# show aruba-central
Central admin state                     : enabled

Central location                        : device-uswest4-d2.central.arubanetworks.com
VRF for connection                      : default
Shared Token                            : N/A
Central connection status               : connected

Central source                          : activate
Central source connection status        : connected
Central source last connected on        : Thu Oct 20 14:44:59 UTC 2022
System time synchronized from Activate  : True

Activate Server URL                     : devices-v2.arubanetworks.com
CLI location                            : N/A
CLI VRF                                 : N/A

Source IP                               : 10.1.3.4
Source IP Overridden                    : False

Central support mode                    : disabled
```

Question: What is the current support mode?

Answer: Disabled.

Question: What does it mean?

Answer: With Aruba support mode enabled, Aruba Central will not perform configuration sync and it will be possible to make local configuration changes. After disabling the support mode, Aruba Central will push the latest version of the configuration (known to Aruba Central) to the device again. You will use this command later in this task.

Apply a Template with a Configuration Error

In the next steps you will load a template that contains a configuration error.

8. In Aruba Central, navigate to Context: **Groups / campus-sw-edge-tpl** > Navigation: **Devices** > Top: **Switches** > Right-top: **Config**
9. Click **Templates**.

10. From the ACAF Student Files, open lab14.01-sw-edge-tpl-troubleshoot-Jxxxx.txt. Replace xxxx with your switch product code. You have registered this product code earlier in this lab.

Example files:

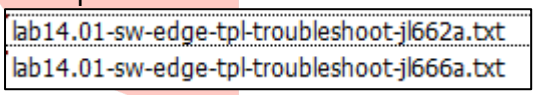

11. Copy the contents of the file.
12. In Aruba Central, **edit** the template **sw-edge** using the pencil icon.
13. **Replace** the contents of the template with the contents of the troubleshoot text file.
14. Click **Save**.

Verify problem on the switch

In the next steps you will verify that the switch has a problem connecting to Aruba Central.

15. Switch to the MGMT PC SSH connection to sw-edge1.
16. Wait about one minute for the template to get pushed to the device.
17. Attempt to reach an internet IP. This should fail now, because of the template error.

```
ping 8.8.8.8
```

```
sw-edge1(config)# ping 8.8.8.8
PING 8.8.8.8 (8.8.8.8) 100(128) bytes of data.
From 10.1.3.4 icmp_seq=1 Destination Host Unreachable
From 10.1.3.4 icmp_seq=2 Destination Host Unreachable
From 10.1.3.4 icmp_seq=3 Destination Host Unreachable

--- 8.8.8.8 ping statistics ---
5 packets transmitted, 0 received, +3 errors, 100% packet loss, time 4024ms
```

18. Review the Aruba Central status

```
show aruba-central
```

```
sw-edge1(config)# show aruba-central
Central admin state                    : enabled

Central location                       : device-uswest4-d2.central.arubanetworks.com
VRF for connection                     : default
Shared Token                           : N/A
Central connection status              : connected
...
```

19. Repeat the command every 30 seconds. The connection should be disconnected after about a minute.

```
show aruba-central
```

```
sw-edge1(config)# show aruba-central
Central admin state                     : enabled

Central location                        : device-uswest4-d2.central.arubanetworks.com
VRF for connection                      : default
Shared Token                            : N/A
Central connection status               : connection_failure
...
```

Verify Configuration Access

Now verify your access to the configuration mode and the available commands.

20. While in configuration mode, check the available options.

```
?
```

```
sw-edge1(config)# ?
  aaa                           Configure Authentication, Authorization and
                                Accounting feature
  access-list                   Access control list (ACL)
  alias                         Create a short name for the specified
                                command(s).
  allow-unsafe-updates          Allow non-failsafe updates of programmable
                                devices
  allow-unsupported-transceiver Allow unsupported transceivers
  apply                         Apply a configuration record
  aruba-central                 Configure Aruba-Central
  banner                        Customize login banner
  bfd                           Enable Bidirectional Forwarding Detection
                                (BFD)
  bluetooth                     Configure Bluetooth wireless management
...
```

21. Attempt to create a VLAN

```
vlan 1000
exit
```

```
sw-edge1(config)# vlan 1000
sw-edge1(config-vlan-1000)# exit
```

Question: Did this work?

Answer: Yes, when the Aruba Central connection is disconnected, local configuration changes can be made. These local changes will be overwritten by the Aruba Central version again when the connection is restored.

22. **Optional step**: You may attempt to troubleshoot the configuration issue yourself. If you get stuck, you can move to the next page to implement the solution.

First attempt to correct the issue

You will now discover the configuration issue that was introduced by the template.

23. Review the IP routing table

```
show ip route
```

```
sw-edge1(config)# show ip route

Displaying ipv4 routes selected for forwarding

Origin Codes: C - connected, S - static, L - local
              R - RIP, B - BGP, O - OSPF
Type Codes:   E - External BGP, I - Internal BGP, V - VPN, EV - EVPN
              IA - OSPF internal area, E1 - OSPF external type 1
              E2 - OSPF external type 2

VRF: default

Prefix              Nexthop          Interface     VRF(egress)       Origin/
Distance/    Age
                                                                     Type      Metric
-------------------------------------------------------------------------------
-------------------
0.0.0.0/0           10.1.3.254       vlan3         -                 S         [1/0]
00h:03m:15s
10.1.3.0/24         -                vlan3         -                 C         [0/0]
-
10.1.3.4/32         -                vlan3         -                 L         [0/0]
-

Total Route Count : 3
```

Question: What is the default route next hop?

Answer: The default route 0.0.0.0/0 points to 10.1.3.254.

24. Review the route commands in the running configuration.

```
show run | include route
```

```
sw-edge1(config)# show run | include route
ip route 0.0.0.0/0 10.1.3.254
```

25. Since you have access to the configuration, you can apply the correct command to the configuration.

```
ip route 0.0.0.0/0 10.1.3.1
exit
```

```
sw-edge1(config)# ip route 0.0.0.0/0 10.1.3.1
sw-edge1(config)# exit
```

26. Verify you can reach the internet again. Ping 8.8.8.8, use 180 repetitions.

TIP: When you expect longer command output, such as a ping with many repetitions, your command output will be stopped due to the paging system.

Example:

```
-- MORE --, next page: Space, next line: Enter, quit: q
```

You can disable the paging before running your command to prevent your command from being stopped.

```
no page
```

```
no page
ping 8.8.8.8 repetitions 180
```

```
sw-edge1# no page
sw-edge1# ping 8.8.8.8 repetitions 180
PING 8.8.8.8 (8.8.8.8) 100(128) bytes of data.
76 bytes from 8.8.8.8: icmp_seq=1 ttl=114 (truncated)
76 bytes from 8.8.8.8: icmp_seq=2 ttl=114 (truncated)
76 bytes from 8.8.8.8: icmp_seq=3 ttl=114 (truncated)
76 bytes from 8.8.8.8: icmp_seq=4 ttl=114 (truncated)
...
```

Question: What do you observe?

Answer: The ping works again.

27. Leave the ping running, wait about 1-3 minutes. The switch will re-establish the connection to Aruba Central in this period.

```
76 bytes from 8.8.8.8: icmp_seq=23 ttl=114 (truncated)
76 bytes from 8.8.8.8: icmp_seq=24 ttl=114 (truncated)
76 bytes from 8.8.8.8: icmp_seq=25 ttl=114 (truncated)
76 bytes from 8.8.8.8: icmp_seq=26 ttl=114 (truncated)
76 bytes from 8.8.8.8: icmp_seq=27 ttl=114 (truncated)
From 10.1.3.4 icmp_seq=28 Destination Host Unreachable
From 10.1.3.4 icmp_seq=29 Destination Host Unreachable
From 10.1.3.4 icmp_seq=31 Destination Host Unreachable
From 10.1.3.4 icmp_seq=32 Destination Host Unreachable
From 10.1.3.4 icmp_seq=34 Destination Host Unreachable
From 10.1.3.4 icmp_seq=35 Destination Host Unreachable
```

Question: What happens with the ping after about 1-2 minutes? (This may occur before earlier as well; it depends on when the connection to Aruba Central was restored.)

Answer: It stopped again with a *destination host unreachable* message.

Question: What happened?

Answer: When internet access was restored, the switch connected to Aruba Central and received the configuration version from Aruba Central again. This re-introduced the issue in the switch configuration.

Second attempt to correct the issue.

You will now try this again, but now you will use the correct procedure. To give you time to troubleshoot, you enable Aruba Support mode. This prevents the configuration push, even when the Aruba Central connection is active.

28. Enable aruba support mode

```
aruba-central support-mode
```

```
sw-edge1# aruba-central support-mode
```

29. Change the default route again.

```
config
ip route 0.0.0.0/0 10.1.3.1
```

```
sw-edge1# config
sw-edge1(config)# ip route 0.0.0.0/0 10.1.3.1
```

30. Verify Aruba Central is connected again (this may take a few minutes).

```
show aruba-central
```

NOTE: You can force the connection attempt by disabling and enabling the Aruba Central feature on the switch.

```
sw-edge1(config)# aruba-central
sw-edge1(config-aruba-central)# disable
sw-edge1(config-aruba-central)# enable
sw-edge1(config-aruba-central)# exit
```

As you have now 'found and corrected' the issue on the switch (for now), the next step is to apply this change to the template in Aruba Central.

31. In Aruba Central, correct the template. Navigate to Context: **Groups / campus-sw-edge-tpl** > Navigation: **Devices** > Top: **Switches** > Right-top: **Config**
32. Click **Templates**.
33. You can either change the configuration yourself in the template or upload the known-good text from the lab drive.

34. From the ACAF Student Files, open **lab07.02-sw-edge-tpl-jxxxx.txt**. (Replace Jxxxx with your product code).
35. Copy the contents of the file.
36. In Aruba Central, edit the template sw-edge using the pencil icon.
37. **Replace** the contents of the template.
38. Click **Save**.

Verify Configuration Push and the Resulting Status: In Sync

39. On the MGMT PC, switch to the SSH connection to sw-edge1.
40. Disable Aruba support mode. This will trigger the template configuration push.

```
no aruba-central support-mode
```

41. In Aruba Central, navigate to Context: **Groups / campus-sw-edge-tpl** > Navigation: **Devices** > Top: **Switches** > Right-top: **List**
42. In the Config Status column, the value should be **In Sync**.
43. On the MGMT PC, switch to the SSH connection to sw-edge1.

```
show aruba-central
```

```
sw-edge1(config)# show aruba-central
Central admin state                     : enabled

Central location                        : device-uswest4-d2.central.arubanetworks.com
VRF for connection                      : default
Shared Token                            : N/A
Central connection status               : connected
...
```

44. Verify the connection to Aruba Central is still active.

Task 3: Site Level Troubleshooting

When an issue occurs at a site, it may be useful to see what configuration changes were made to site devices in recent history. In this task you will learn how to see when configuration changes were made, on what devices they were made, and what those changes were.

Objectives

- Explore site level change logs in Aruba Central.

Steps

1. In Aruba Central, navigate to Context: **Site / site-campus-main** > Navigation: **Overview** > Top: **Site Health**
2. Under the Summary Statistics, check the **Change Log** tile.
3. At the top-right, change the period history to **1 day**.

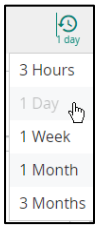

4. The change log shows configuration, firmware, and reboot events that were initiated from Aruba Central.

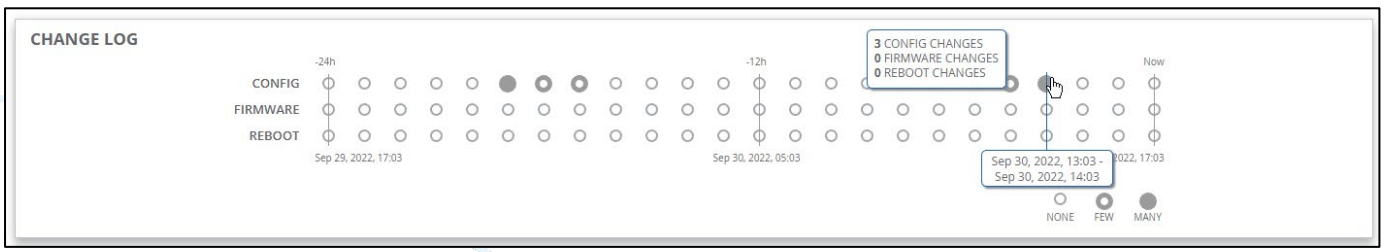

In the previous lab activity (Monitoring) you disabled the guest WLAN. This was registered by Aruba Central as a configuration change.

5. Hover your mouse over the config row over specific dots (days and time of day). An overlay will appear and will show how many configuration changes were made.
6. Click the **dot** that contains a config change.

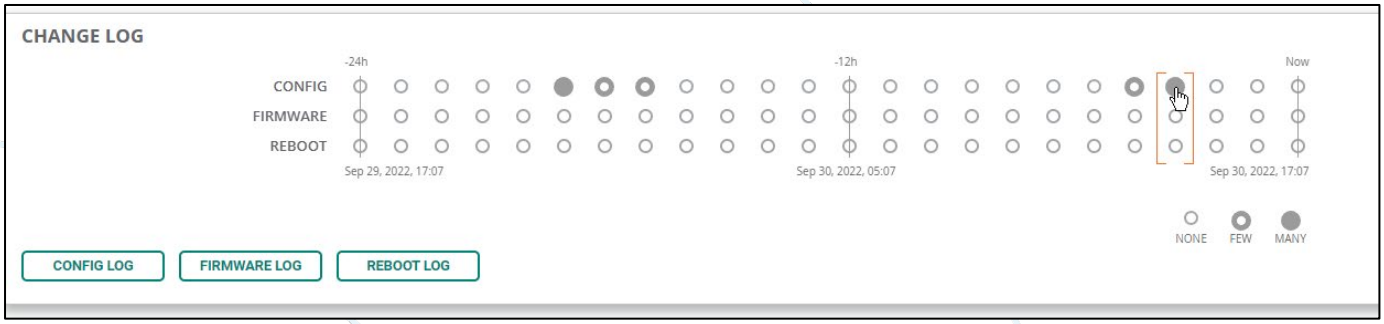

7. At the bottom-left you see the config log. Click the **Config Log**.
8. Expand the entry (or entries) in the log to see the details.

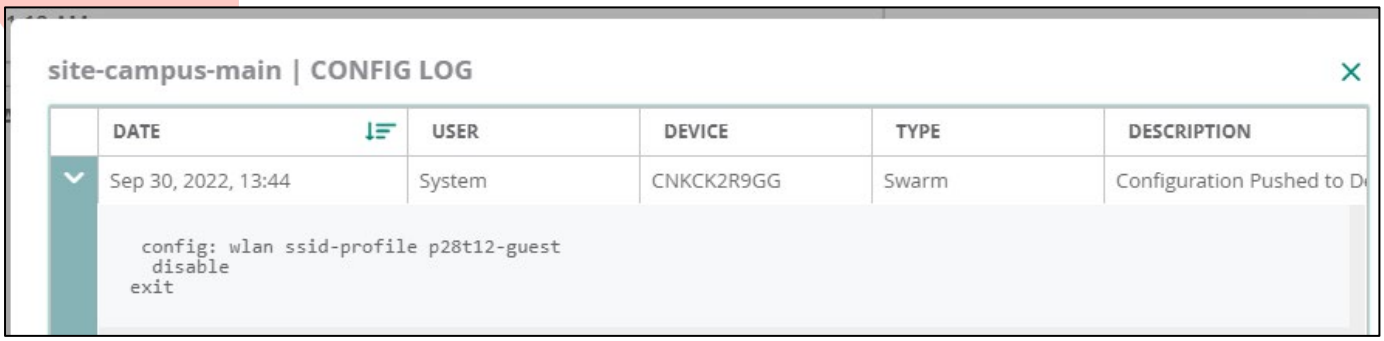

9. This may help you to collect data about the problem at the site.

Review the Alert System Acknowledgement

In the previous lab you enabled alerts when an AP is disconnected. With the switch template issues that were introduced in one of the previous tasks, your AP was disconnected for a while, and then it reconnected. Aruba Central will notice this and automatically acknowledge this alert.

10. In Aruba Central, navigate to Context: **Global** > Navigation: **Alerts & Events**
11. Enable Show Acknowledged Alerts.
12. Review the **AP Disconnected** Alert.

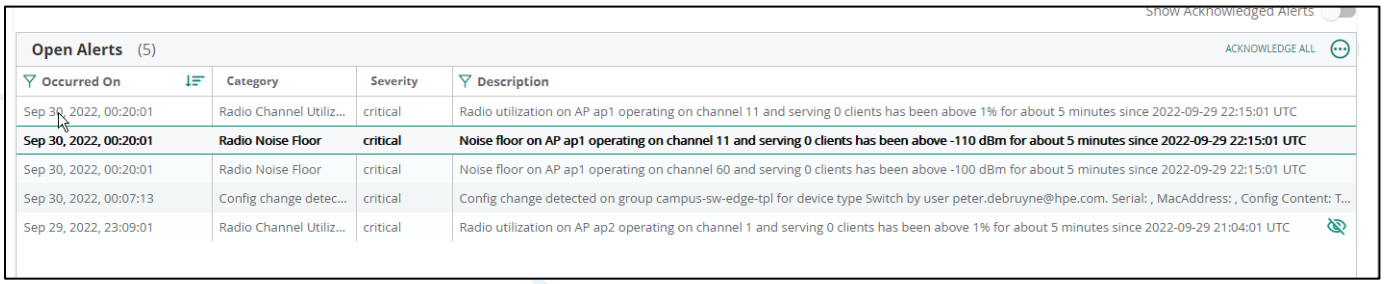

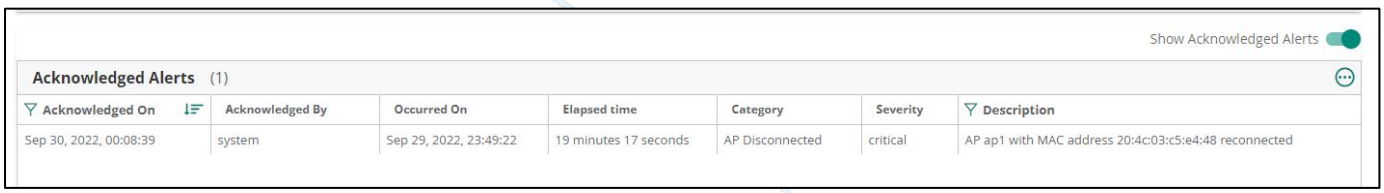

Question: What is the user that acknowledged the AP Disconnected alert?

Answer: System. When Aruba Central detects that the AP is online again, it automatically acknowledges the alert.

Question: What is the elapsed time?

Answer: This is the time it took you to correct the switch connection to Aruba Central.

You have completed the Lab!

Lab Appendix A: Lab Diagrams

Overview

In this appendix you can find the topology diagrams used in this training.

Base Lab Topology

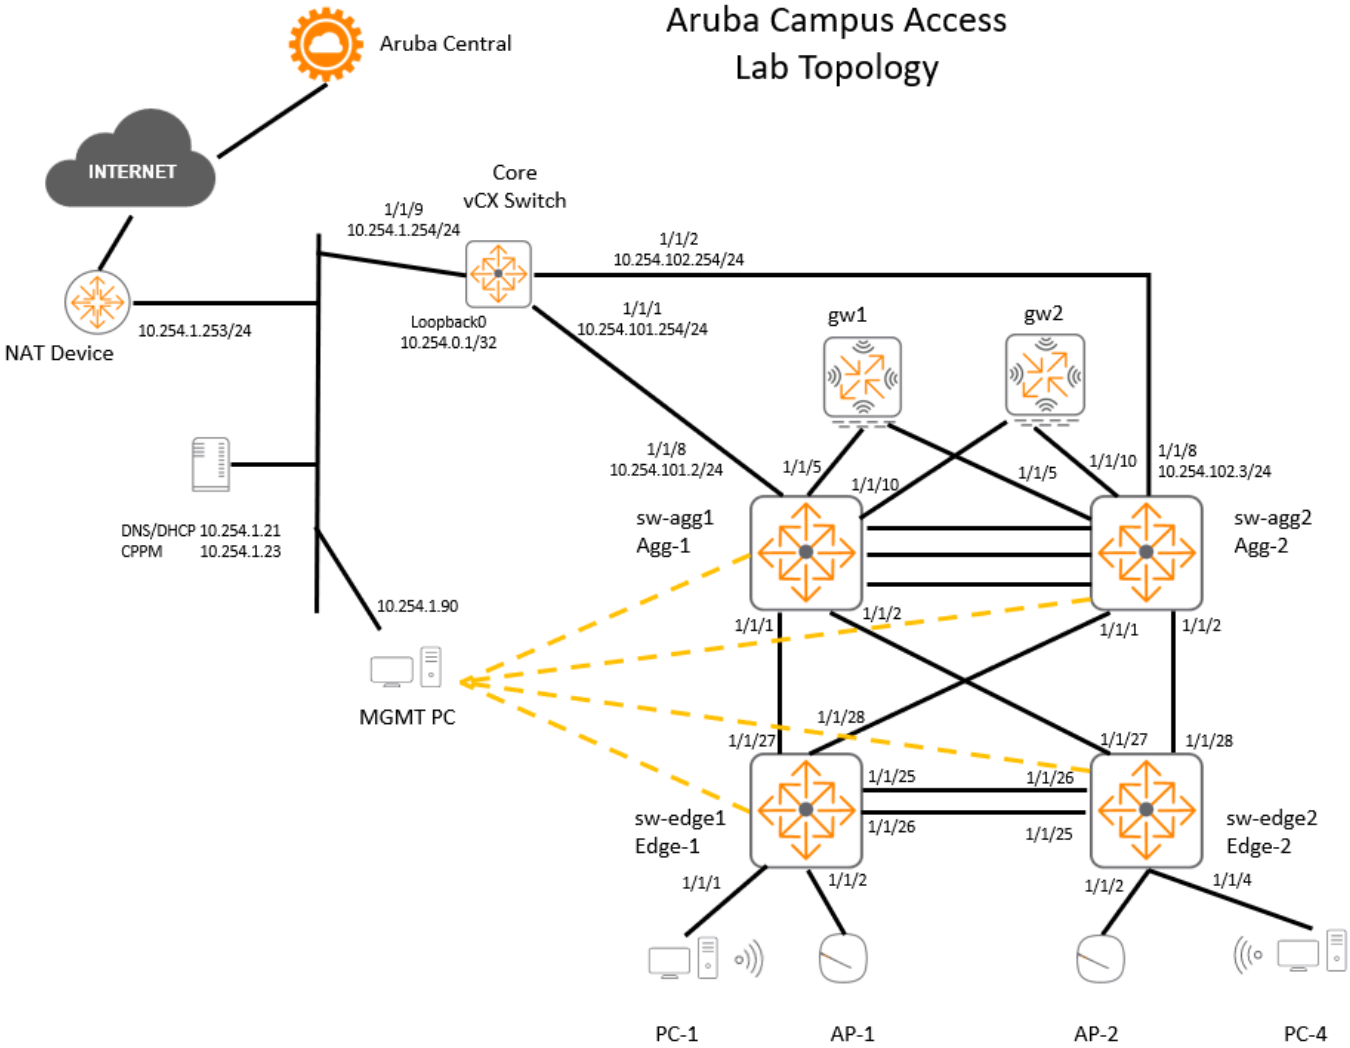

Base IP Topology and VLAN List

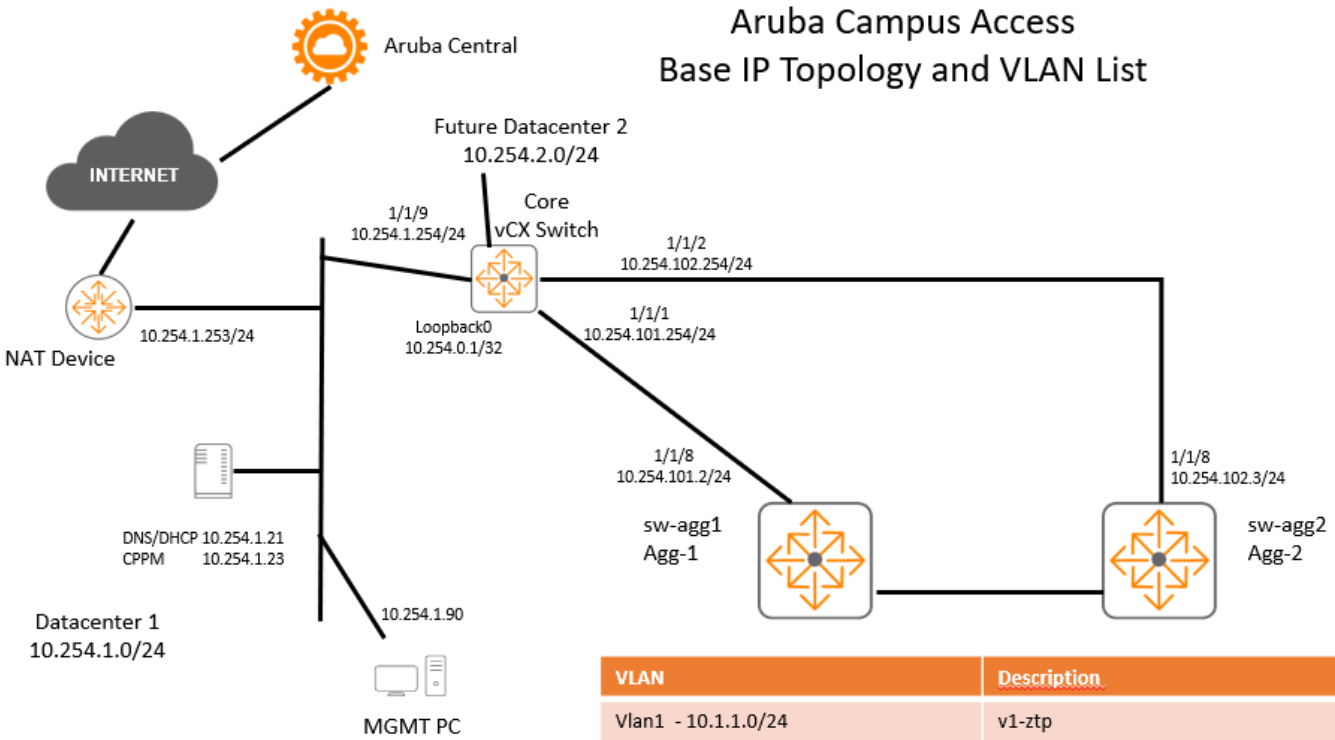

VLAN	Description
Vlan1 - 10.1.1.0/24	v1-ztp
Vlan 3 – 10.1.3.0/24	v3-mgmt
Vlan 4 – 10.1.4.0/24	v4-ap
Vlan 11 – 10.1.11.0/24	v11-employee
Vlan 12 – 10.1.12.0/24	v12-employee
Vlan 13 – 10.1.13.0/24	v13-voice
Vlan 14 – 10.1.14.0/24	v14-iot
Vlan 15 – 10.1.15.0.24	v15-guest

Device OOBM Subnet and IP addresses

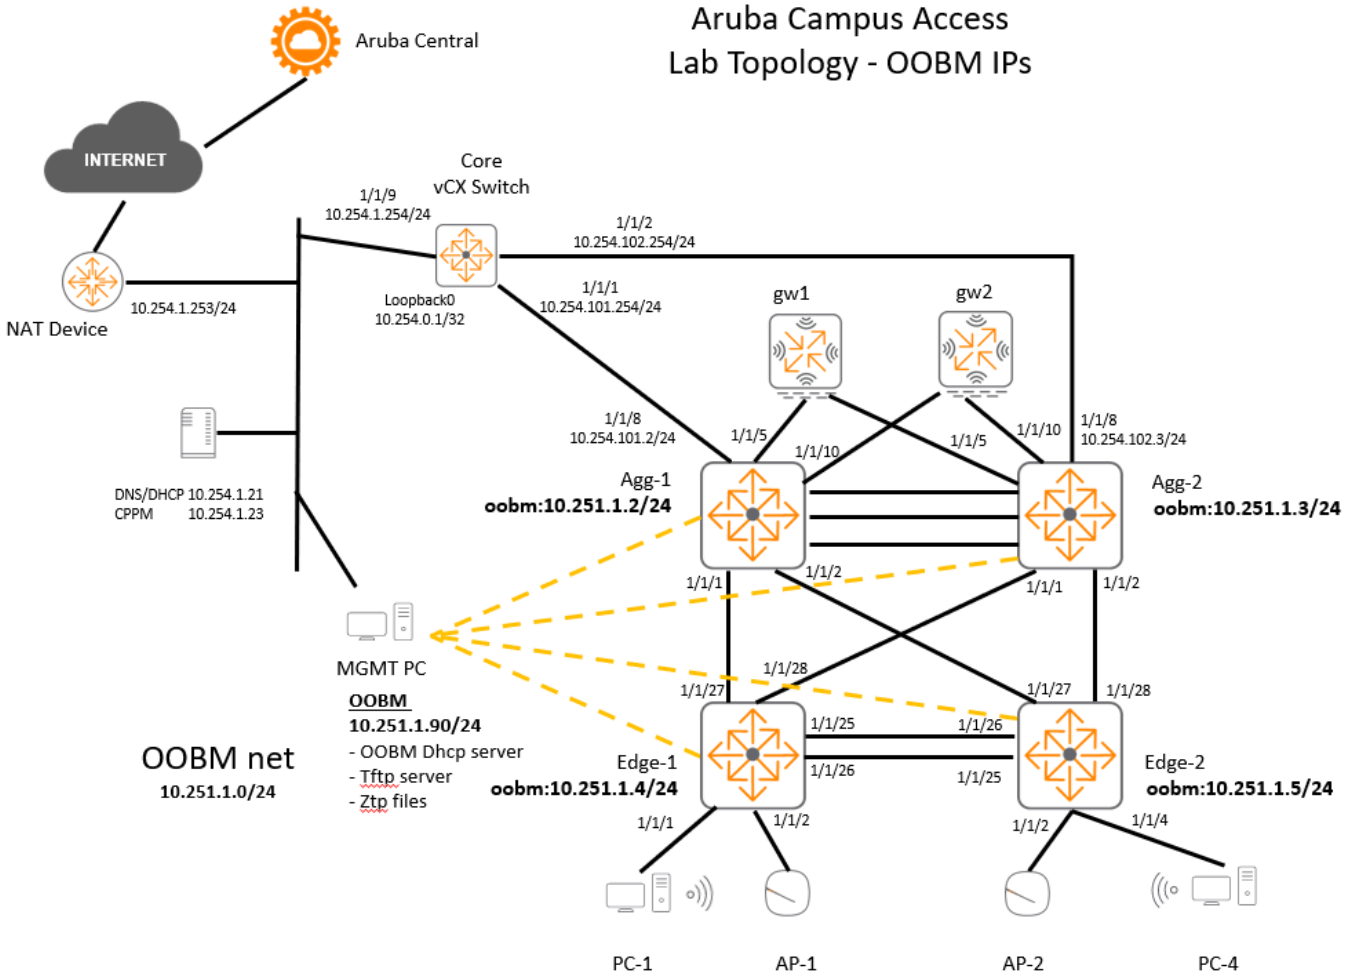

Lab 02.02 - Link Aggregation

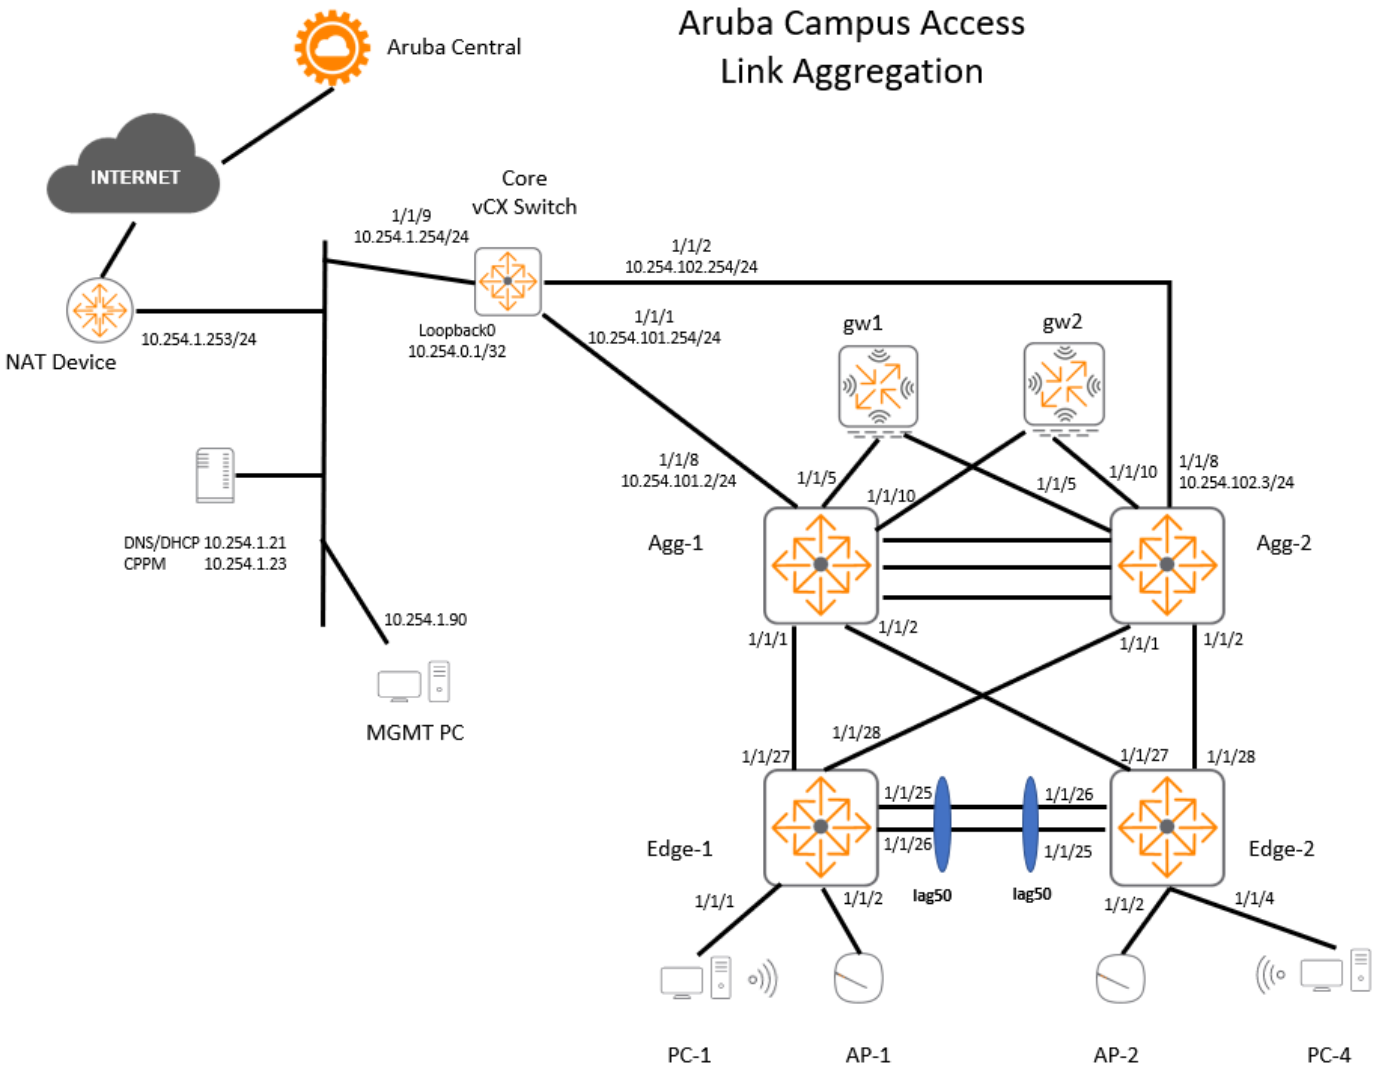

Lab 03 – Static Routes and OSPF

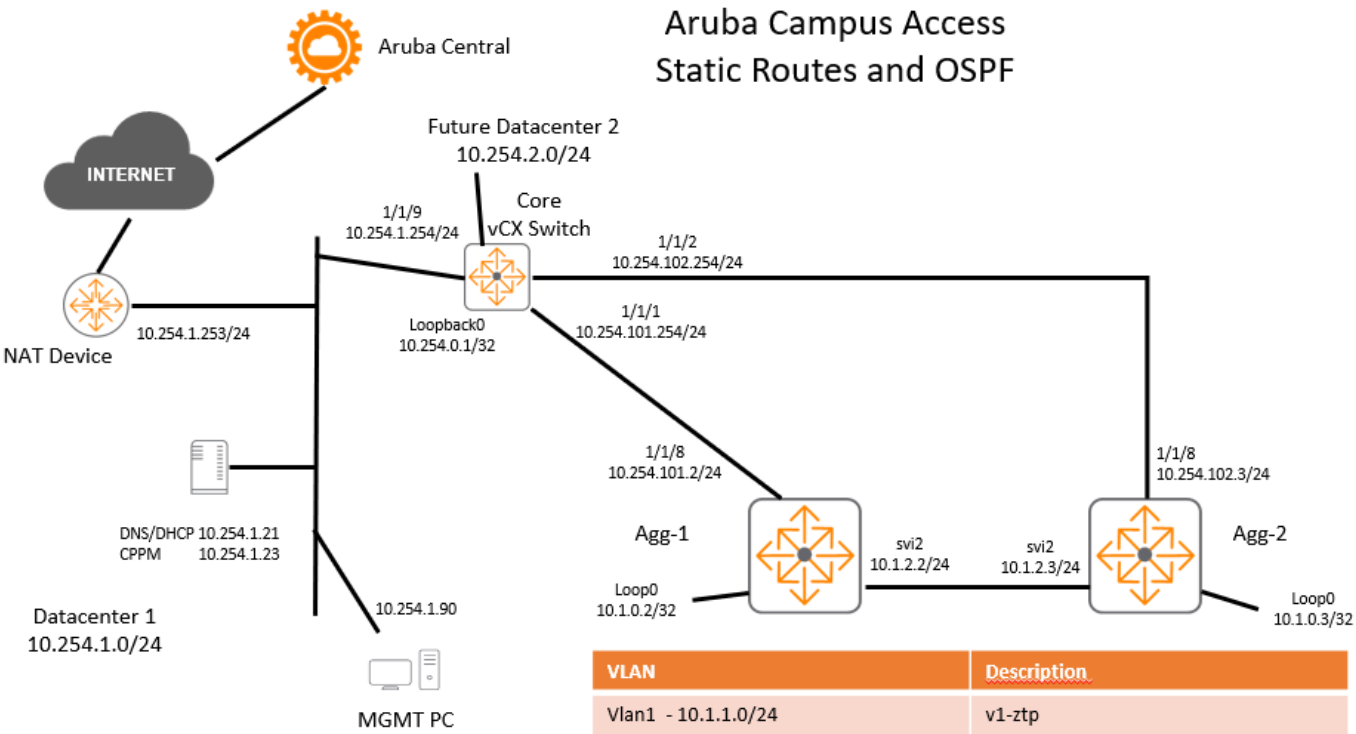

VLAN	Description
Vlan1 - 10.1.1.0/24	v1-ztp
Vlan 3 – 10.1.3.0/24	v3-mgmt
Vlan 4 – 10.1.4.0/24	v4-ap
Vlan 11 – 10.1.11.0/24	v11-employee
Vlan 12 – 10.1.12.0/24	v12-employee
Vlan 13 – 10.1.13.0/24	v13-voice
Vlan 14 – 10.1.14.0/24	v14-iot
Vlan 15 – 10.1.15.0.24	v15-guest

Lab 04.03 VSX

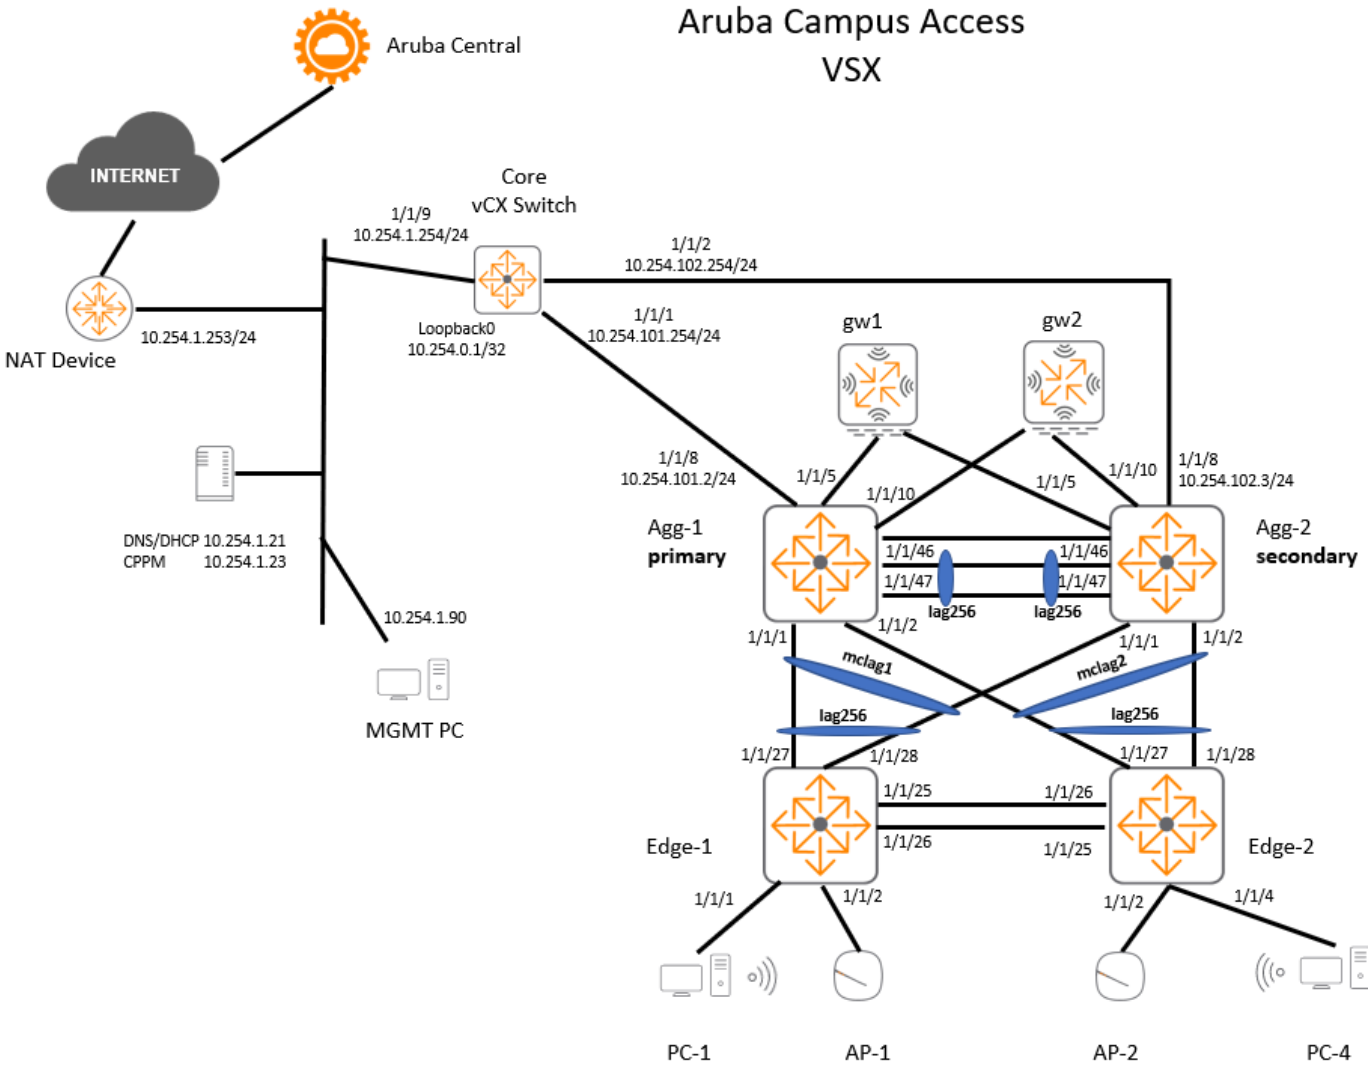

6280 AMERICA CENTER DR. SAN JOSE, CA 95002
TEL: 408.227.4500 | FAX: 408.227.4550
www.ARUBANETWORKS.com

EDU-ACAF-RLABS-v22.41