aruba®

NETWORKS

Implementing ArubaOS-CX Switching

Remote Labs

TRAINING MANUAL

EDU-ICX-RLABS-v20.211

Implementing ArubaOS-CX Switching

20.211

Lab Guide

February 2022

Implementing ArubaOS-CX Switching Lab Guide rev 20.211

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

http://www.arubanetworks.com/open_source

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

Warranty
This hardware product is protected by the standard Aruba warranty of one year parts/labor. For more information, refer to the ARUBACARE SERVICE AND SUPPORT TERMS AND CONDITIONS.
Altering this device (such as painting it) voids the warranty.

SKU: EDU-ICX-RLABS-v20.211
February 2022

Implementing ArubaOS-CX Switching Lab Guide
Table of Contents

aruba
a Hewlett Packard Enterprise company

EDUCATION SERVICES

Implementing ArubaOS-CX Switching Lab Guide

Lab 01: Base Configuration- Initial Lab Setup

Lab Diagram

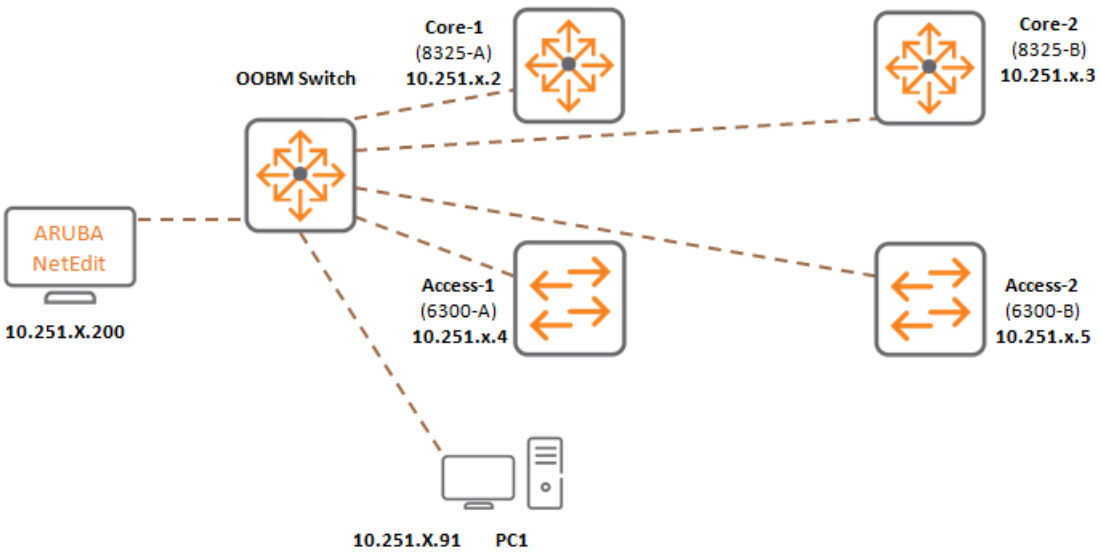

Overview

In this lab activity all switches will be factory reset and the Out-Of-Band Management (OOBM) will be configured.

At the end of the configuration, a configuration checkpoint will be made.

Objectives

- Configure the OOBM network
- Prepare a basic configuration checkpoint

Task 1: Factory Reset of Devices

Objectives

- Factory reset
- Remove all checkpoints

Note that in case AOS-CX is completely factory default, the switch will prompt to change the admin password at first login.

This will also happen after the 'erase all zeroize' procedure. This procedure will not only clear the startup configuration, but it will also clear any other custom files, such as configuration checkpoints.

Steps

1. Open a console connection to the 6300A (Access-1). Login using admin, no password.

NOTE: The initial hostname may be different from the output below, this depends on the remote lab setup.

```
6300 login: admin
Password:

Last login: 2019-12-08 05:26:49 from the console
User "admin" has logged in 2 times in the past 30 days
6300 #
```

NOTE: In case the switch prompts to change the password at this point, it is already factory default, so there is no need to perform the 'erase all zeroize' of the next step. Move on to the next switch, the password will be set in the next task.

2. Clear all configuration files and snapshots.

```
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.
P52-6300A-Table12#
```

3. Open a console connection to 6300B (Access-2). Login using admin, no password.
4. Clear all configuration files and snapshots (skip this step if the switch prompts you to change the admin password).

```
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.
P52-6300B-Table12#
```

5. Open a console to 8325A (Core-1) and repeat the factory reset process.

NOTE: If the switch shows a new password prompt, the switch is already at factory default, so you can move on to the next step.

6. Open a console to 8325B (Core-2) and repeat the factory reset process.

NOTE: If the switch shows a new password prompt, the switch is already at factory default, so you can move on to the next step.

Task 2: Configure the Out of Band Management (OOBM) for Access1

Objectives

- Initial password
- Make a checkpoint of the factory default configuration
- Disable all interfaces to prevent remote lab network loops
- Configure the OOBM management IP address
- Make the base configuration checkpoint

Steps

7. Switch to the console connection of the 6300A (Access-1).
8. Login using admin, no password.
9. Since the switch is factory default, the software will show a prompt to change the admin password. Set the admin password to **aruba123.**

```
6300 login: admin
Password:


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

NOTE: If you entered a different password, the password can be changed with these commands:

```
6300# configure terminal
6300(config)# user admin password plaintext aruba123
6300(config)# end
```

10. Make a checkpoint named 'icx-factory-default'. This checkpoint is not used in the future labs, but will allow you to switch back to the factory default state without the 'erase all zeroize' command and reboot.

```
6300# copy running-config checkpoint icx-factory-default
```

11. Enter the configuration mode and configure the hostname and OOBM IP.
Ask your instructor for your TABLE number (your kit/student number). Your table number can be found on the remote lab dashboard.

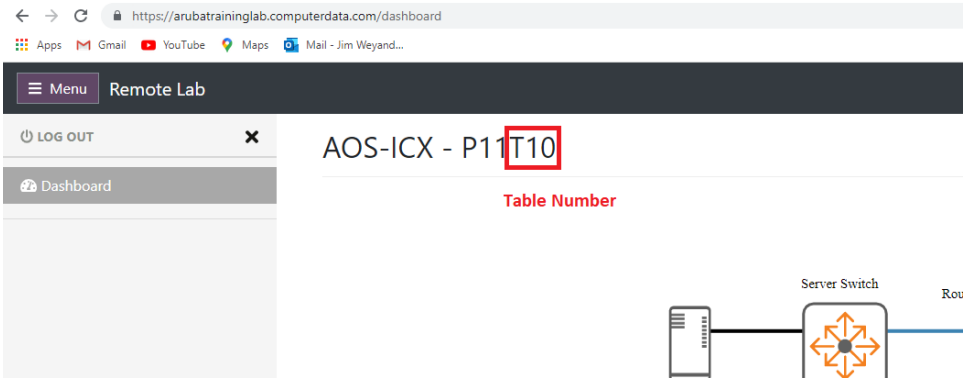

IMPORTANT: Change x to your table number.

IMPORTANT: The output of the commands was done on Table 12. Your output should show your own Table number.

```
6300# configure terminal
6300(config)# hostname ICX-Tx-Access1
ICX-Tx-Access1(config)# interface mgmt
ICX-Tx-Access1(config-if-mgmt)# ip static 10.251.x.4/24
ICX-Tx-Access1(config-if-mgmt)# default-gateway 10.251.x.254
ICX-Tx-Access1(config-if-mgmt)# exit
```

EXAMPLE: If the above would be applied to Table12, this would be the result:

```
6300# configure terminal
6300(config)# hostname ICX-Tx-Access1
ICX-Tx-Access1(config)# interface mgmt
ICX-Tx-Access1(config-if-mgmt)# ip static 10.251.12.4/24
ICX-Tx-Access1(config-if-mgmt)# default-gateway 10.251.12.254
ICX-Tx-Access1(config-if-mgmt)# exit
```

12. Disable all switch ports. Since the remote lab contains several redundant links between all the switches, all interfaces are disabled initially. You will enable the interfaces as needed in later labs.

```
ICX-Tx-Access1(config)# interface 1/1/1-1/1/28
ICX-Tx-Access1(config-if-<1/1/1-1/1/28>)# shutdown
ICX-Tx-Access1(config-if-<1/1/1-1/1/28>)# exit
ICX-Tx-Access1(config)# end
```

13. Verify access on the OOBM network with a ping to the ClearPass host (10.254.1.23).

```
ICX-Tx-Access1# ping 10.254.1.23 vrf mgmt
PING 10.254.1.23 (10.254.1.23) 100(128) bytes of data.
108 bytes from 10.254.1.23: icmp_seq=1 ttl=62 time=1.14 ms
108 bytes from 10.254.1.23: icmp_seq=2 ttl=62 time=0.875 ms
108 bytes from 10.254.1.23: icmp_seq=3 ttl=62 time=0.930 ms
108 bytes from 10.254.1.23: icmp_seq=4 ttl=62 time=1.04 ms
108 bytes from 10.254.1.23: icmp_seq=5 ttl=62 time=1.09 ms

--- 10.254.1.23 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4057ms
rtt min/avg/max/mdev = 0.875/1.017/1.147/0.108 ms
```

NOTE: If the ping was not successful, review your settings.

14. You have now completed the base config and verified connectivity. This configuration will be saved as the base configuration checkpoint for future labs. List current checkpoints. Only 'icx-factory-default' snapshot should be in the list after the zeroize factory reset.

NOTE: There may be an automatic snapshot if 5 minutes have passed since your last configuration change. This is the checkpoint post-configuration feature that is enabled by default.

```
ICX-Tx-Access1# show checkpoint list
icx-factory-default
```

15. Make a checkpoint of the current configuration as ACSP-lab-base.

```
ICX-Tx-Access1# copy running-config checkpoint icx-lab01-base
Configuration changes will take time to process, please be patient.
ICX-Tx-Access1#
```

16. Verify the checkpoint is now in the list (another checkpoint may have been created by the system already, this can be ignored).

```
ICX-Tx-Access1# show checkpoint list
CPC20191215142246
icx-lab01-base
icx-factory-default
ICX-Tx-Access1#
```

17. Verify the contents of the checkpoint.

```
ICX-Tx-Access1# show checkpoint icx-lab01-base
Checkpoint configuration:
!
<output omitted>
```

NOTE: If you made a mistake you can remove the checkpoint with the command:

```
erase checkpoint <NAME>
```

NOTE: If you need to see the configuration changes between two checkpoints or between a checkpoint and the running-configuration, the 'checkpoint diff' command can be used. Examples:

```
IAS-Tx-Access2# checkpoint diff icx-factory-default icx-lab01-base
IAS-Tx-Access2# checkpoint diff icx-lab01-base running-configuration
```

NOTE: Sometimes it is handy to copy and paste text directly from the lab guide. If you are unable to paste directly into your device or PC, CTRL-ALT-SHIFT will open a clipboard. Paste the content into the clipboard, CTRL-ALT-SHIFT to close the clipboard and then paste the contents into the device or PC.

Task 3: Configure the OOBM for Access2/Core1/Core2

Objectives

– Apply the initial configuration to the Access2, Core1 and Core2 switches

Steps

6300B Base config

18. Switch to the console connection of the 6300B (Access-2).
19. Login with admin/ no password.
20. Configure the admin password as **aruba123.**

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

21. Enter these configuration commands.

```
6300# copy running-config checkpoint icx-factory-default
Configuration changes will take time to process, please be patient.
6300# configure terminal
6300(config)# hostname ICX-Tx-Access2
ICX-Tx-Access2(config)# interface mgmt
ICX-Tx-Access2(config-if-mgmt)# ip static 10.251.x.5/24
ICX-Tx-Access2(config-if-mgmt)# default-gateway 10.251.x.254
ICX-Tx-Access2(config-if-mgmt)# exit
ICX-Tx-Access2(config)# interface 1/1/1-1/1/28
```

```
ICX-Tx-Access2(config-if-<1/1/1-1/1/28>)# shutdown
ICX-Tx-Access2(config-if-<1/1/1-1/1/28>)# exit
ICX-Tx-Access2(config)# end
```

22. Verify access to the ClearPass host on the management network.

```
ICX-Tx-Access2# ping 10.254.1.23 vrf mgmt
PING 10.254.1.23 (10.254.1.23) 100(128) bytes of data.
108 bytes from 10.254.1.23: icmp_seq=1 ttl=62 time=1.27 ms
108 bytes from 10.254.1.23: icmp_seq=2 ttl=62 time=1.20 ms
108 bytes from 10.254.1.23: icmp_seq=3 ttl=62 time=1.18 ms
108 bytes from 10.254.1.23: icmp_seq=4 ttl=62 time=1.05 ms
108 bytes from 10.254.1.23: icmp_seq=5 ttl=62 time=1.14 ms

--- 10.254.1.23 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4004ms
rtt min/avg/max/mdev = 1.057/1.173/1.278/0.072 ms
```

23. Save the configuration checkpoint 'icx-lab01-base'.

```
ICX-Tx-Access2# copy running-config checkpoint icx-lab01-base
Configuration changes will take time to process, please be patient.
```

8325A Base config

24. Switch to the console connection of the 8325A (Core-1).
25. Login with admin/ no password.
26. Configure the admin password as **aruba123.**

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

27. Enter these configuration commands.

NOTE: The command 'system interface-group 1 speed 10g' provides the switch with the instruction of handling port speed for this group of ports.

NOTE: On the 8xxx platforms:

- all ports are L3 routed ports and shutdown by default
- the HTTPS server is disabled on the in-band network (vrf 'default') by default
- the REST API access is read-only by default.

The command instructions below will enable HTTPS on the 'default' VRF (the default global routing table) and allow read-write access to the REST API.

```
8325# copy running-config checkpoint icx-factory-default
Configuration changes will take time to process, please be patient.
8325# configure terminal
8325(config)# hostname ICX-Tx-Core1
ICX-Tx-Core1(config)# interface mgmt
ICX-Tx-Core1(config-if-mgmt)# ip static 10.251.x.2/24
ICX-Tx-Core1(config-if-mgmt)# default-gateway 10.251.x.254
ICX-Tx-Core1(config-if-mgmt)# exit
ICX-Tx-Core1(config)# https-server rest access-mode read-write
ICX-Tx-Core1(config)# https-server vrf default
ICX-Tx-Core1(config)# system interface-group 1 speed 10g
Changing the group speed will disable all member interfaces that
do not match the new speed.

Continue (y/n)? y
ICX-Tx-Core2(config)# end
```

28. Test access to the ClearPass server on the management network.

```
ICX-Tx-Core1# ping 10.254.1.23 vrf mgmt
```

29. Save a configuration checkpoint name 'icx-lab01-base'.

```
ICX-Tx-Core1# copy running-config checkpoint icx-lab01-base
Configuration changes will take time to process, please be patient.
```

8325B Base config

30. Switch to the console connection of the 8325B (Core-2).
31. Login with admin/ no password.
32. Configure the admin password as **aruba123.**

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

33. Enter these configuration commands:

```
8325# copy running-config checkpoint icx-factory-default
Configuration changes will take time to process, please be patient.
8325# configure terminal
8325(config)# hostname ICX-Tx-Core2
ICX-Tx-Core2(config)# interface mgmt
ICX-Tx-Core2(config-if-mgmt)# ip static 10.251.x.3/24
ICX-Tx-Core2(config-if-mgmt)# default-gateway 10.251.x.254
ICX-Tx-Core2(config-if-mgmt)# exit
ICX-Tx-Core2(config)# https-server rest access-mode read-write
ICX-Tx-Core2(config)# https-server vrf default
ICX-Tx-Core2(config)# system interface-group 1 speed 10g
Changing the group speed will disable all member interfaces that
do not match the new speed.

Continue (y/n)? y
ICX-Tx-Core2(config)# end
```

34. Test access to the ClearPass server on the management network.

```
ICX-Tx-Core2# ping 10.254.1.23 vrf mgmt
```

35. Save a configuration checkpoint name 'icx-lab01-base'.

```
ICX-Tx-Core2# copy running-config checkpoint icx-lab01-base
Configuration changes will take time to process, please be patient.
```

You have completed Lab 1!

Lab 02: NetEdit- Managing and Monitoring AOS-CX Switches

Lab Diagram

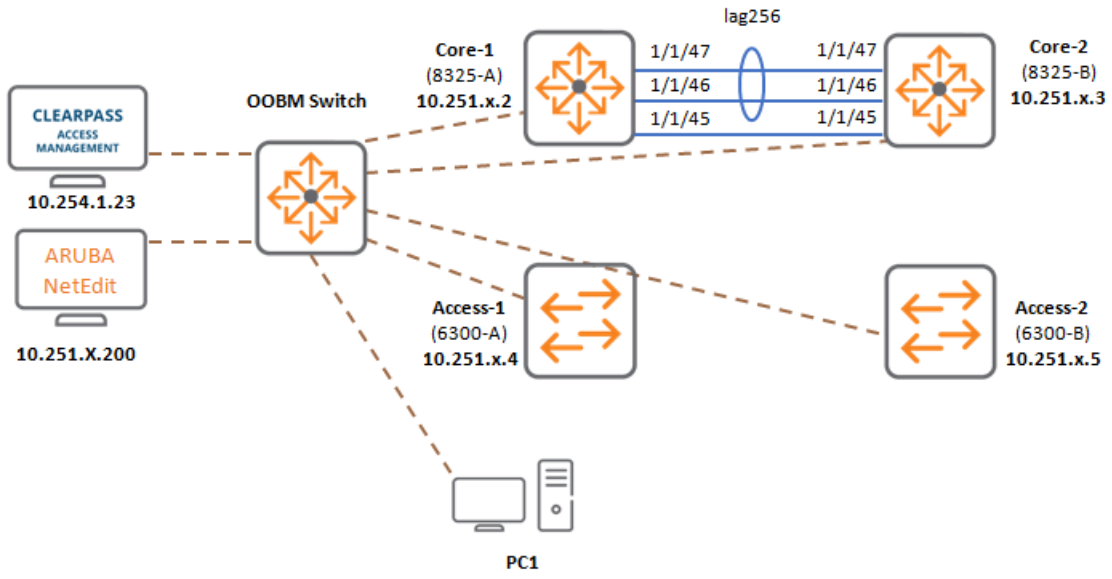

Overview

In this lab activity, NetEdit will be used to manage and monitor the switches.

First the switches will be discovered. Next a configuration plan will be pushed to the switches.

To verify if a configuration change has the expected result, it is possible to define validation rules in NetEdit. This feature will be explored in the lab.

The lab will show how devices in NetEdit can be logically grouped using attributes.

NetEdit supports conformance rules, so an administrator can have automated configuration checks and alerts if a configuration is not compliant.

Objectives

- Learn how to use the NetEdit interface

- Perform device discovery using NetEdit
- Learn how to use a plan to deploy configurations to the switches
- Understand and configure plan validation rules
- Understand device grouping using attributes
- Configure device conformance rules

Task 1: Verify Lab Start Configuration

Objectives

- If you have just completed Lab 01 - Base configuration, you can skip this task and move to the next task.
- If you have completed any other lab and are performing this lab again, complete these steps to get the base configuration on the devices.

Steps

1. Open a console connection to the 6300A. Login using admin, password aruba123.

```
ICX-Tx-Access1# copy checkpoint icx-lab01-base running-config
Configuration changes will take time to process, please be patient.
ICX-Tx-Access1#
```

2. Open a console connection to the 6300B. Login using admin, password aruba123.

```
ICX-Tx-Access2# copy checkpoint icx-lab01-base running-config
Configuration changes will take time to process, please be patient.
ICX-Tx-Access2#
```

3. Open a console connection to the 8325A. Login using admin, password aruba123.

```
ICX-Tx-Core1# copy checkpoint icx-lab01-base running-config
Configuration changes will take time to process, please be patient.
ICX-Tx-Core1#
```

4. Open a console connection to the 8325B. Login using admin, password aruba123.

```
ICX-Tx-Core2# copy checkpoint icx-lab01-base running-config
Configuration changes will take time to process, please be patient.
ICX-Tx-Core2#
```

Task 2: Access the NetEdit Interface

Diagram

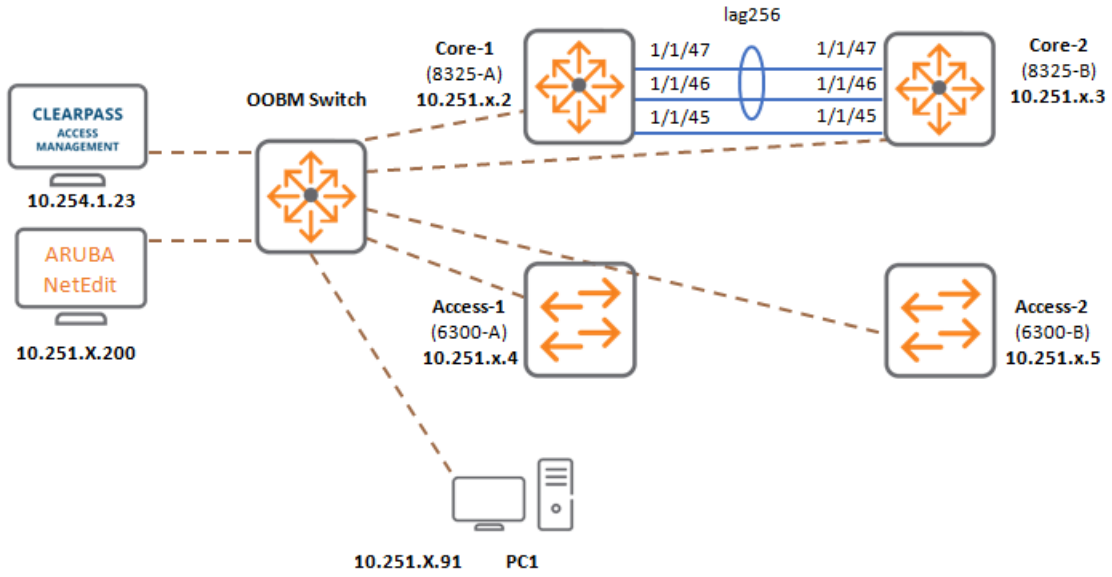

Objectives

– Basic NetEdit navigation

Steps

1. Open a session to PC1 (Management PC - OOBM).

NOTE: The management system has 3 NICs:

- **Do NOT Touch!**: connected to the lab for RDP access via the Remote Lab web GUI. This interface should **never** be disabled or changed.
- **Lab NIC**: This interface is connected to port 1/1/1 on the Access1 switch. It will be used when this station is used as a test host inside the network lab.
- **OOBM**: This interface is connected to the out-of-band-management (OOBM) network. You will use this interface to manage and access any management network relation activity

2. Verify that the NIC OOBM is enabled and the 'Lab NIC' is disabled.

TIP: To open the network settings, click the **Start** button, click **Settings (the gear icon)**. Click **Network & Internet**, click **Ethernet**. Click **Change adapter options**.

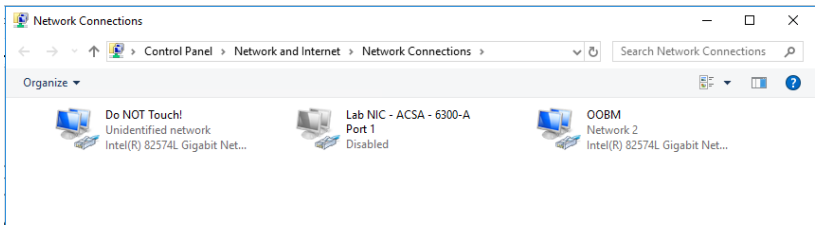

3. Verify that PC1 has an IP address on the OOBM NIC in the 10.251.x.0/24 subnet.

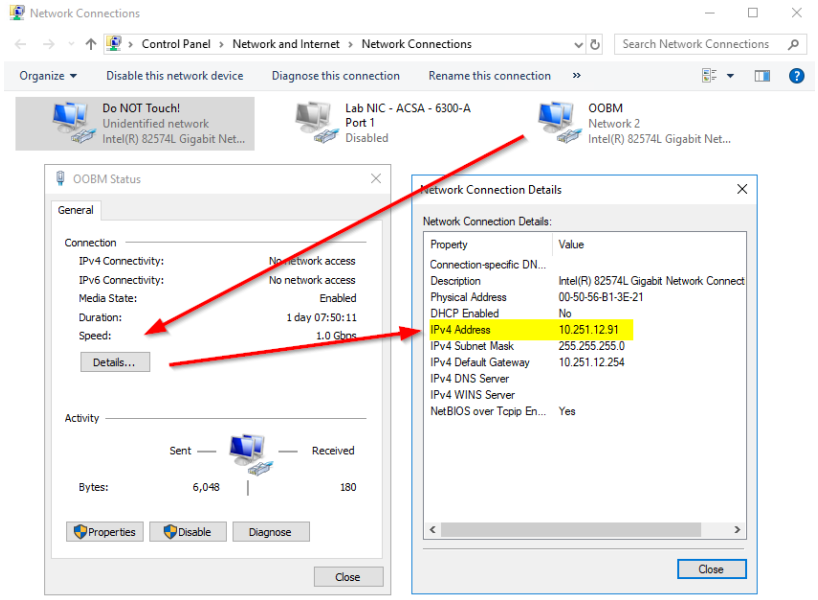

4. On the management station (PC1), open a web browser (Firefox for example) and navigate to the NetEdit host (10.251.x.200). This connection is made via the OOBM network.

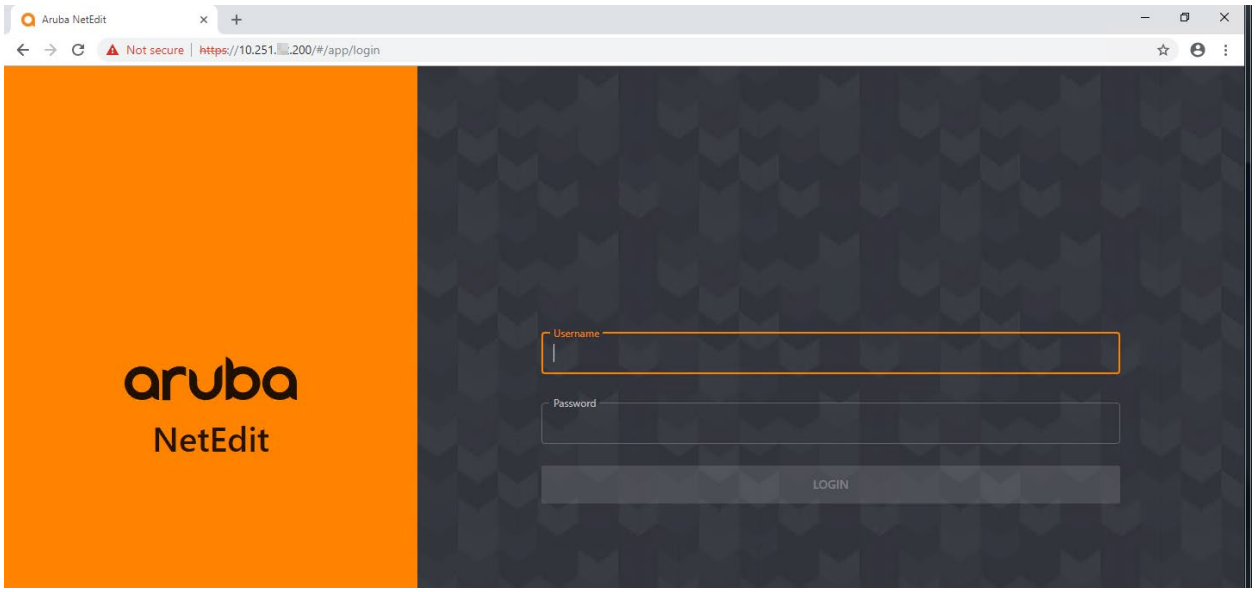

5. Log in using a username of **admin** and no password.
6. After the installation, the first time the admin user connects, a new admin password must be set. Set the password to **aruba123.**

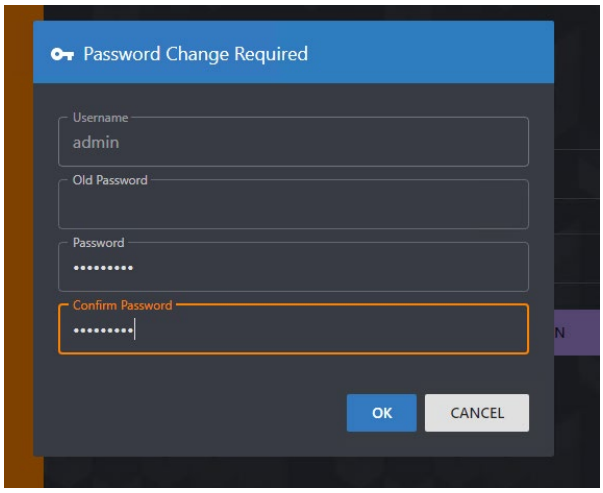

7. The NetEdit overview screen will now be displayed.

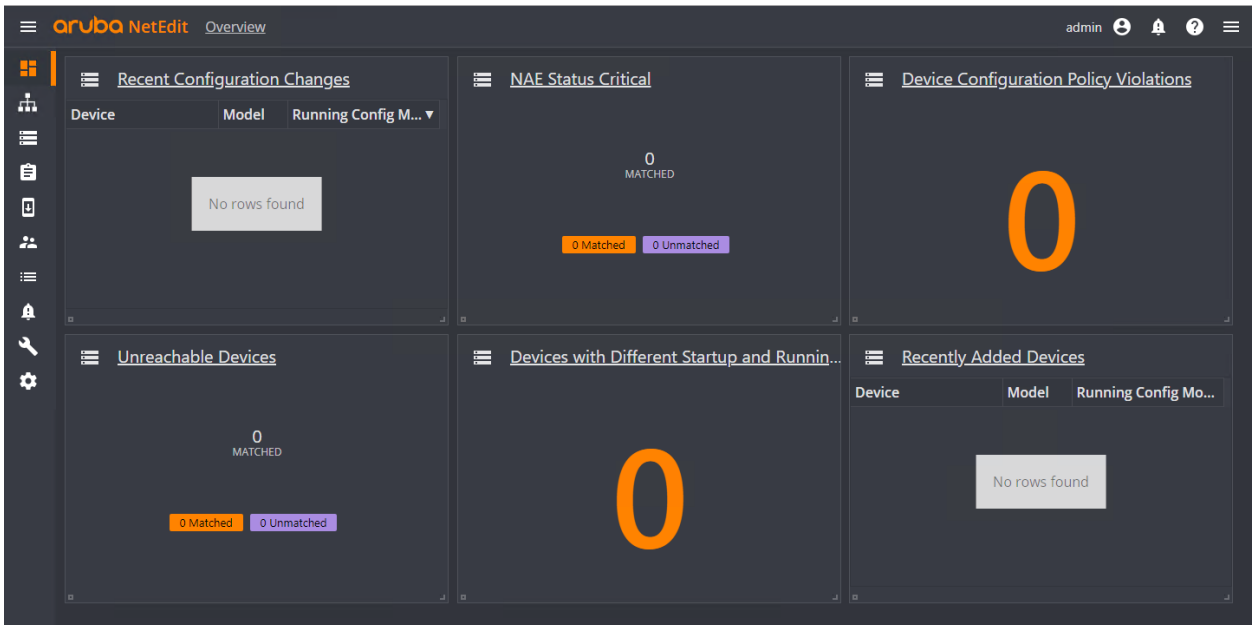

8. The display settings can be changed under **Settings > Appearance**. In the training guide screenshots, the theme has been set to **DARK**. You can use your own preferences during the labs.

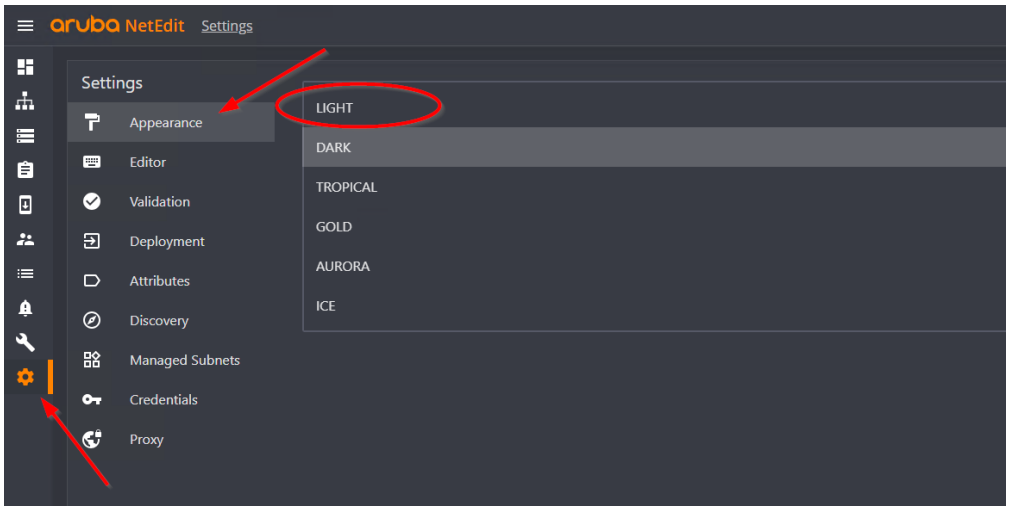

NetEdit Version and Documentation

9. Check the current version of NetEdit. Navigate to the **help** icon on the right-top.

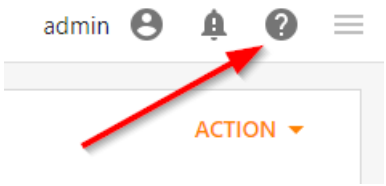

10. Select the **About** option.

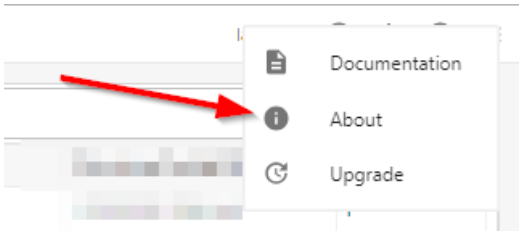

11. Review the current version of NetEdit.

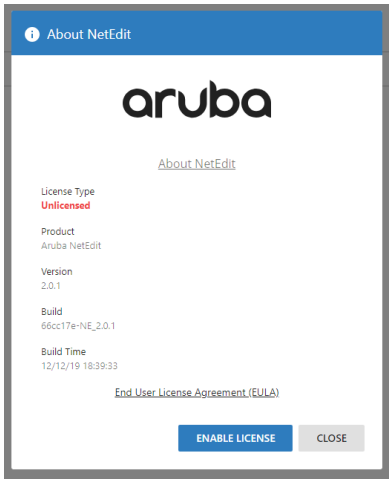

Q: Why is a license needed for NetEdit?

A: A subscription license should be installed as of the first switch (node) that NetEdit needs to manage or monitor. This is recommended for every production node in order to be able to get support for NetEdit. There is a trial version of NetEdit that can handle a maximum of 25 nodes, but there is no support for that version.

12. Close the About screen, open the **Help > Documentation** link

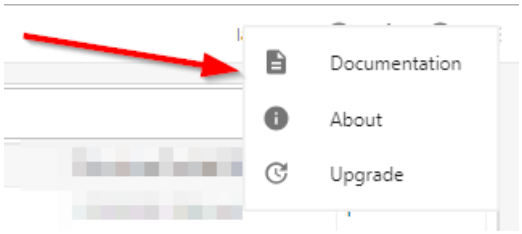

13. The complete NetEdit documentation is included in the product on this link, consult this link when you need more information about NetEdit.
14. Close the documentation webpage by closing the tab in the browser.

Task 3: Device Discovery and Device Information

Objectives

- Discover the switches of your own table
- Set access credentials
- Perform discovery
- Use a seed device address
- Review device details

Steps

NetEdit Device Credentials

1. First, NetEdit must be configured with the correct credentials so it can access the devices using the RESTAPI.
2. Navigate to **Settings > Device Credentials**.

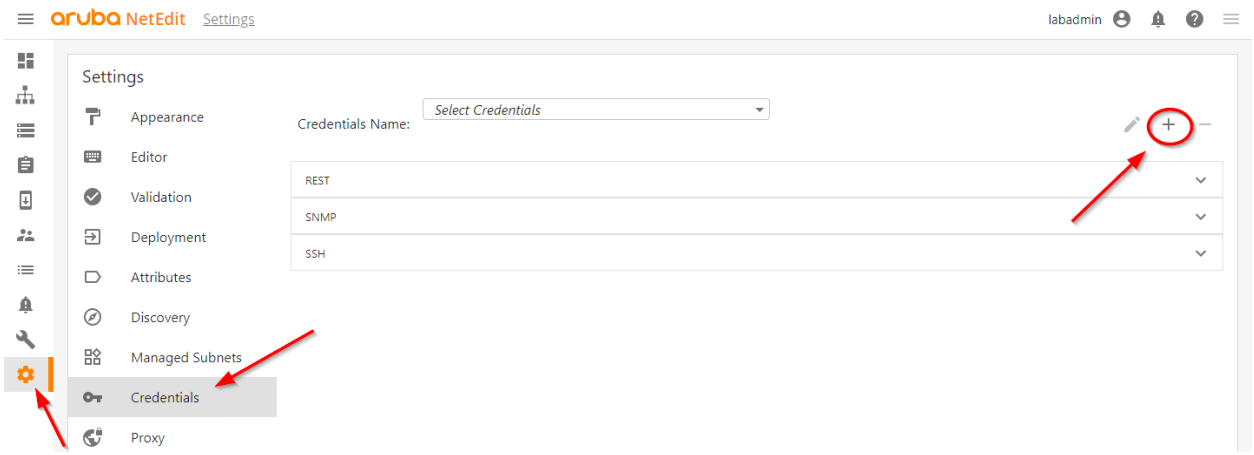

3. Multiple credentials can be used in a network, so NetEdit can be configured with multiple sets of credentials. Add a new set of credentials using these settings.

Name	icx-lab
REST	admin / aruba123
SSH	admin / aruba123

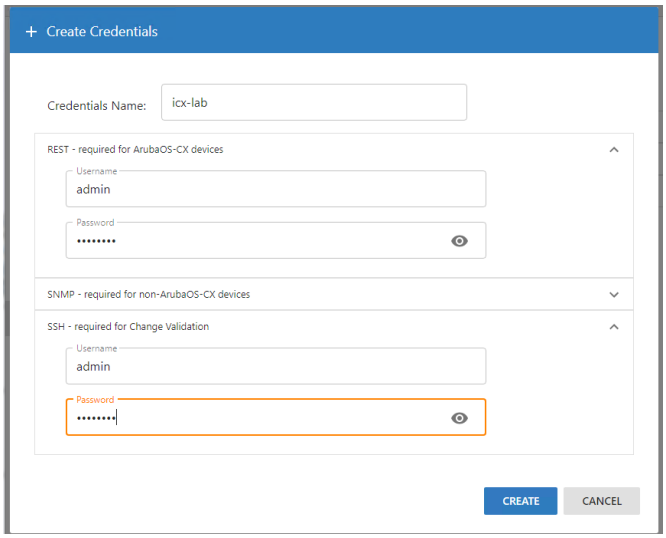

4. Click **Create** to save the new credentials.

NOTE: All configuration changes will be made using the REST API. NetEdit can also collect validation information, such as output of 'show' commands, or perform another command, such as 'ping' to some host. These validation commands are executed over SSH.

5. Once the credentials have been added, they should be visible in the credential dropdown list.

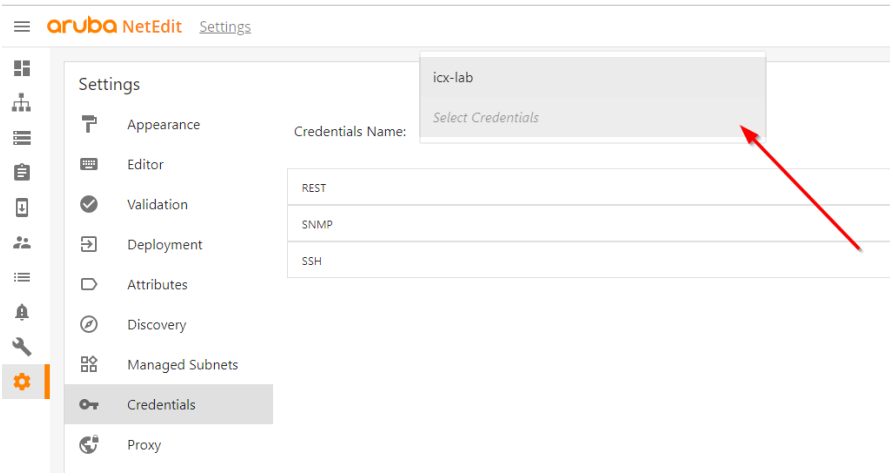

6. If a mistake was made and credentials need to be changed, the 'pen' icon on the right-top can be used.

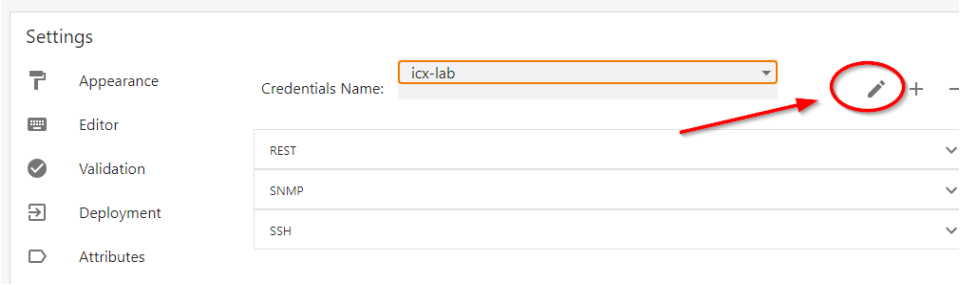

NetEdit Logging

7. Any changes made by an administrator in NetEdit, will be recorded in the NetEdit log system.
8. Navigate to **Logs** and review the changes that have been logged.

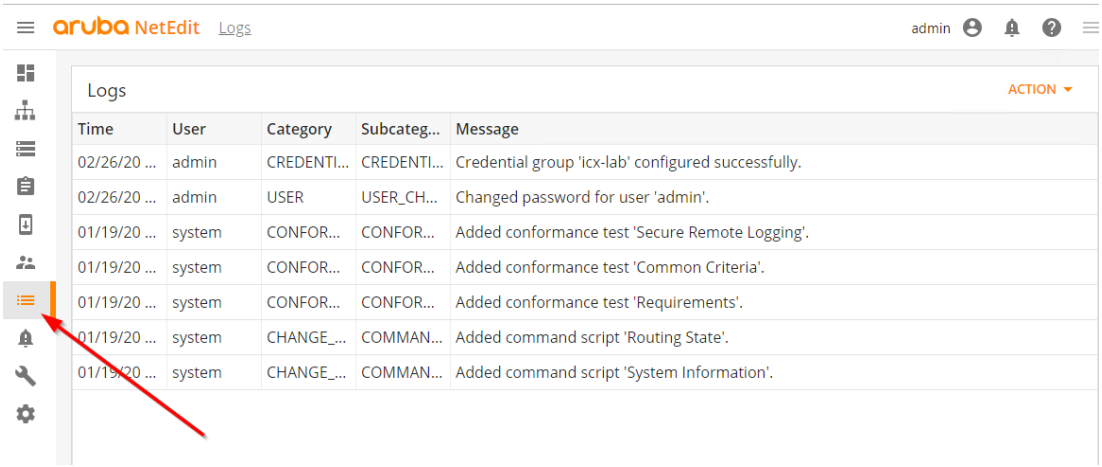

Device Discovery

In this section, the switches will be discovered and added into NetEdit.

9. Select the **Devices** section, from the **Action** menu select **Discover Devices**.

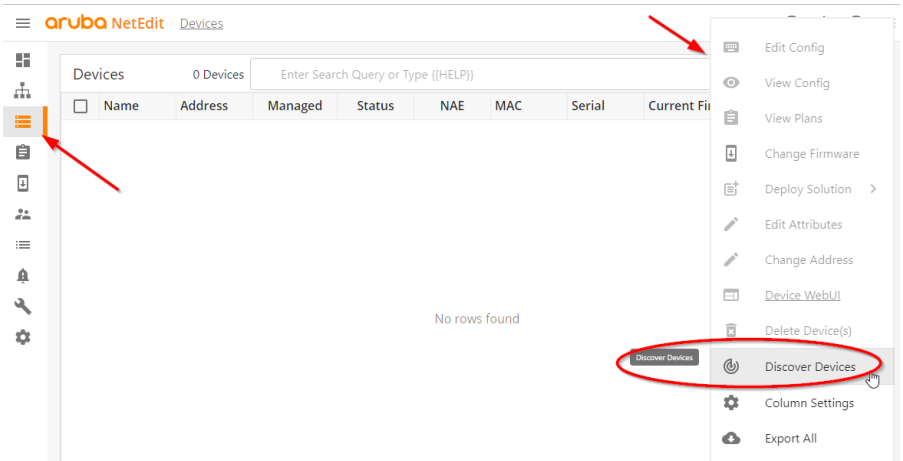

10. Add the management subnet of your table (**10.251.x.0/24**), and select the credentials that were previously defined.

11. In the seed address table, add the OOBM IP addresses of the switches of your table. You may need to scroll down to see the 'seed address' list.

> **NOTE:** In the current state of the lab, the switches do not have an 'in-band' connection, so they cannot discover other switches using LLDP yet. In a network where the switches do have in-band connection, NetEdit would automatically add newly discovered LLDP devices that have an IP address on the configured Subnet.

a. Core1 10.251.x.2
b. Core2 10.251.x.3
c. Access1 10.251.x.4
d. Access2 10.251.x.5

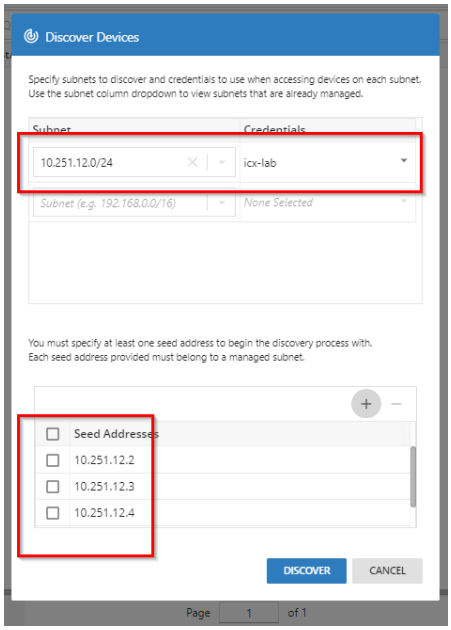

12. Select the checkbox for each Seed Address and click **Discover** to start the discovery.

13. Wait a few moments, the 4 switches should appear in the list.

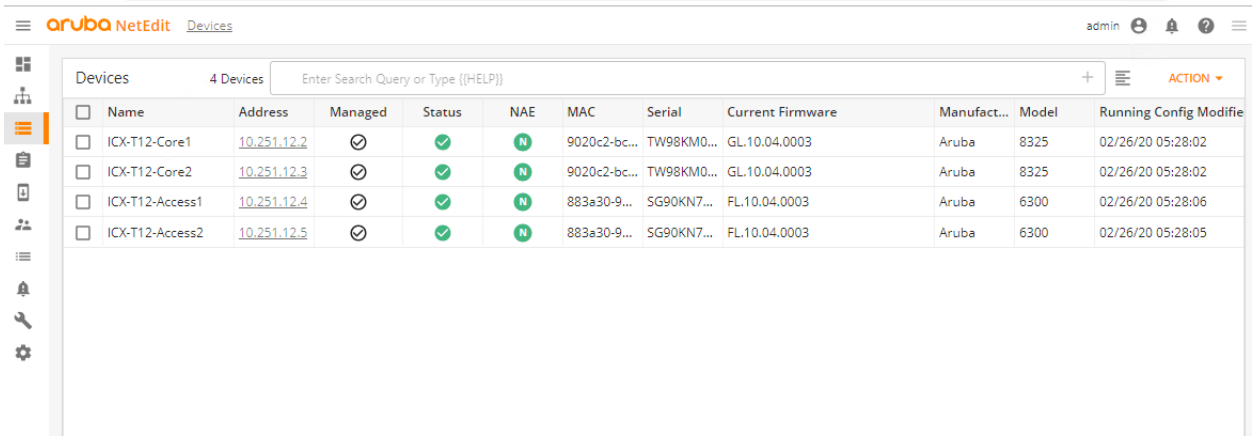

IMPORTANT: Make sure all four switches are added to NetEdit at this point. In case a device cannot be added, check:

- IP address of the OOBM interface on the devices (interface mgmt)
- admin password
- REST API access was enabled for the 8325 switches

Refer to lab 01 for the steps.

Device Details

14. On the **Devices** page, click on the IP address of one of the switches to see the details of the device.

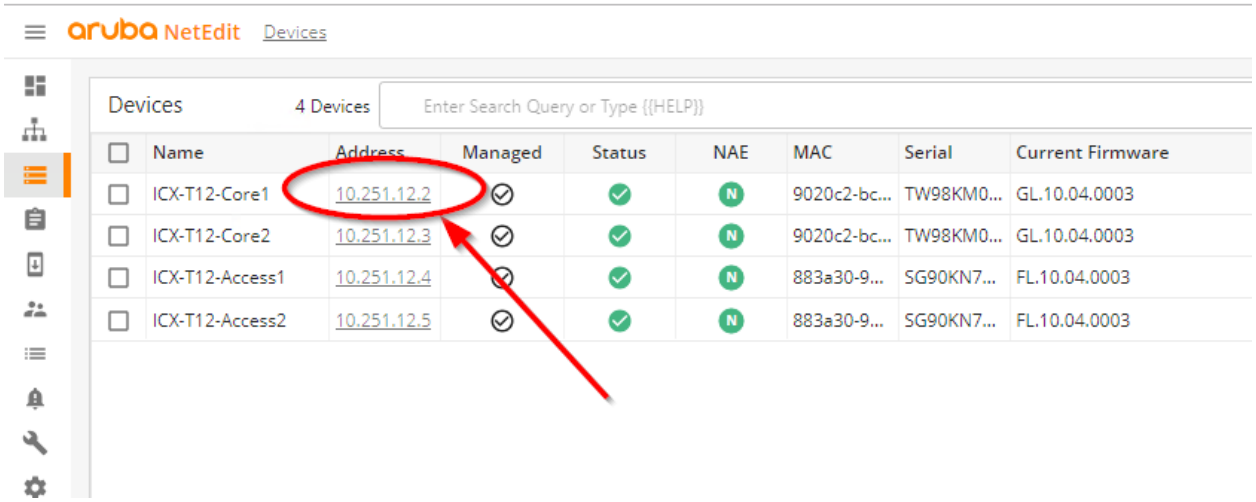

In the details screen, the MAC and serial of the device are listed.

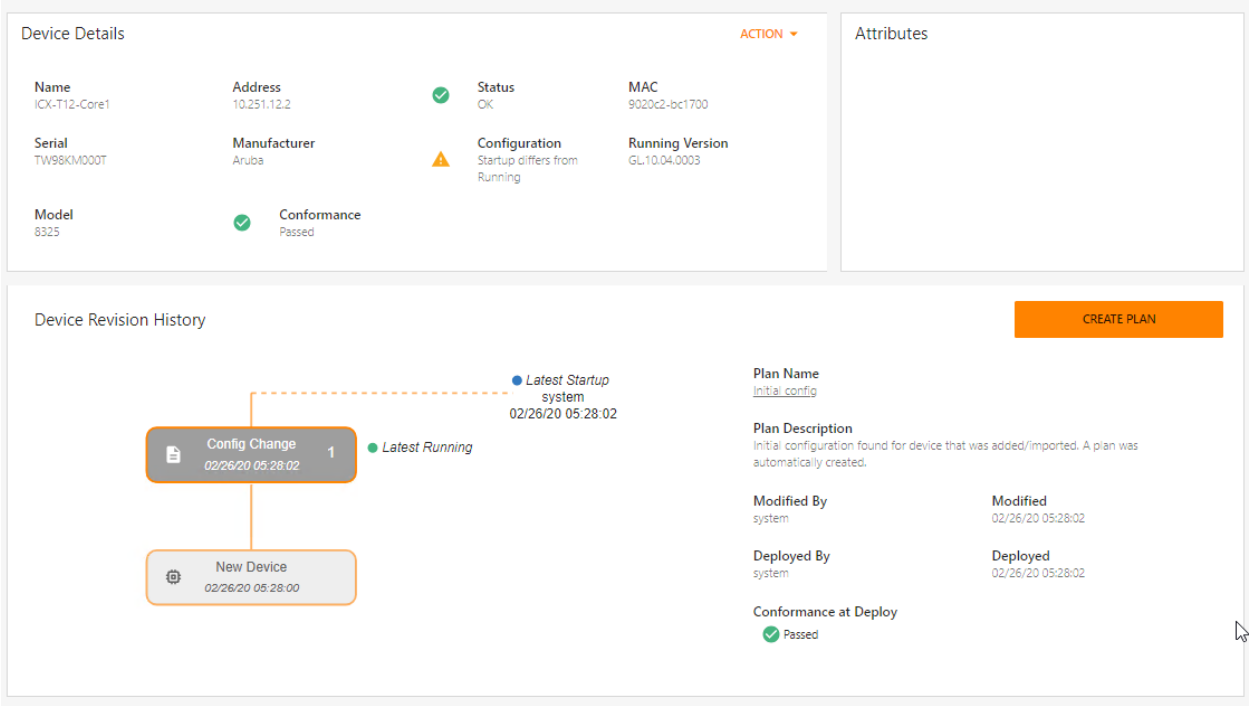

15. The **Action** menu provides the administrator with access to the 'running' and 'startup' configuration, 'firmware' and 'hardware' information and the option to reboot the device.

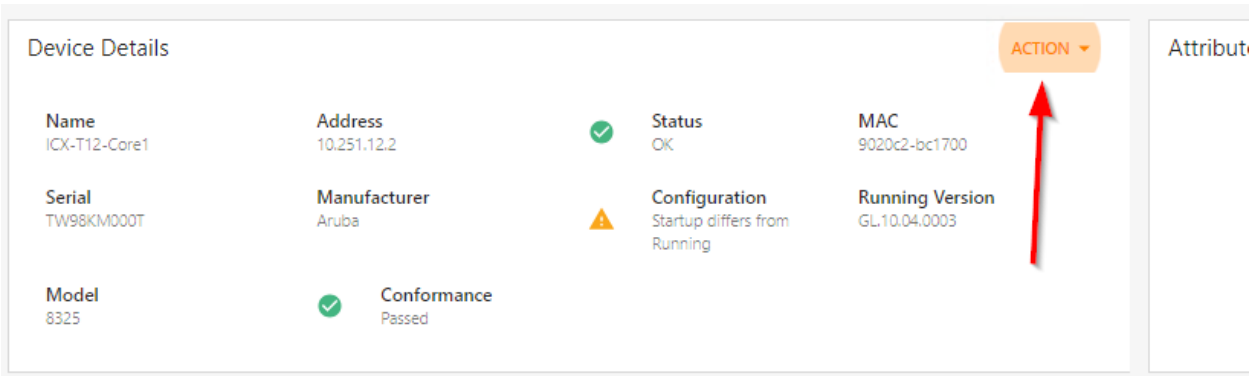

Action list example:

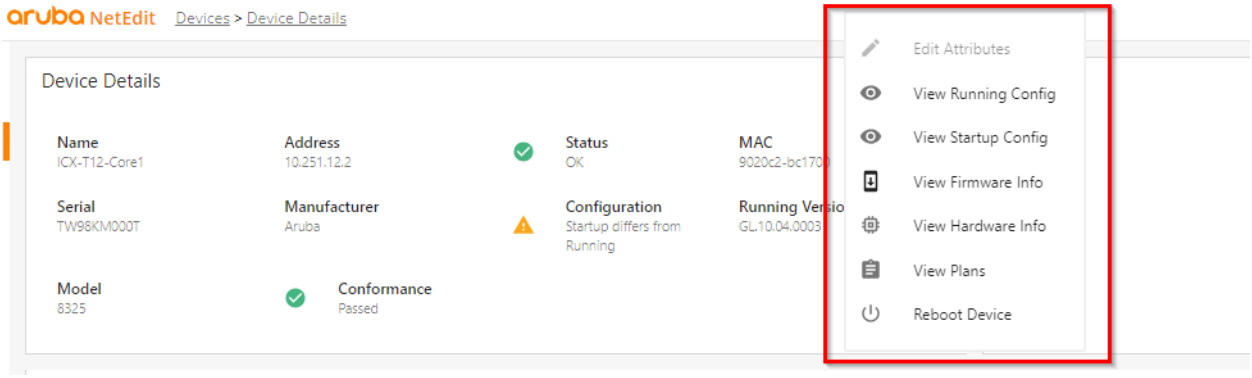
aruba NetEdit
Devices > Device Details
Device Details
Name
ICX-T12-Core1
Address
10.251.12.2
Status
OK
MAC
9020c2-bc170
Serial
TW98KM000T
Manufacturer
Aruba
Configuration
Startup differs from Running
Running Versio
GL.10.04.0003
Model
8325
Conformance
Passed
Edit Attributes
View Running Config
View Startup Config
View Firmware Info
View Hardware Info
View Plans
Reboot Device

Task 4: Configuration Management Using Plans

Diagram

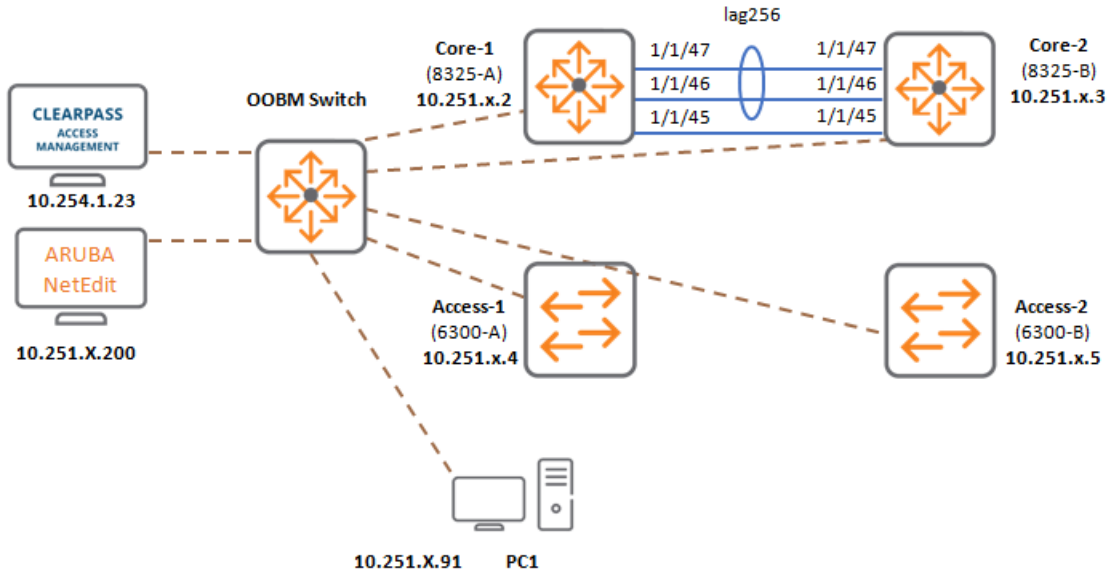

Objectives

- Understand how a configuration plan works
- Deploy a new configuration
- Verify the new configuration
- Push configuration changes to multiple devices

Steps

First Configuration Plan

A configuration plan allows the administrator to send configuration changes to one or more AOS-CX switches.

In these steps, a new VRF will be created on both Core switches a port will be assigned, and an IP address will be set. This VRF configuration will be used in the upcoming VSX lab as the 'keep-alive' link between the VSX core switches.

> **NOTE:** More details on VSX and the 'keep-alive' link are covered in the Module 4.

1. On the management PC, login to NetEdit using the admin account.
2. Navigate to **Devices**, select the **Core1** and **Core2**, and right-click to select **Edit config**.

> **NOTE:** Many lab activities show how **multiple** devices can be configured together. In case the administrator needs to edit a **single** device, simply select one device in the

list and use the 'Edit Config' option.

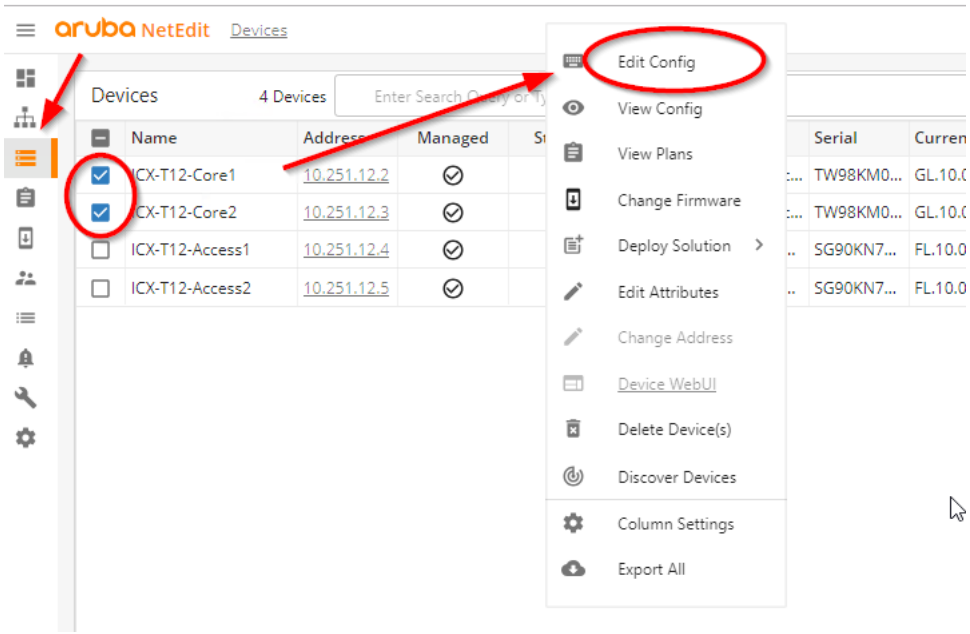

3. Provide a meaningful name for the plan, for example **'vrf ka'**. Click **Create** to define the plan. ('ka' stands for keep-alive.) Depending on the version of NetEdit you may have to click **Return to Plan>Action>Edit Plan & Attributes** to assign a Plan Name. If you skip this step, it will make Task 5 more difficult.

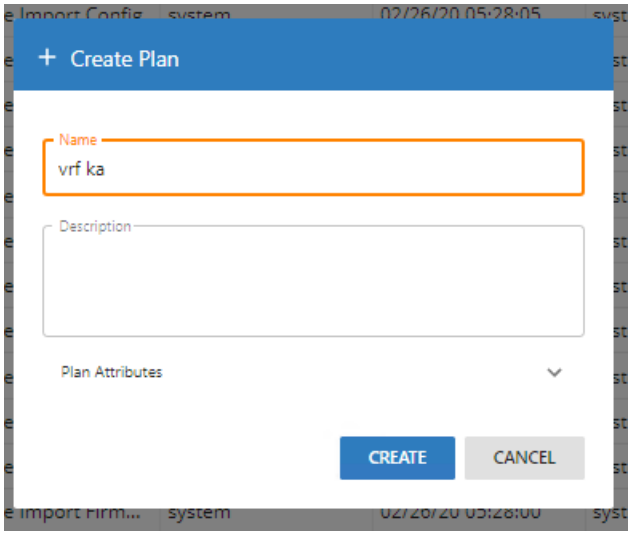

4. NetEdit will now show the **merged** 'running configuration' of the devices.

Any lines that exist on all selected devices will show normally.

Any lines that have unique values, will include some '**PARAMETER**' placeholder in **UPPERCASE**.

5. Move the mouse over the management ip address parameter, shown as **A.B.C.D/M**. NetEdit will show the configured IP addresses per device .

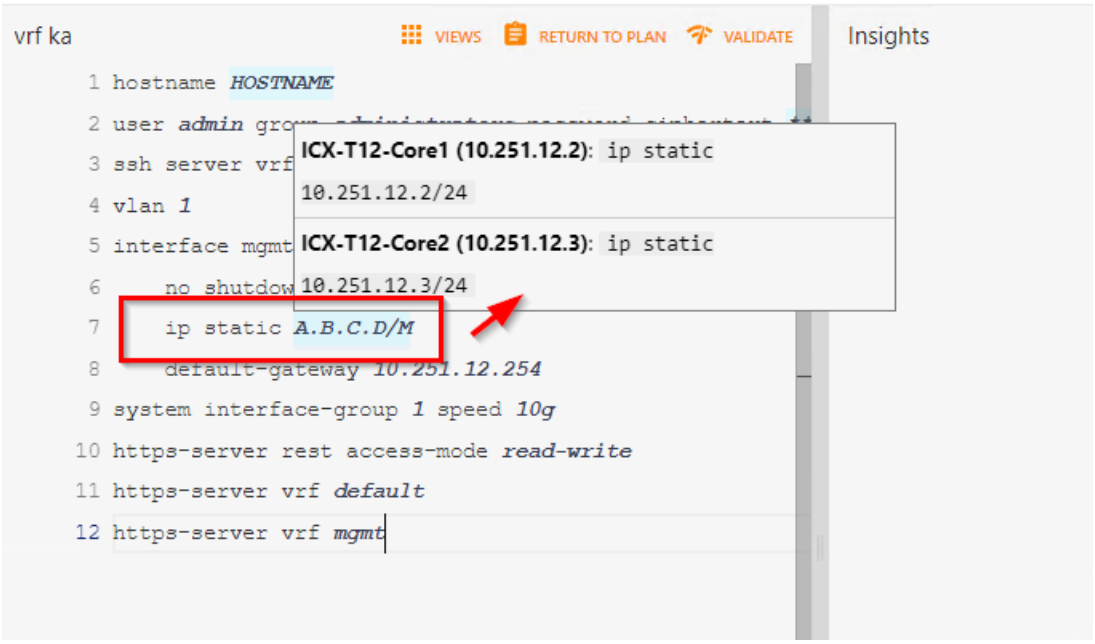

6. The same applies to the **HOSTNAME** parameter.

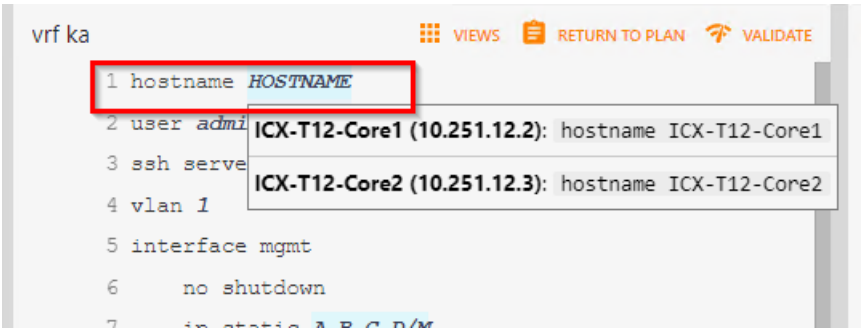

7. Now right-click on the **'HOSTNAME'** parameter, this window shows the actual hostname configuration for each device and allows the administrator to change the values.

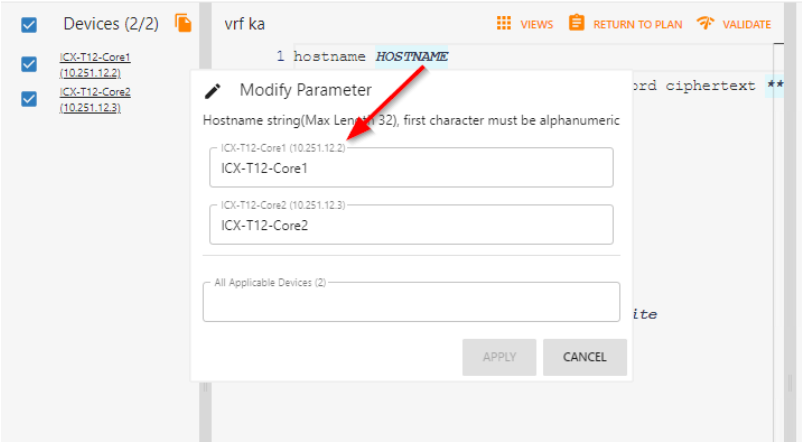

8. Click **Cancel** to close the 'Modify Parameter' window.

Add Configuration Entry

9. Click at the end of the configuration and press **Enter** to get a new line.
10. Start typing '**v**' and notice how the context sensitive help shows the available options. Next type '**r**' (so '**vr**') Use <TAB> to complete the word to **vrf**'.

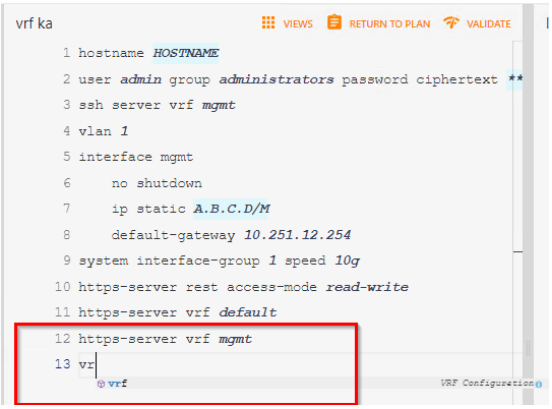

11. Next press <SPACE> and notice the available options for the vrf command.

12. Enter 'KA' and use <ENTER> to submit the command. Notice how NetEdit will automatically move the line to the correct location in the configuration.

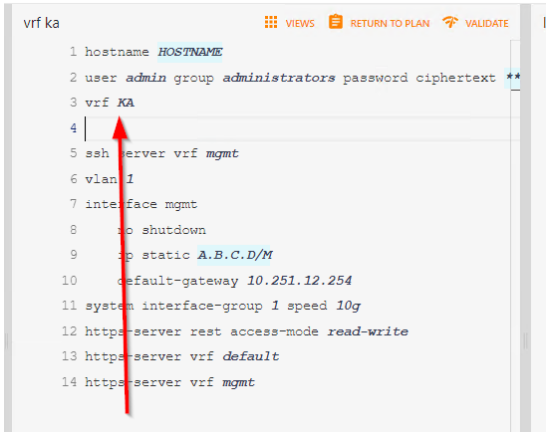

Add Interface Configuration

13. At the end of the configuration, add a line for 'interface 1/1/45'. (int<TAB> 1/1/45). Press <ENTER> so NetEdit moves the line to the correct place in the configuration.

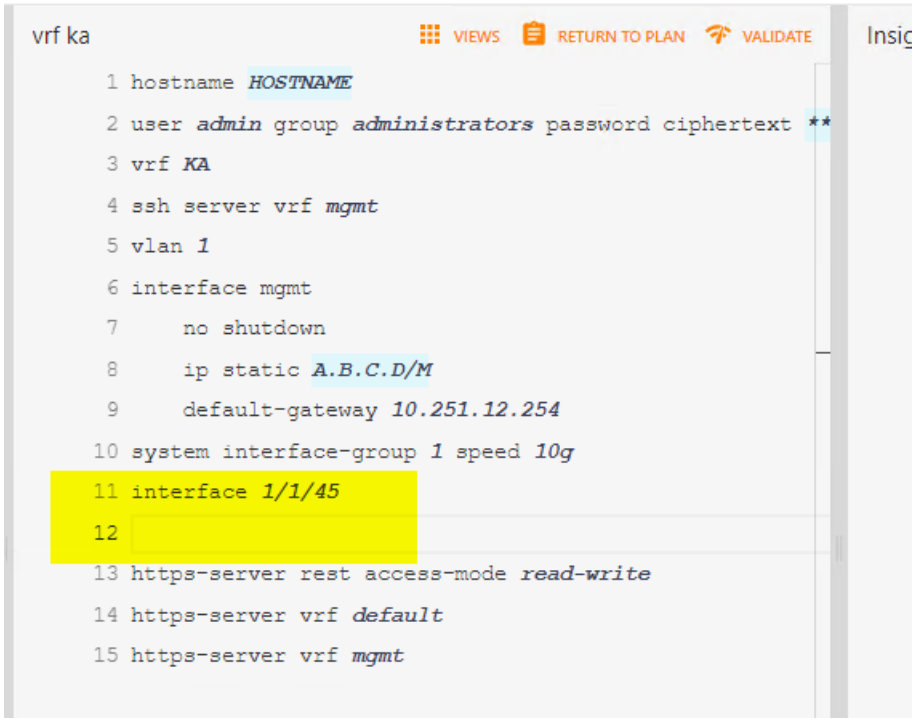

14. On the next line, press <SPACE> to enter the 'interface' context.

IMPORTANT:
The <SPACE> is important here to tell NetEdit that the next command is not a global command, but should be under the interface context. A switch CLI would put you in the 'interface' context, while in the NetEdit CLI you are still at the 'global' level.

15. Enter the command '**vrf attach KA**', press <ENTER> .
16. Enter the command **'no shutdown'**, press <ENTER> .
17. Enter the command '**ip address 192.168.0.0/31**', press <ENTER>.

> **NOTE:** AOS-CX allows you to configure /31 as the subnet mask for an IP interface. For point to point connections (only 2 hosts on the subnet), this is convenient, since you are not wasting the network and broadcast address.

```
10 system interface-group 1 speed 10g
11 interface 1/1/45
12     vrf attach KA
13     no shutdown
14     ip address 192.168.0.0/31
15
16 https-server rest access-mode read-write
```

18. Move the mouse over the '192.168.0.0/31' IP address, right-click to open the 'Modify Parameter' window. Change the **'Core2'** IP to **192.168.0.1/31**, that is the second IP address in this /31 subnet. Click on **Apply** to submit the change.

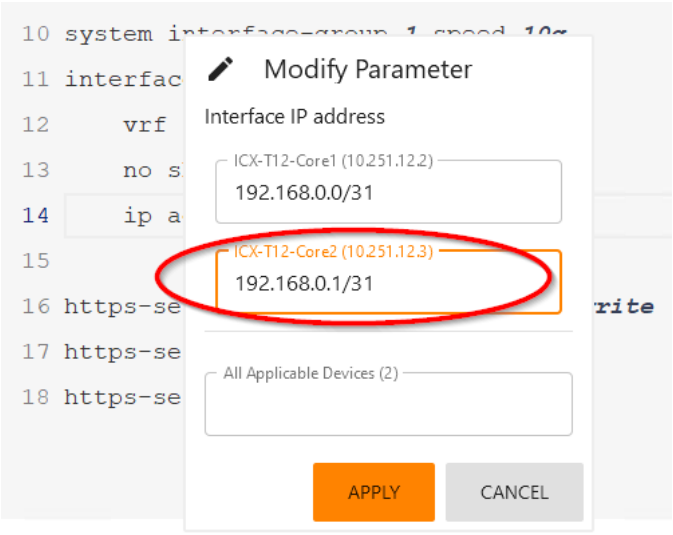

19. Notice how the IP address parameter is now shown as A.B.C.D/M. This is sufficient for the first change, select the **Return to Plan** option to return back to the plan details screen.

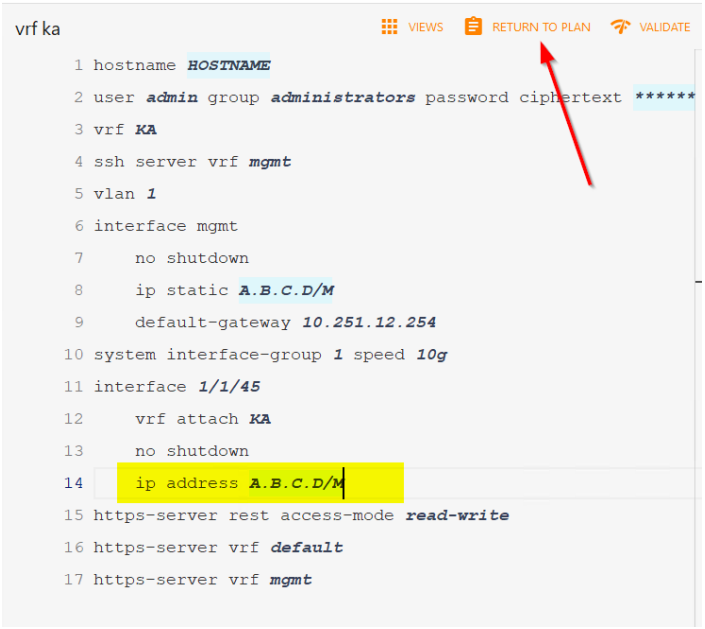

20. To push the configuration change to the device, use the **Deploy** option.

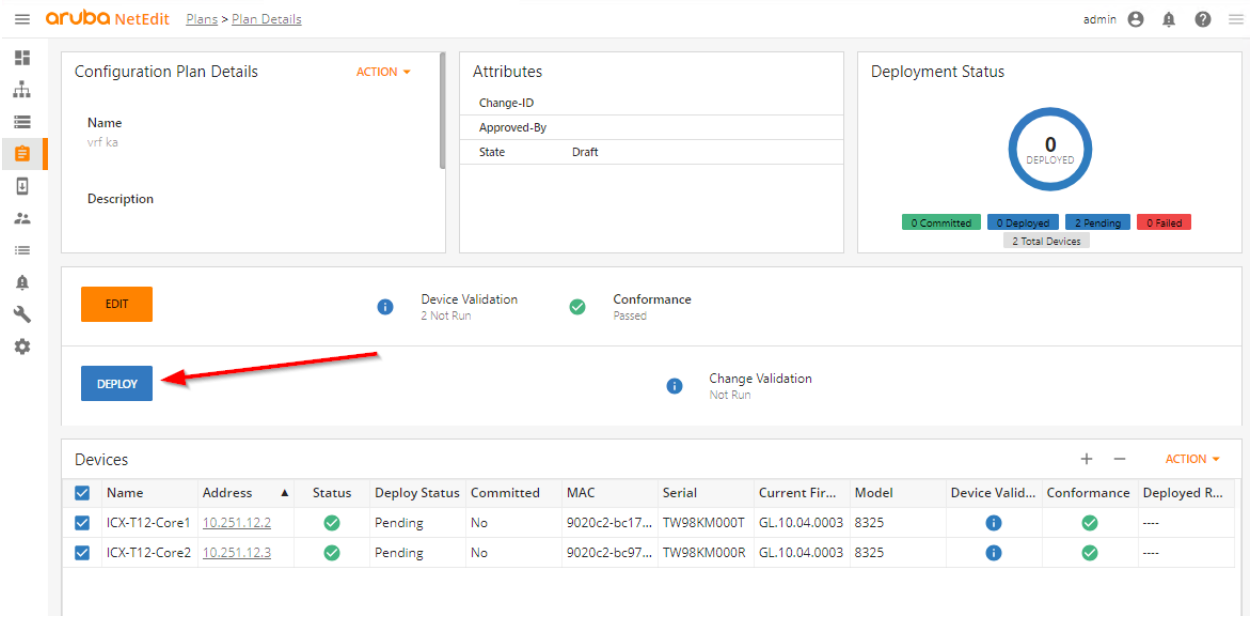

21. Confirm the action dialog with the **Deploy** button.

Q: What happens when the plan is being deployed?

A: The proposed changes are validated with device, after successful validation, this configuration will be activated. It will be in the running configuration of the device.

22. After a few moments, NetEdit will show the 'Commit' and 'Rollback' options. This allows you to either save the configuration changes or undo these changes by reverting to the previous configuration.
23. Verify the changes by clicking the **Change Validation** option.

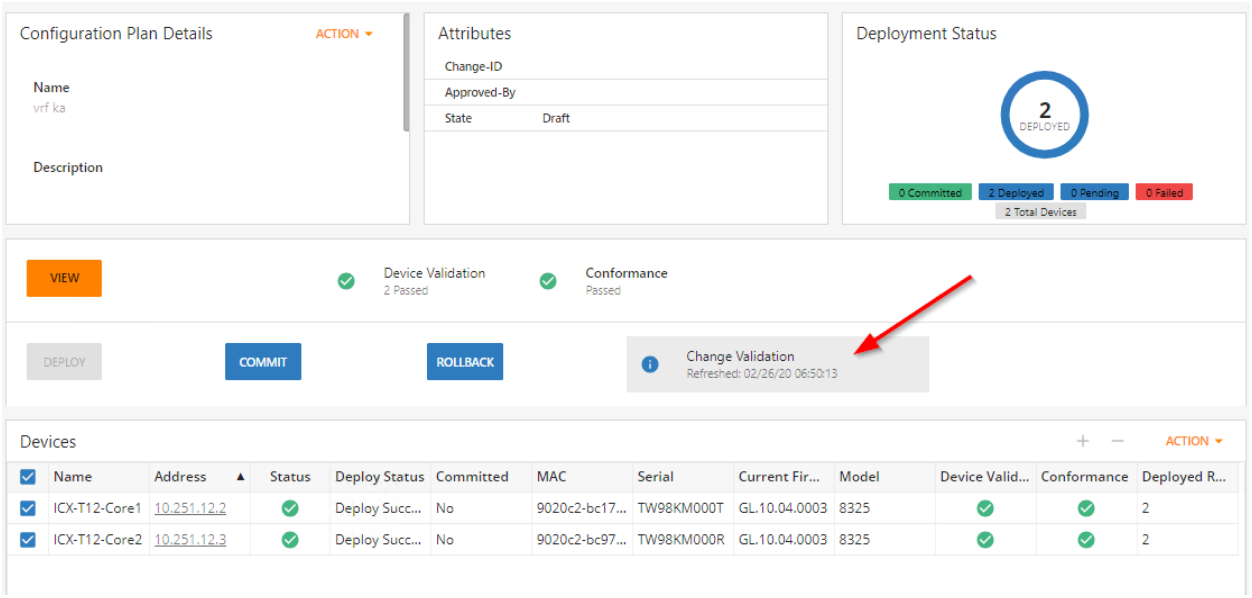

24. Click on the '**>**' icon for the line ICX-Tx-Core1 … 'show running-config' to see the details.

On the left side, the output before the change is shown, on the right side, the output after the change is shown, this should show the configuration changes.

Scroll down in the configuration verify that the vrf KA has been defined, and port 1/1/45 has an IP address.

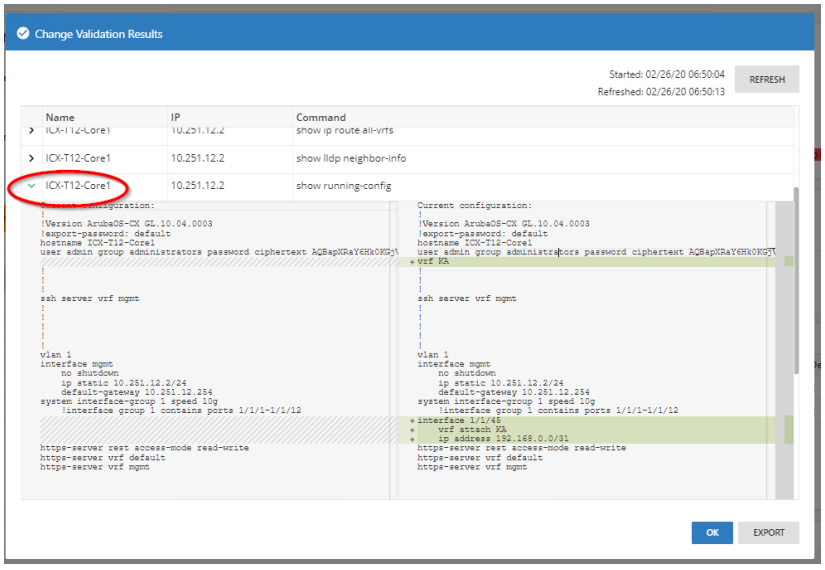

Verify Deployment and Commit the Configuration

25. Click **OK** to return to the management station, in NetEdit, click the **Commit** button to save the configuration on the device.

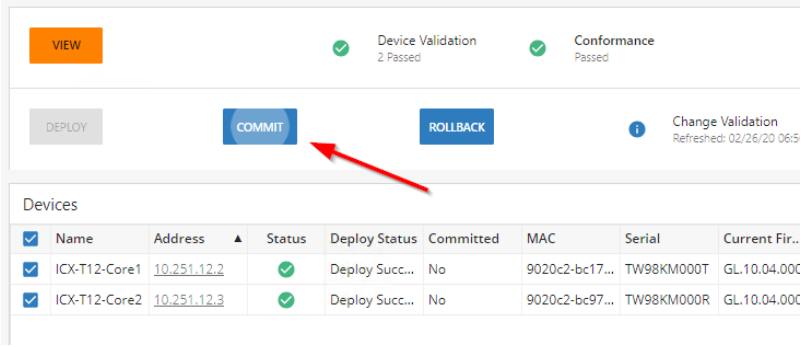

26. NetEdit will prompt a confirmation message, confirm by pressing the **Commit** button.
27. **Optional step**: on the Core1 terminal, review the startup config to verify that the NetEdit configuration was actually pushed to the startup-configuration of the switch.

```
show startup-config
```

Add DNS Configuration using 'Deploy Solution'

This option of NetEdit provides a wizard-like experience to configure essentials settings on the devices, such as DNS or NTP. In these steps, the NTP and DNS will be configured on **all four switches**.

28. Open NetEdit, navigate to **Devices**.
29. Select all 4 devices, from the **Action** menu, select **Deploy Solution** and pick **DNS Management** from the list..

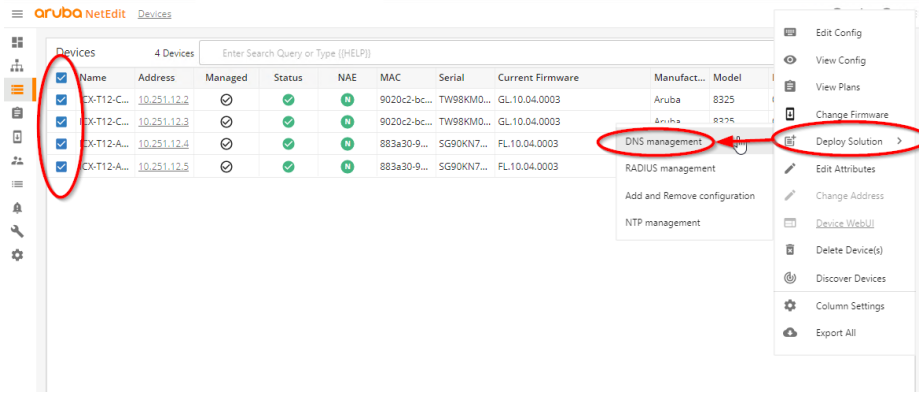

30. In the **Deploy Solution screen**, enter these settings:

DNS server IPs	10.254.1.21
Domain List	(leave blank > no value)
VRF	mgmt

31. Click **Create.** It will take a few moments to prepare the plan.
32. In the next window, click **Deploy**, confirm with **Yes**.

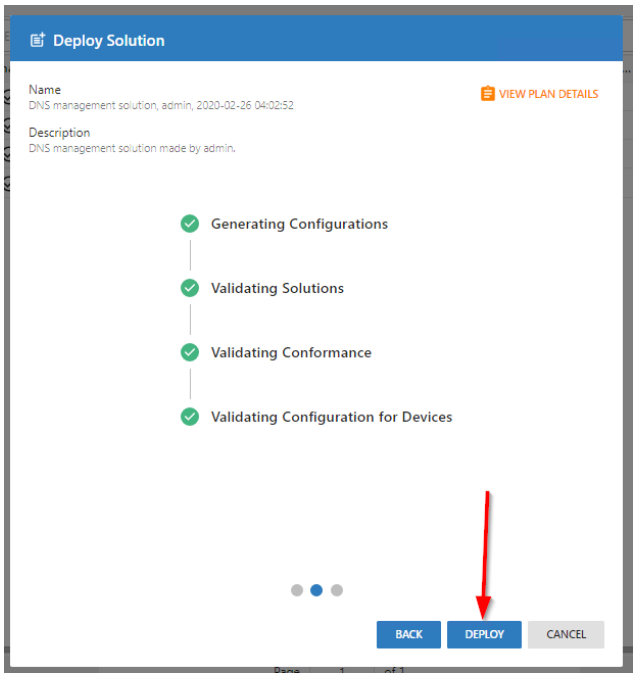

33. Click **Close** to close the window.

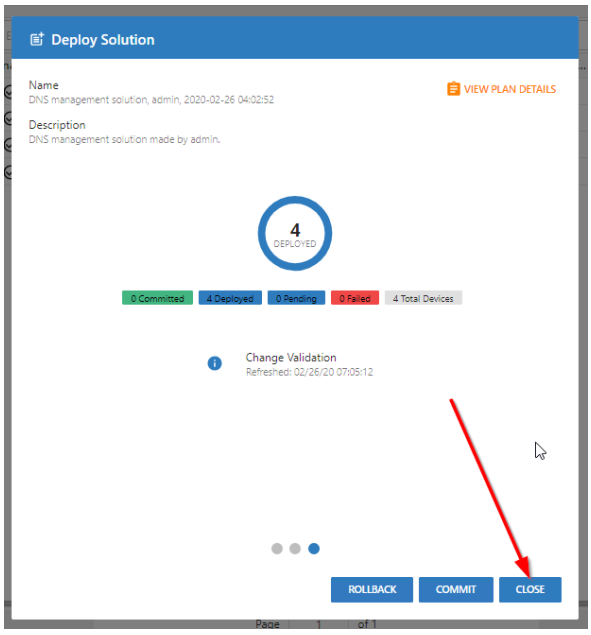

Add NTP Configuration Using 'Deploy Solution'

Use the next deployment solution to configure NTP.

34. Under NetEdit > **Devices**, select all four devices.
35. Navigate to the **Action** menu, select **Deploy Solution** and pick **NTP Management** from the list.
36. In the **Deploy Solution screen**, enter these settings:

Preferred Servers	**10.253.1.15**
Additional NTP servers	**10.253.1.15**
Burst mode	**iburst**
VRF	**mgmt**

Leave all other settings default.

37. Click **Create.**
38. In the next screen, click **Deploy**; confirm with **Yes.**
39. After a few moments, the deployment should complete. Close the window by selecting the **Close** button.

Second Configuration Plan - Multiple Devices

In this section, the configuration of the two core switches will be updated.

- Configure an IP address on VLAN 1 on both Core switches.
- Complete the time configuration by configuring the time zone on the switches.
- Define a Link Aggregation between both Core Switches .

40. In NetEdit, navigate to **Devices** and select both Core switches.
41. Next click on **Action** and select **Edit Config**.
42. Provide a name for the config: **'core interswitch link'** and click **Create**.

Configure the VLAN 1 IP address for both Core switches.

43. At the end of the config, press <ENTER> to get a new line, define the VLAN 1 L3 interface and press <ENTER>. NetEdit will automatically move the 'interface VLAN 1' to the right place in the configuration.

```
interface vlan 1
```

44. Press <SPACE> to enter the interface context and assign the IP address of Core1. The IP for Core2 will be set in the next step.

```
 ip address 10.x.1.2/24
```

NOTE: Make sure to include a SPACE at the beginning of the line after the 'interface vlan 1' to be within the interface context.

NOTE: Make sure to press <ENTER> after entering the IP address. The IP address should become highlighted in blue, that indicates NetEdit has recognized it as an IP address and gives you the multi-edit 'right-click' option.

45. Right-click on the VLAN 1 IP address to access the 'Modify Parameter' window. Assign IP address **10.x.1.3/24** to the **'Core2'**. Click **Apply** to confirm.

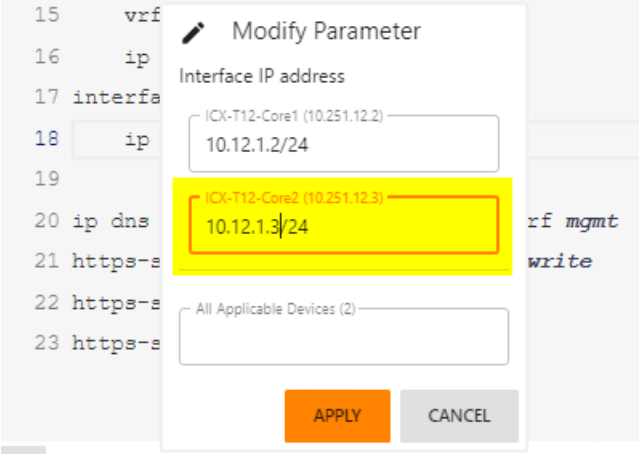

46. Move the mouse over the A.B.C.D/M parameter of VLAN 1 to verify the IP address of both Core switches.

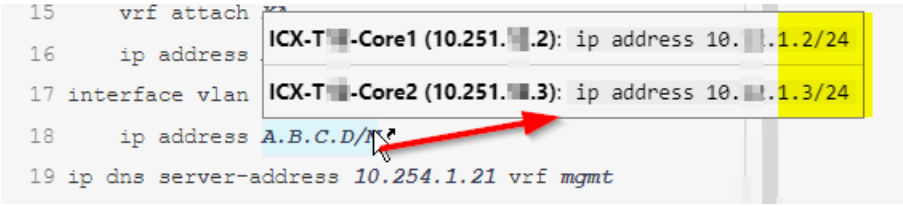

Define a Link Aggregation with two Interfaces Between the two Core Switches

47. At the end of the configuration, insert a new line for '**interface lag 256**'. Press <ENTER> after the line so NetEdit can move it to the correct location.

```
interface lag 256
```

48. Press <SPACE> to enter the LAG context, make the LAG:

- enabled (no shutdown)
- switched port (no routing)
- VLAN trunk and allow all VLANs
- LACP active

```
no shutdown
no routing
vlan trunk allowed all
lacp mode active
```

49. Next configure two ports as member ports of the LAG, and enable the ports.

NOTE: Make sure to keep using the <SPACE> to enter the interface context.

```
interface 1/1/46
    no shutdown
    lag 256
interface 1/1/47
    no shutdown
    lag 256
```

50. At the top of the screen, click **Validate** to check if all of the commands entered are correct. If the validation failed, review and correct your input.

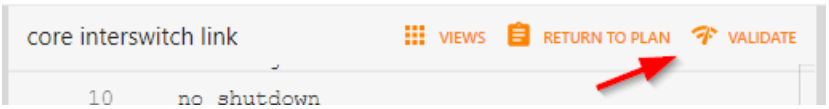

51. Return to the Plan using the **Return to Plan** button
52. In the Plan Details screen, use **Deploy** to push the changes to the 2 Core switches. Click on **Deploy** to confirm.
53. Once the deployment has completed, click **Commit** to save the configuration on the devices, confirm by clicking **Commit** again.

Optional Steps: Verification

These steps are *optional* and can be done if time permits. Check with your instructor.

In the next steps, the changes made by the plan will be verified on the switches.

The verification steps will check:

- VLAN 1 IP address
- VRF KA configuration
- Link aggregation status

54. **Optional Step**: On the terminal of Core1, verify the link aggregation interface status to Core2 has state 'Up'.

```
ICX-Tx-Core1# 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     ASFNCD  90:20:c2:bc:17:00 65534  256  up
```

```
1/1/47  lag256     48    1     ASFNCD  90:20:c2:bc:17:00 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     ASFNCD  90:20:c2:bc:97:00 65534  256
1/1/47  lag256     48    1     ASFNCD  90:20:c2:bc:97:00 65534  256
```

End of optional steps.

Third Configuration Plan - All Devices

In these steps, you will push some final management settings to all four devices.

55. In NetEdit, navigate to **Devices**, select all four devices (Core1/Core2/Access1/Access2).
56. From the **Action** menu, select **Edit Config.**
57. For the plan name, enter '**management'.** Then click **Create**. Depending on the version of NetEdit you may have to click **Return to Plan>Action>Edit Plan & Attributes** to assign a Plan Name. If you skip this step, it will make Task 5 more difficult.

    NOTE: At this point, the configurations of the core and access switches contain several differences. In the plan editor, lines that are unique to some devices will be marked with the hostnames of those devices to handle this.

58. Add these lines to the configuration. This will apply the correct time zone for the remote lab and it will increase the admin session timeout to 12 hours.

```
clock timezone us/eastern
cli-session
   timeout 43200
```

59. Select **Return to plan**, click **Deploy,** and confirm with **Deploy**.
60. Once the plan completes the deployment, click **Commit** and confirm by clicking **Commit** again.

Task 5: Configuration Plan Validation

Objectives

- Understand plan validation
- Review plan validation settings
- Define new validation settings

Steps

In the previous task some configuration changes were pushed to the switches and these changes could be manually verified on the switch cli (optional step at the end of the previous task).

In this task, the same validation checks will be included in NetEdit.

This means NetEdit will be connecting to the devices using SSH and will collect the output of the defined validation commands.

Review Default Plan Validation

1. Open NetEdit, navigate to **Plans**, open the plan called **'core interswitch link'** (click this name).
2. Click the link **Change Validation** to see the default validation steps.

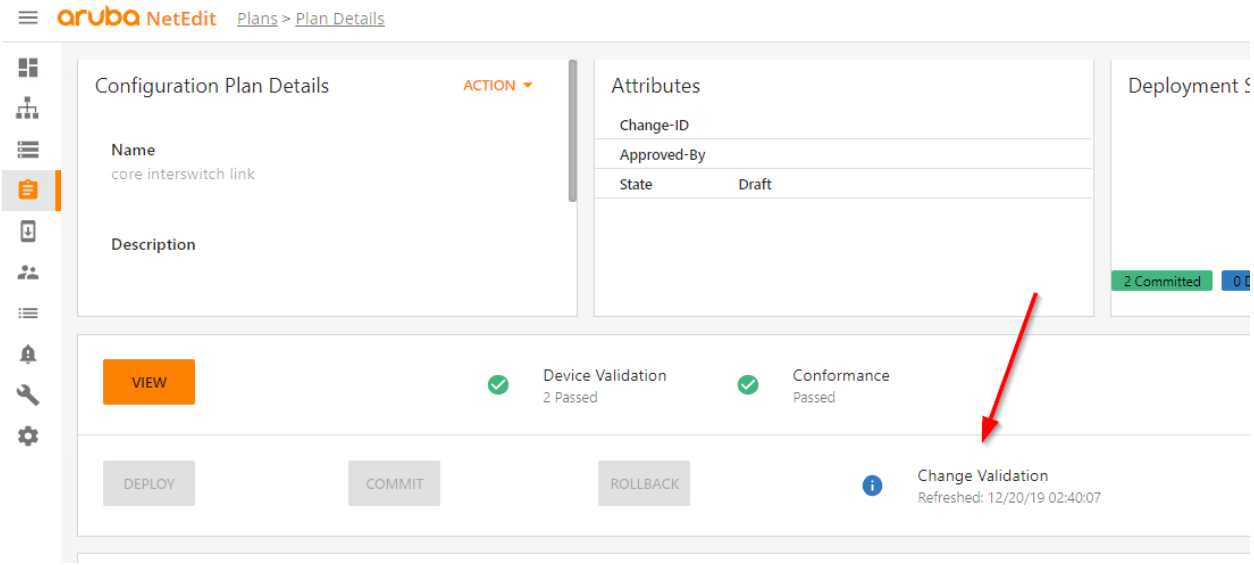

3. The list of commands will be shown here. For each command, NetEdit will collect the data **before** and **after** the configuration change has been made (**Deploy**). When any changes are found, the entry will be marked in green.

In this lab, the 'show bgp all-vrf all summary' command will have no changes, so these remain 'black'. The 'show interface brief' will be marked green.

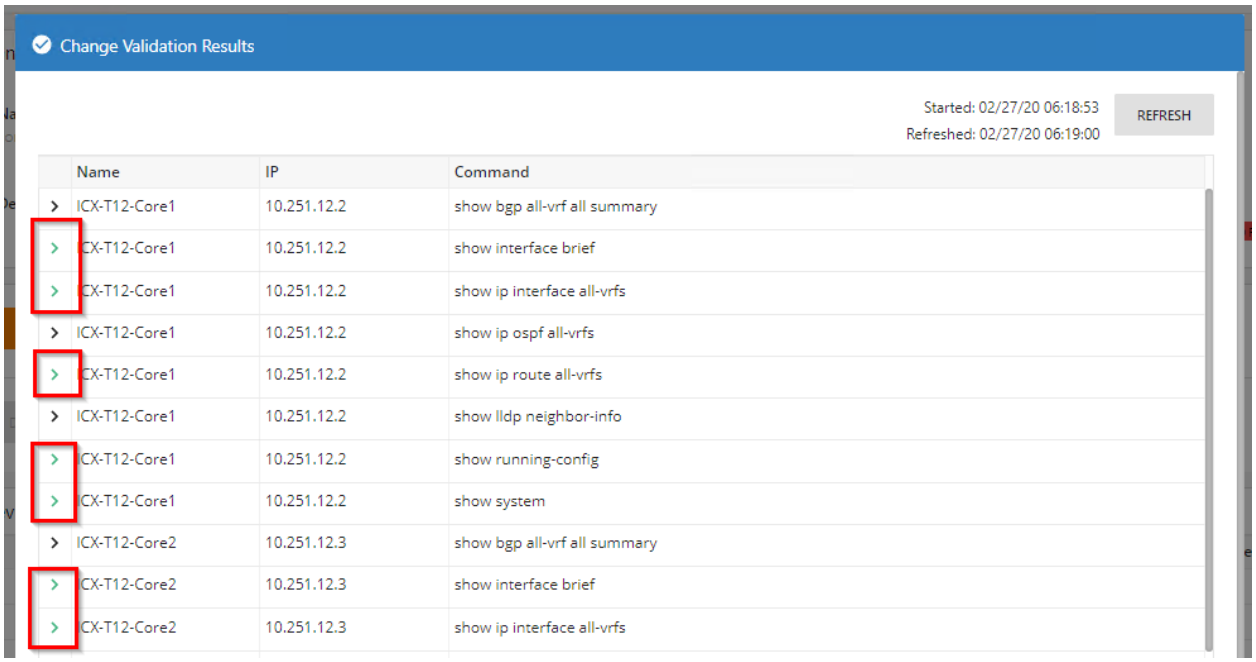

4. Expand the ICX-Tx-Core1 - 'show interface brief' section. The left side will show the before state, on the right side, the state after the configuration change will be shown.

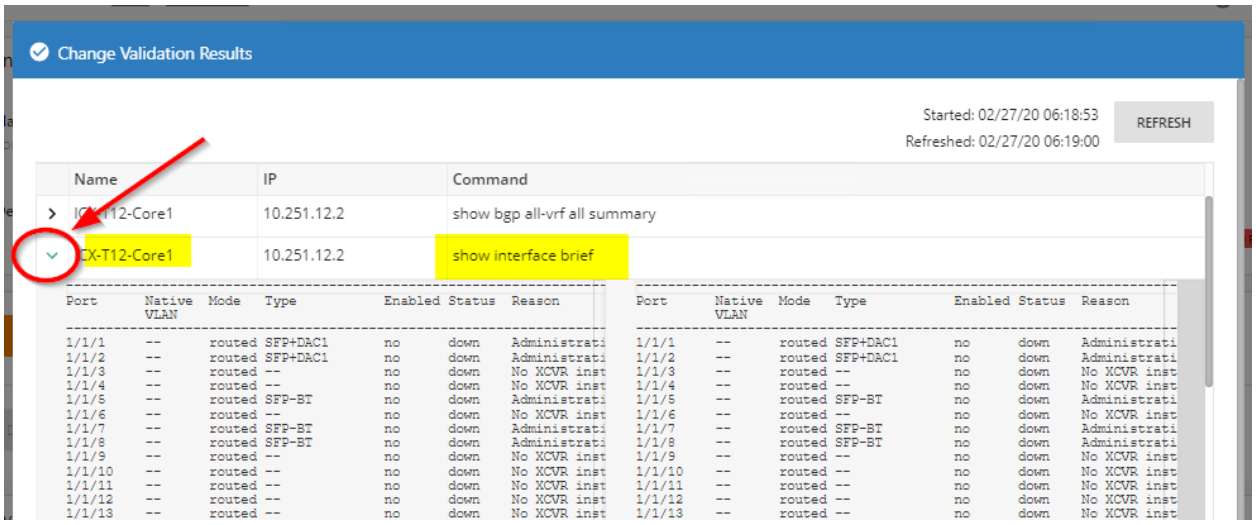

Scroll down to check the state of the interfaces. You should see the two interfaces coming UP, and the new VLAN 1 and LAG interface.

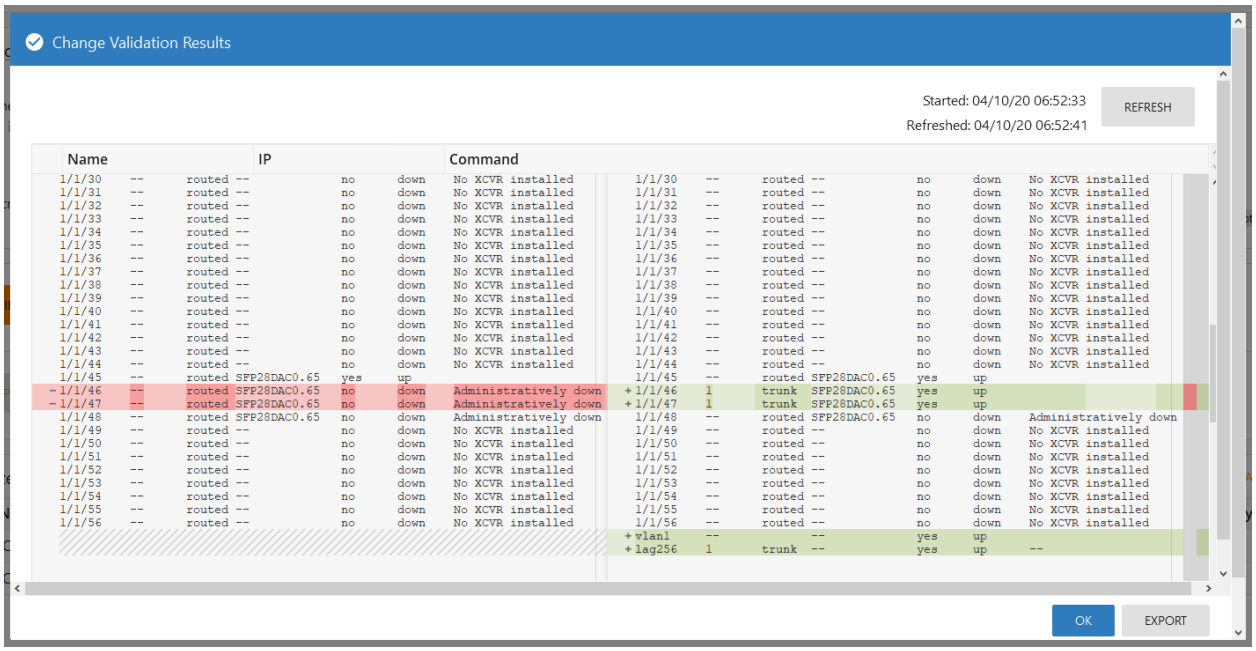

5. Review the LLDP neighbors by clicking on **show lldp neighbor-info** link.

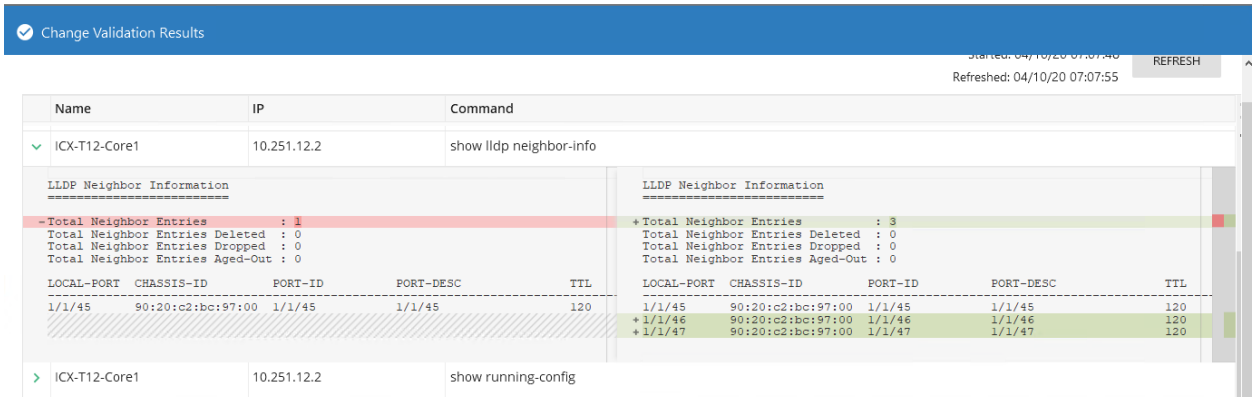

NOTE: The 'after' output is collected a few seconds after the configuration change was done. Sometimes, a protocol will need more time to show useful output. In case the LLDP neighbors would not be in the list yet, it is possible to manually refresh the 'after' output state. This is done in the next step.

6. On the right-top of the screen, note the Refreshed time. Next use the **Refresh** button to collect the data again.

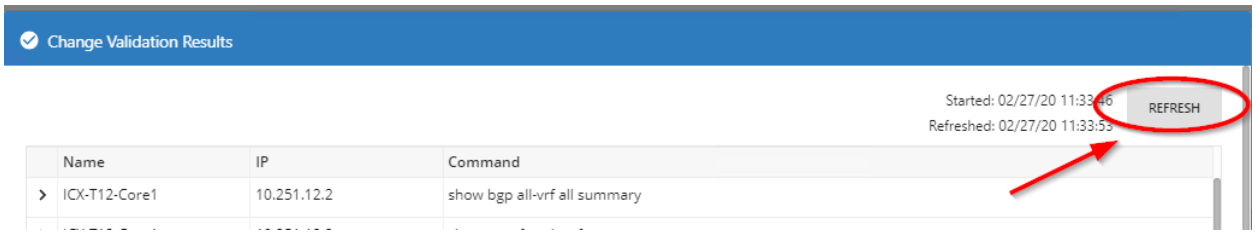

7. After the refresh has completed, notice that the Refreshed time has been updated. NetEdit has now connected to both switches and collected all this output.
8. NetEdit will collect the validation output for all the member devices of the Plan. Scroll down in the list, and notice that all these outputs are also available for the Core2 device.

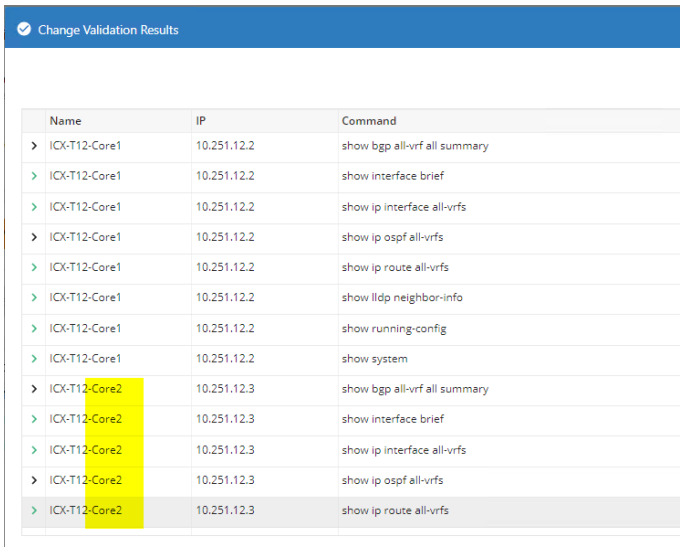

9. Click **OK** to close the validation results.

Review Plan Validation Settings

In these steps, the default validation checks will be reviewed.

10. In NetEdit, navigate to **Settings > Validation > Change Validation.**

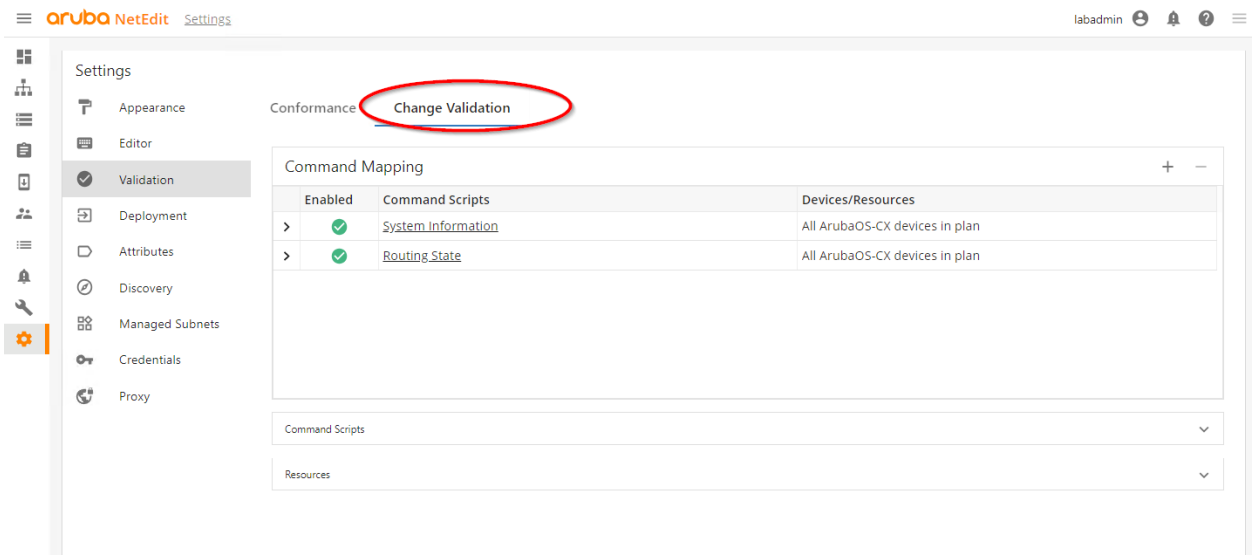

The Validation shows three sections:

- Command Mapping: Enables execution of a set of 'Command Scripts' to a set of devices. This allows the administrator to send a command set to a limited set of devices. E.g.: routing commands should only be sent to the core devices.
- Command Scripts: In this section the actual 'show' commands are defined. This section does not define on which devices the commands will be executed.
- Resources: This allows integration with devices that are not managed by NetEdit.

11. Expand **Command Scripts** and click **System Information**.

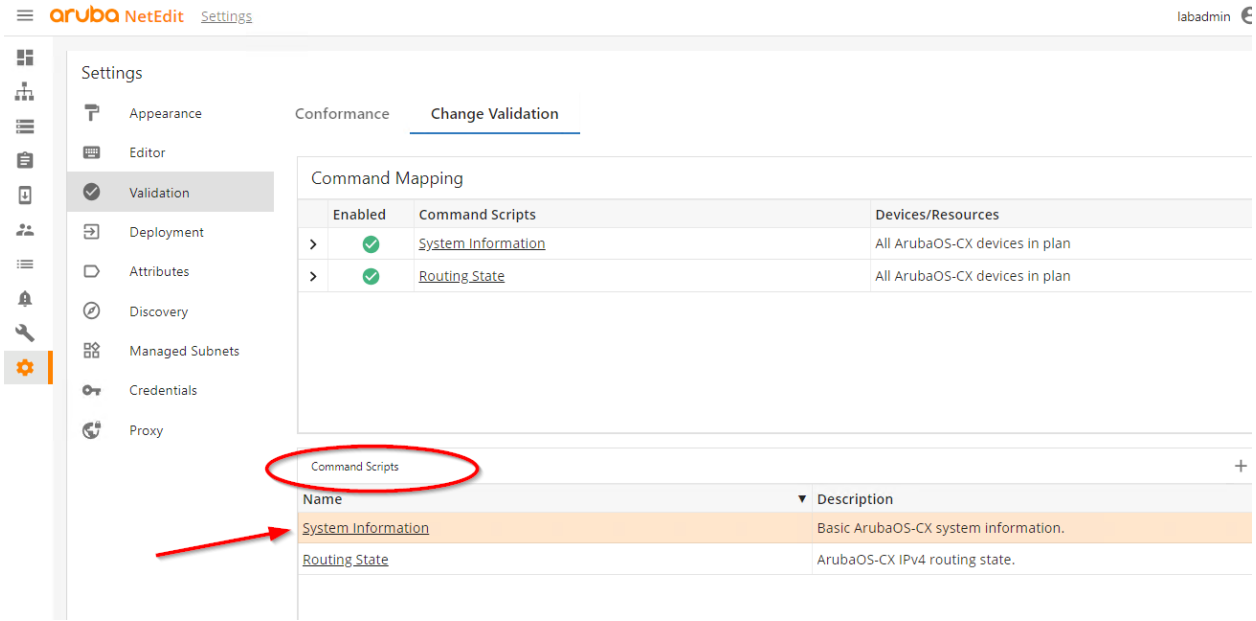

12. Review the **Commands** under System Information.

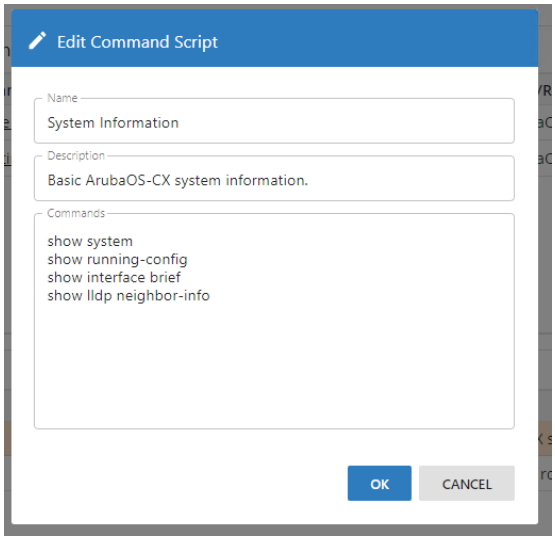

13. Click **OK** to close the screen.

Review the Command Mappings

The Command Mapping shows which command scripts will be executed by which devices.

14. Review the default mappings.

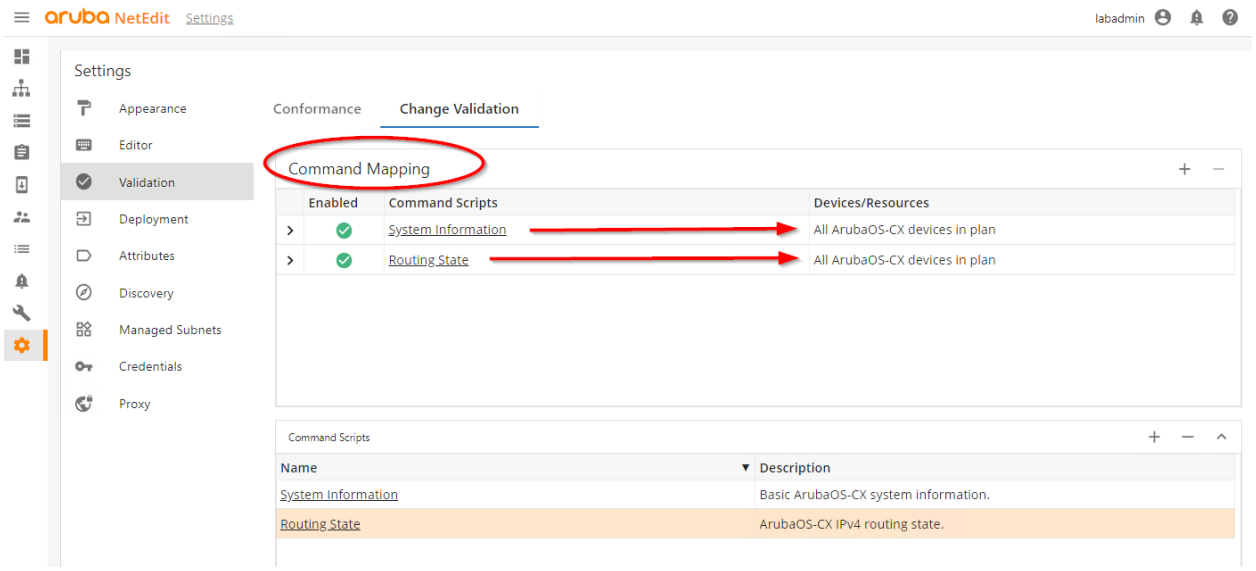

Q: On which devices will the 'System Information' commands be executed?

A: All CX devices in the plan.

Define New Validation Settings

In the next steps, you will add a command to the validation of 'System Information' validation. Then you will make a change with a plan to test the new validation.

15. Under 'Command Scripts', click **System information** and add a line for the command '**show lacp interfaces'.**

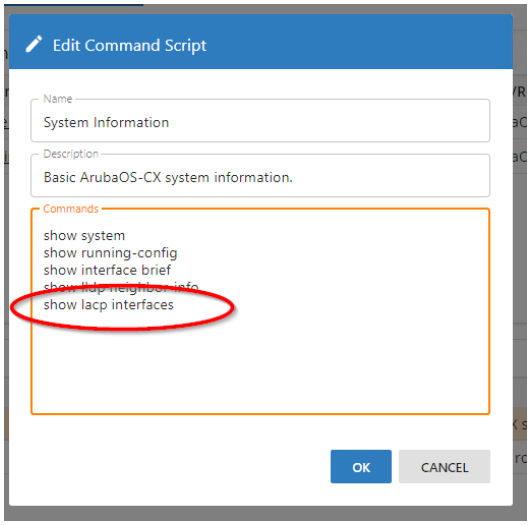

16. Click **OK** to save the change.

Test the New Validation

> **NOTE:** The new validation command will only take effect on **new** configuration plans. So, opening the Validation of an existing plan and using Refresh will not include this new command.

17. In NetEdit, navigate to devices, select Core1 (only Core1 in this example), right-click, and select **Edit Config**.
18. For the Plan name, enter **'test validation**. Click **Create** to continue.
19. Shutdown interface 1/1/46: remove **'no '** (including space) from the **'no shutdown'** line of the respective interface.

```
interface 1/1/46
    shutdown
```

NOTE: Make sure to remove the word 'no ' (including space) from the 'no shutdown' line from the NetEdit configuration, resulting in just 'shutdown'.

```
interface 1/1/46
    shutdown
    lag 256
```

20. Select **Return to plan**, select **Deploy**, and confirm with **Deploy**.
21. Check the Validation results by clicking on the **Change Validation** button, the new command **'show lacp interfaces'** should now be included. Verify that the interface 1/1/46 is shown as **'down'**.

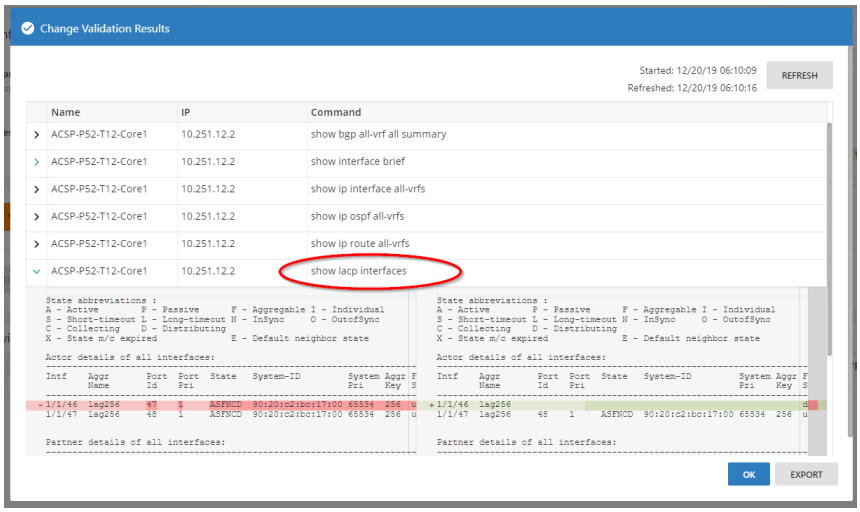

22. Click **OK** to close the Results window.
23. Revert the change with the **Rollback** option: confirm by clicking **Rollback.**

IMPORTANT: In the next steps you will make a checkpoint that will be used for the rest of the course labs, it is important to have a valid LAG state in that saved checkpoint.

Save New Configuration Checkpoint

24. Open a terminal connection to each switch (Core1/Core2/Access1/Access2).
25. Save a new configuration checkpoint on each switch.

```
ICX-Tx-Core1# copy running-config checkpoint icx-lab02-netedit
```

```
ICX-Tx-Core2# copy running-config checkpoint icx-lab02-netedit
```

```
ICX-Tx-Access1# copy running-config checkpoint icx-lab02-netedit
```

```
ICX-Tx-Access2# copy running-config checkpoint icx-lab02-netedit
```

IMPORTANT: This checkpoint is required for the next Lab activities, make sure to verify that the checkpoint is available on each of the four devices.

This completes the NetEdit configuration validation task

You have completed Lab 2!

Task 6: Organizing Devices Using Attributes (Optional)

If time permits, you may complete the next optional tasks. Check with your instructor.

Diagram

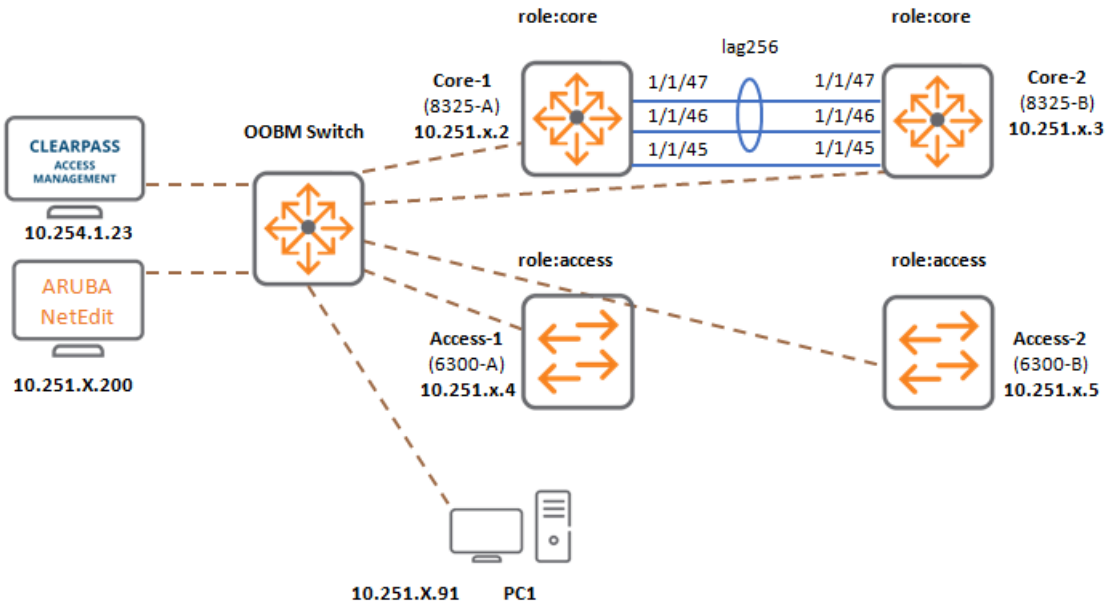

Objectives

- Define new attributes based on role, location
- Review attributes under device details
- Use attributes in queries and save a query
- Filtered validation

NetEdit allows the administrator to group devices based on features or location. The administrator can define attributes on the devices and assign values to these attributes.

These attributes can then be used in:

- Device views and plan configuration
- Configuration validation

By default, devices do not have any attributes defined. It is also possible to assign attributes to a plan. For a plan, NetEdit has predefined attributes for Change Management, such as:

- Change-ID
- Approved-By
- State

Steps

Define New Attributes Based on Role or Location

1. Open NetEdit and navigate to **Settings > Attributes**.
2. Select **Device**, and then select the '**+**' sign to add a new Attribute.

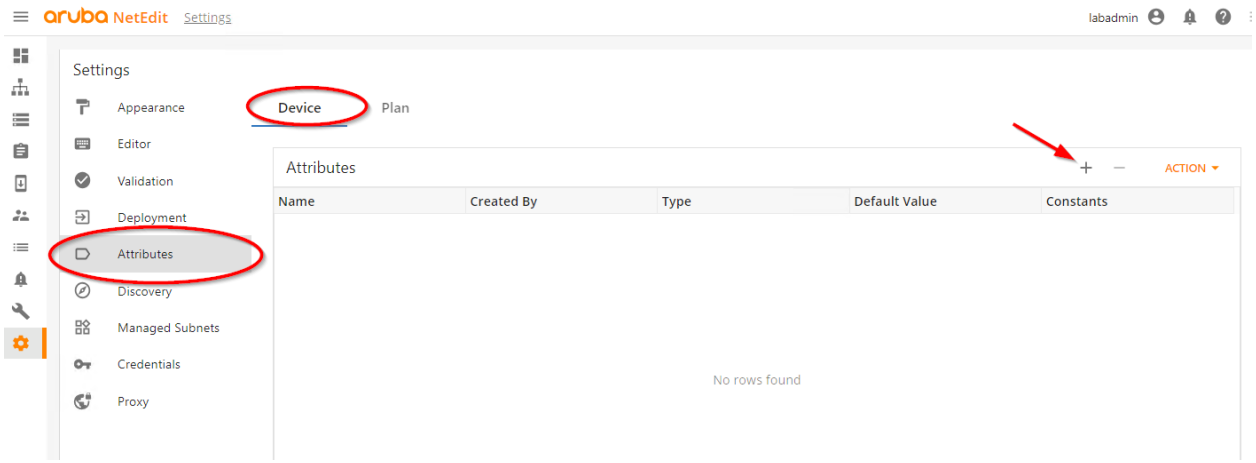

3. Enter the values listed below and click **Create.**

- Name role
- Type LIST
- List Values core,access,undefined
- Default value undefined

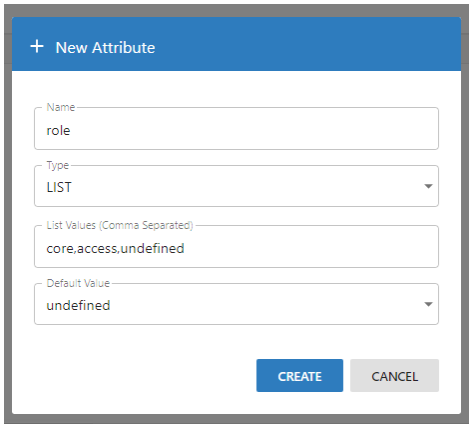

NOTE: In a similar way, a location attribute could be added to indicate the physical site information of the device.

Assign Attributes to Devices

4. Navigate to **Devices**, select the 2 core switches, and from the **Action** menu, select **Edit Attributes**.

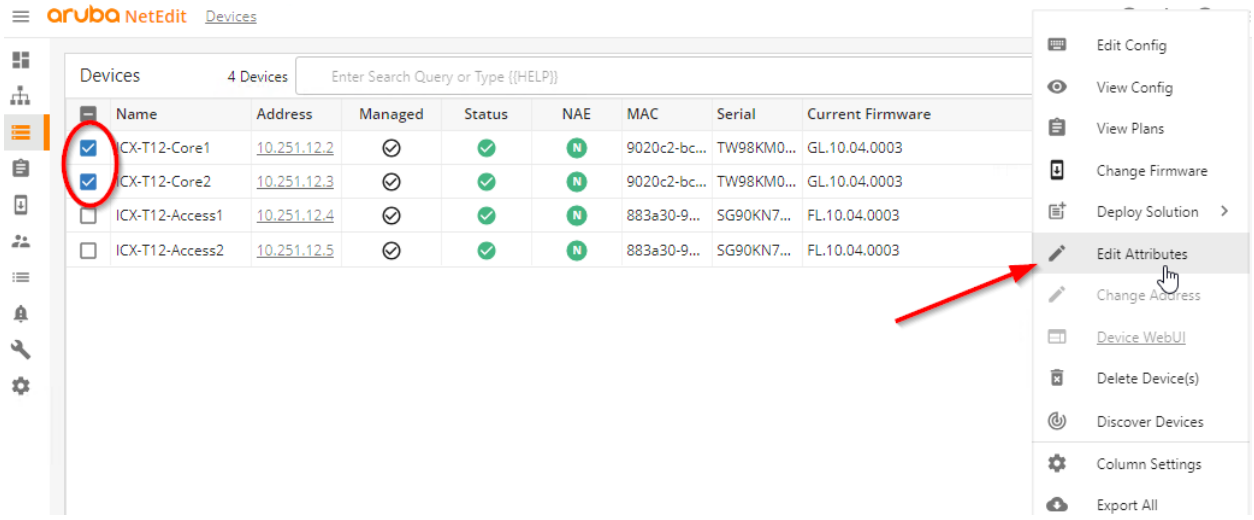

5. Select the **'role'** attribute and set it to **'core'**. Then click **Save**.

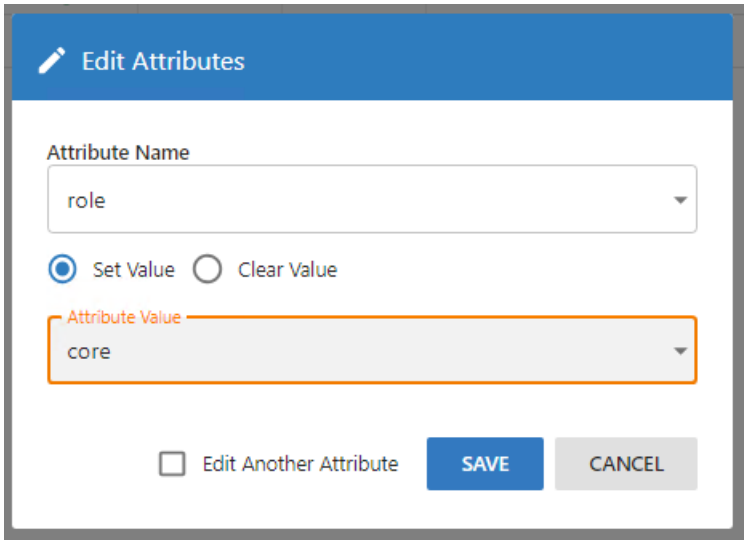

6. Once the attribute has been saved, uncheck the core switches and select both Access switches.
7. Repeat the steps to add the 'role' attribute to the two Access switches. Set the value for 'role' to 'access' and click **Save**.

Verify

8. In the device list, open the switch **Core1** by clicking its IP address (10.251.X.2). In the device details, the new attribute will be shown.

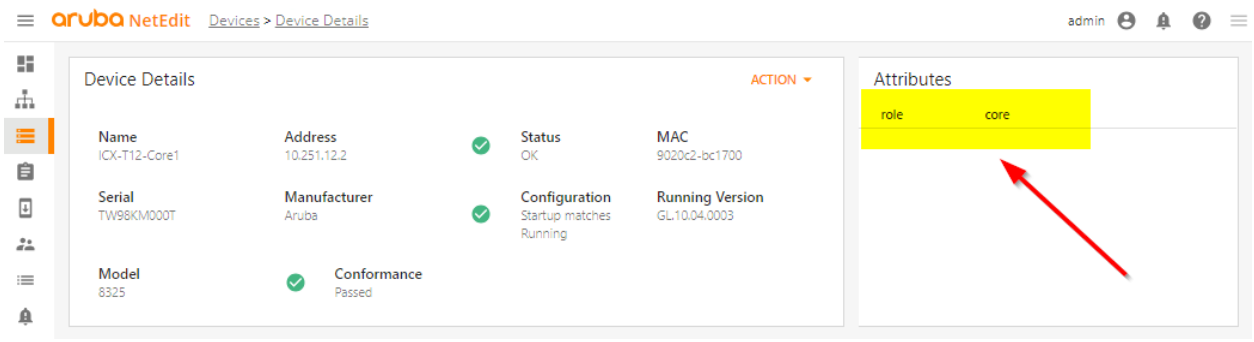

Use Attributes in Queries and Save a Query

These attributes can now be used to filter the devices in the device list.

9. Navigate to **Devices**.
10. At the top of the screen, a query filter can be applied.

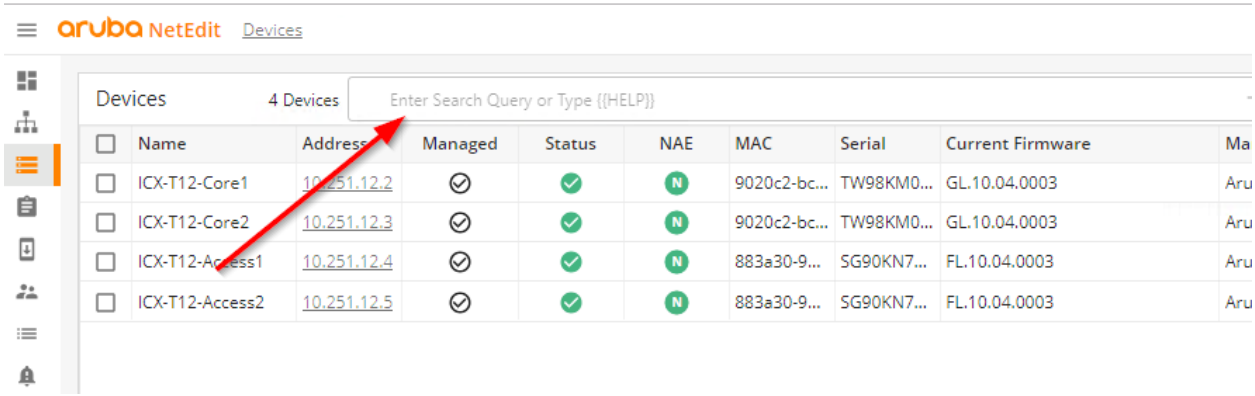

11. Enter **{{HELP}}** in the query field and press **<ENTER>.**
12. A help webpage will be shown, review the help options.
13. Close the help page and return to the NetEdit page.
14. Enter **'role:core'** in the query field and press <ENTER>.
15. The device list should only show the devices with that role assigned to them (“core”).

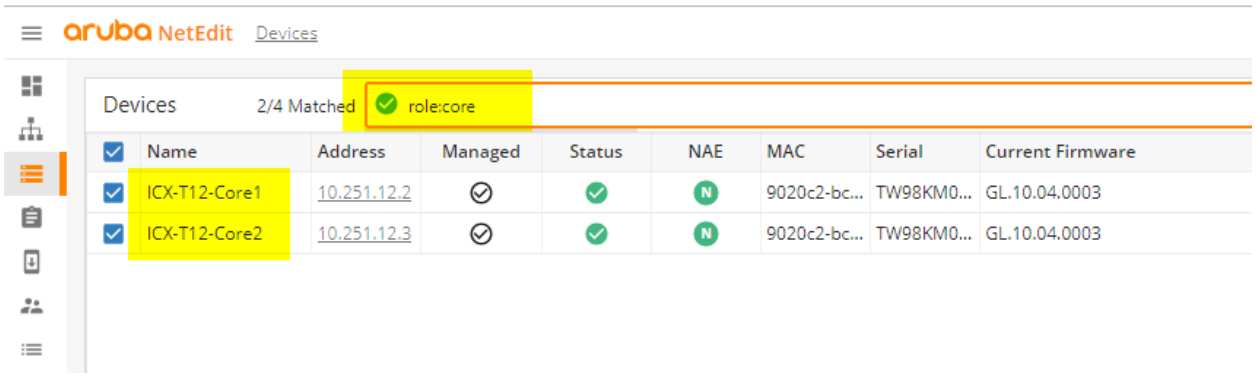

Save Query Filter

Since the administrator may be using the same filter frequently, the query can be saved for future use.

16. On the right side of the filter, click the **'+'** sign to add a saved query.

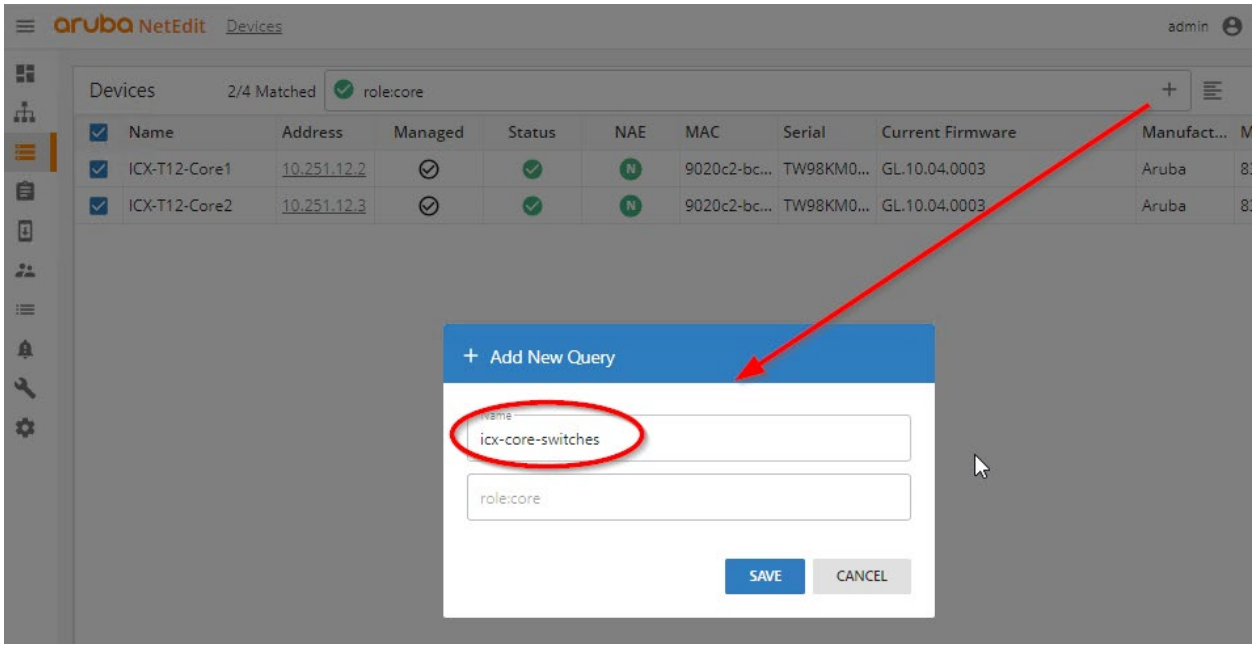

17. Enter **'icx-core-switches'** for the query name. The query will be copied from your currently active query filter, in this example **role:core**.
18. Click **Save** to save the query.
19. Change the 'Search Query' field to '**role:access**'.
    a. Verify the two access switches are shown.
    b. Save the query as **'icx-access-switches'**.

Using the Saved Filters

20. On the right-side of the filter, select the **Menu** icon to see the saved query filters.

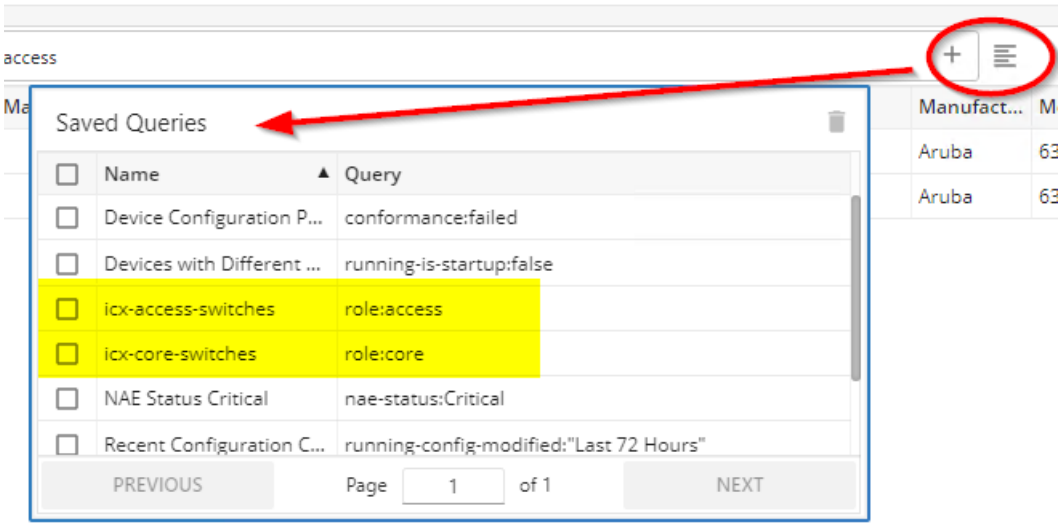

21. Double-click on the **'icx-core-switches'**. This should activate the filter and show the 2 core switches.

NOTE: The device queries can be used to restrict validation tests to a subset of devices.

Optional Task 7: Configuration Plan Conformance

Objectives

- Review default conformance settings
- Introduce a configuration error on Core2
- Define example conformance settings
- Verify operation of conformance settings

NetEdit allows the administrator to set verification rules to ensure the configuration conforms with the set policy.

This can be used for:

- Security conformance: ensure the configuration is compliant with security policies
- Configuration conformance: ensure the configuration contains key base configuration elements.

Steps

Review Default Conformance Settings

1. Login to NetEdit, navigate to **Settings > Validation > Conformance.**

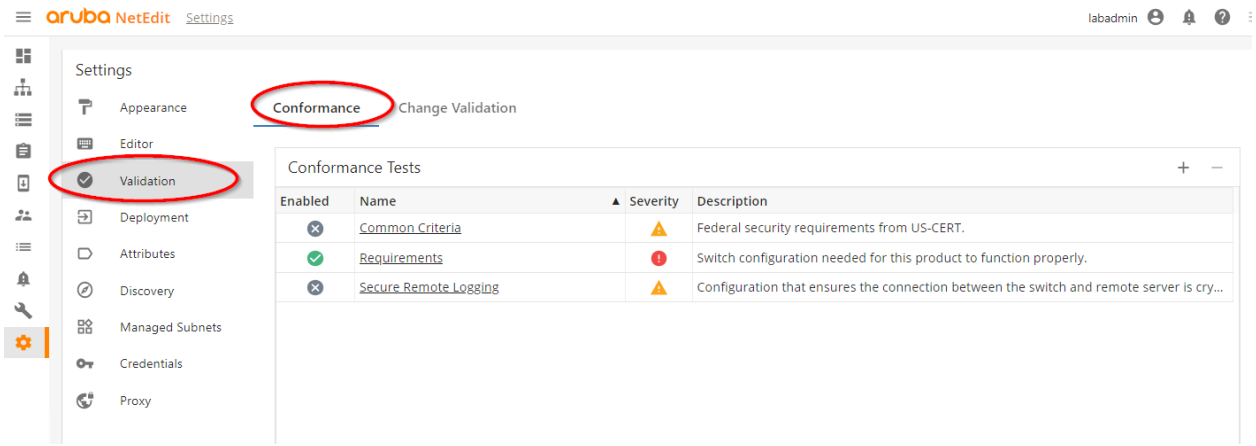

2. By default, NetEdit has three conformance rule sets.

- **Common Criteria**: security requirements set by US-CERT
- **Requirements**: basic configuration to ensure the REST API is enabled.
- **Secure Remote Logging**: verifies remote logging in the configuration

Only the 'Requirements' rule set is enabled by default.

3. Click the **Requirements** entry and review the checks performed by this rule set.

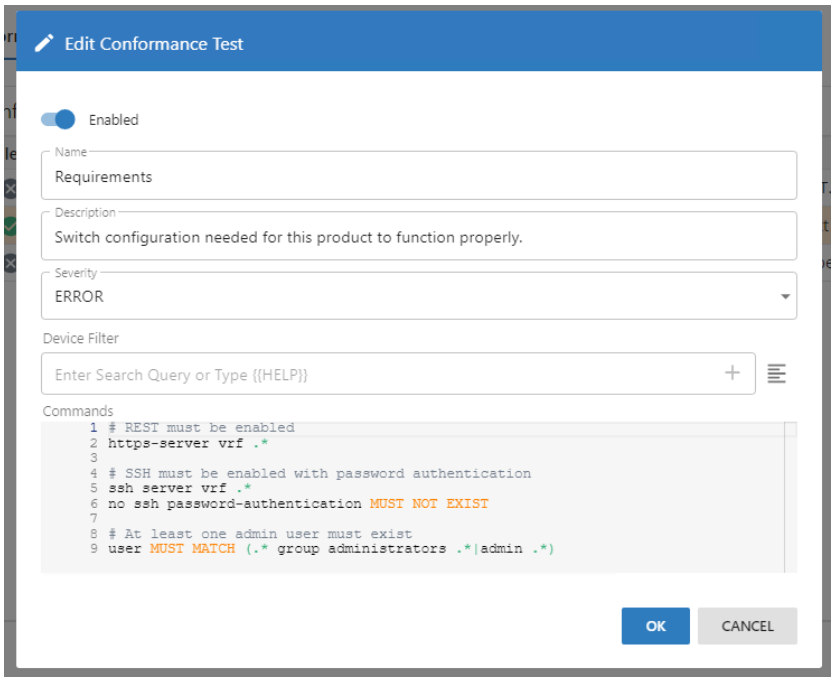

4. Click **OK** to close the 'requirements' conformance test.

Introduce a Configuration Error on Core2

NetEdit will connect every 5 minutes to each managed switch and collect the configuration. So even when an administrator makes a direct configuration change using SSH or console on the device, NetEdit will be aware of this change or changes within 5 minutes. This also means that any configuration/policy error made directly on the device can be quickly discovered by NetEdit.

In the next steps, you will make a 'local' configuration error on Core2, where this 'error' should be discovered by NetEdit at the end of this task.

The lab uses the SNMP system location as a test configuration.

5. Open a terminal connection to Core2. Enter configuration mode and apply a 'bad' SNMP system location.

```
ICX-Tx-Core2# configure terminal
ICX-Tx-Core2(config)# snmp-server system-location bad-location
```

6. Verify the bad configuration is in the running configuration.

```
ICX-Tx-Core2(config)# show run | include snmp
snmp-server system-location bad-location
ICX-Tx-Core2(config)# end
ICX-Tx-Core2#
```

NOTE: NetEdit will poll the device configuration every 5 minutes. This means it may take up to 5 minutes before this change on the switch is detected by NetEdit and

visible in the NetEdit UI.
Therefore, you should **not** attempt to check this running configuration in NetEdit yet; just continue with the next steps first.

Define New Conformance Settings for SNMP System-location

In this section, a new conformance test will be created to ensure an SNMP system location is configured and it must meet the correct naming convention.

The lab is using the SNMP system-location since this configuration option is not impacting any other lab activities, but the same conformance logic could be applied to any command.

If any administrator makes a mistake in a configuration plan with regards to the SNMP system-location command, the conformance test will report an error about the change.

7. In the conformance screen, add a new conformance test by clicking the '**+**' icon.

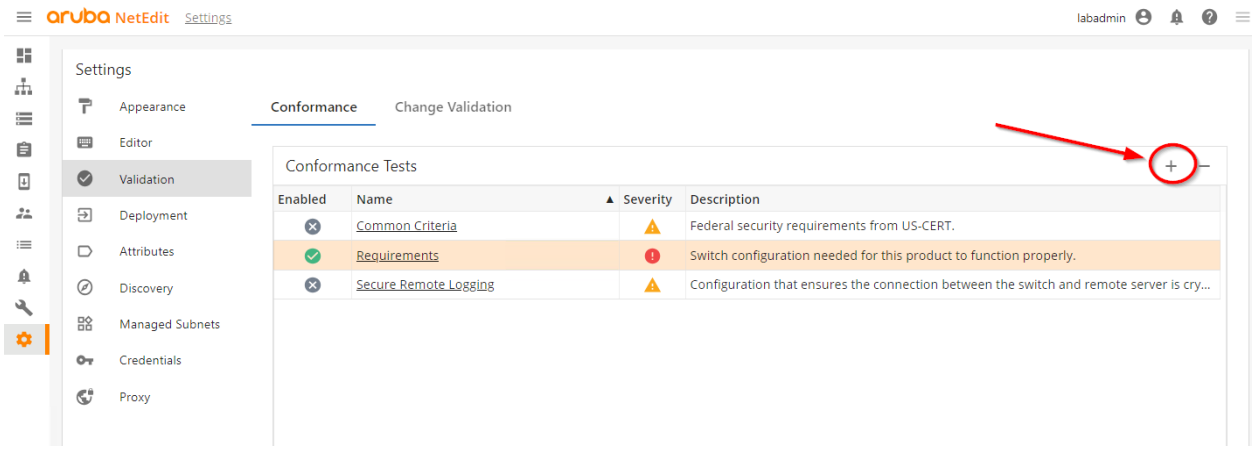

8. For a name, enter **'icx-snmp'**.
9. Severity level can be set to:

- **'ERROR'**: displayed as **red** in the configuration
- **'WARNING'**: displayed as **orange** in the configuration

Leave the severity to **'ERROR'**.

10. Device Filter: This allows the administrator to apply certain conformance checks on specific device. In this example, only the Core switches will be checked.

Device Filter: **role:core**

11. In the **'commands'** field, enter the following command. Note that regular expression syntax can be used. The '.' character indicates 'any' character and the '*' character means the previous character can be repeated 0 or more times.

- snmp-server system-location icx-.*

12. The resulting conformance test should look like this:

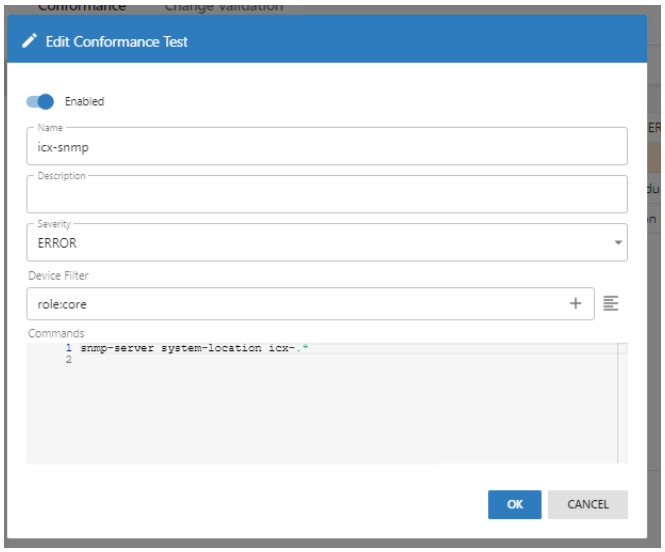

13. Click **Add** to save the new test.

NOTE: To learn more about the syntax for conformance checks, consult the online help and navigate to 'Conformance'.

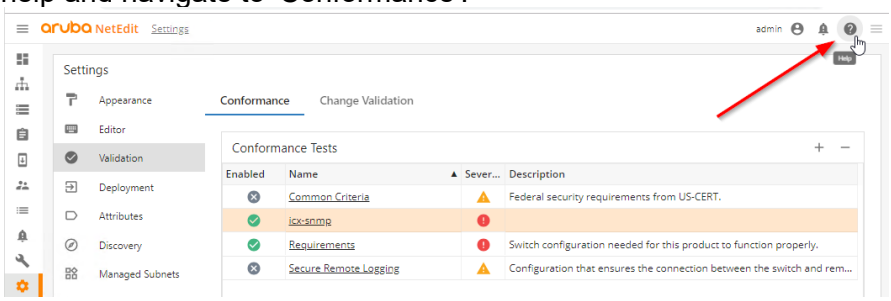

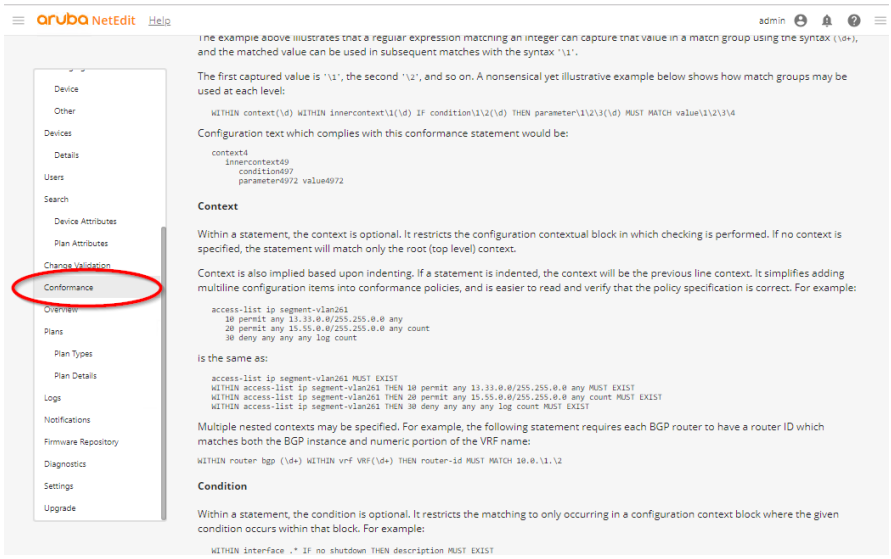

Verify Operation of Conformance Settings for OOBM

In the next steps, the new conformance set will be tested by making some mistakes in a new configuration plan.

14. In NetEdit, navigate to **Devices**. Select the **'Core1'** switch. Navigate to **Action > Edit Config**. Set some text for the plan, for example '**conformance test**', then click **CREATE**.
15. Add a new line for the SNMP system location: set it to '**bad-location**'

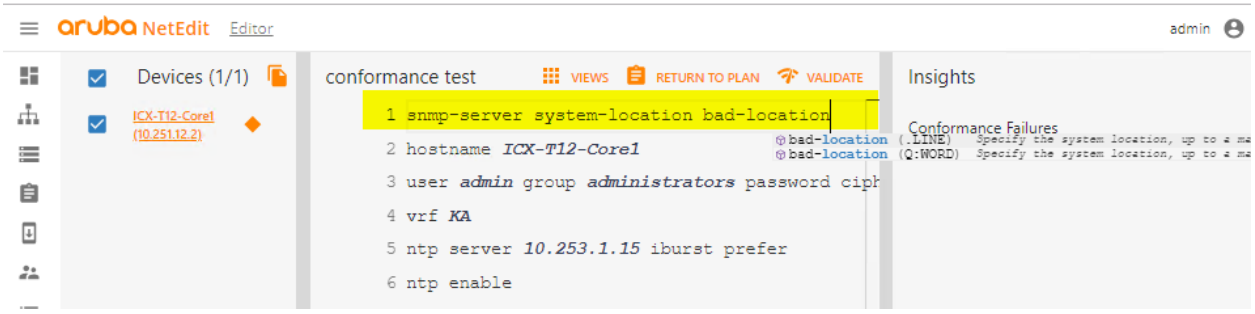

16. Press <ENTER>, NetEdit will move the line to the appropriate place in the configuration, and will also verify the conformance check at this point. Since the line does not conform, it is shown based on the error level, in this case 'ERROR', so in red.

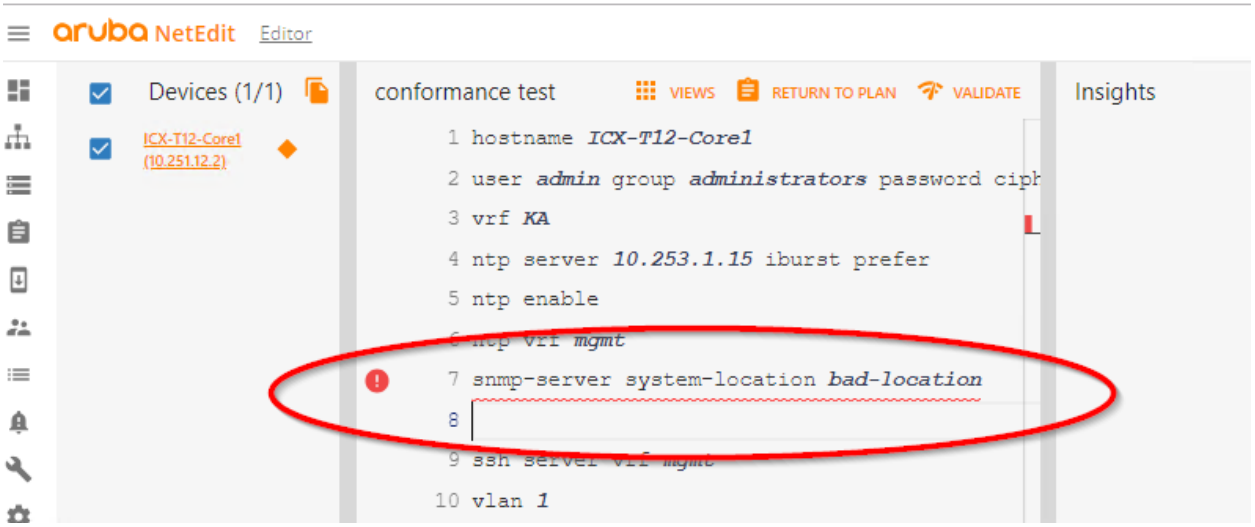

17. Move the mouse over the red icon before the line to see the details.

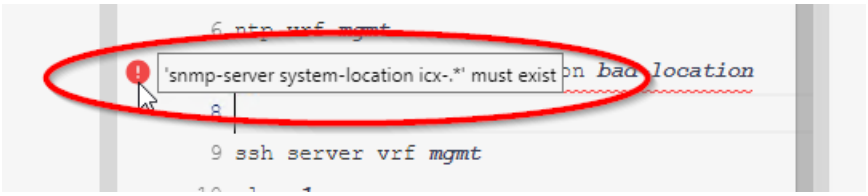

18. Click the line containing the error, the 'Insights' on the right-side will provide more details about which conformance test and lines apply to this line (note that you must expand the two sections, Conformance Failures and icx-snmp, to see the full results).

19. Click **Return to Plan**. Notice that the Plan Details shows 1 conformance error as well.

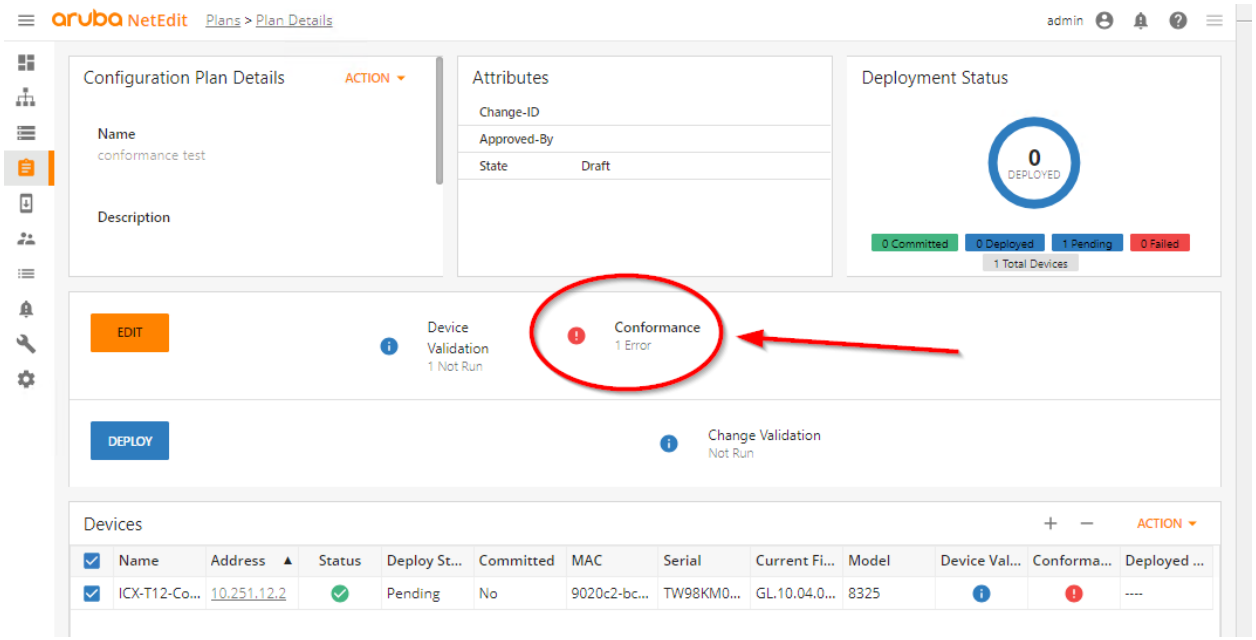

20. Click **Conformance** to review the error. Expand the **icx-snmp** entry.

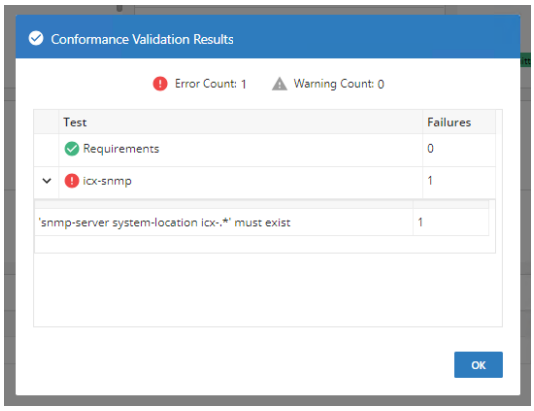

21. Click **OK** to close the Conformance window.

Deploy Configuration

Now the error will be pushed to the devices and you will explore how NetEdit can display the devices that are not conforming the rules.

22. Next click **Deploy** to push the configuration to the device, confirm with **Deploy.** Wait a few moments until the deployment completes.

Now the device has a configuration with a conformance error. When an administrator wants to review if there are any devices with conformance failures, a built-in filter can be used in the device view.

23. Navigate to **Devices**, open the saved filter list and **double-click** the **'conformance:failed'** filter.

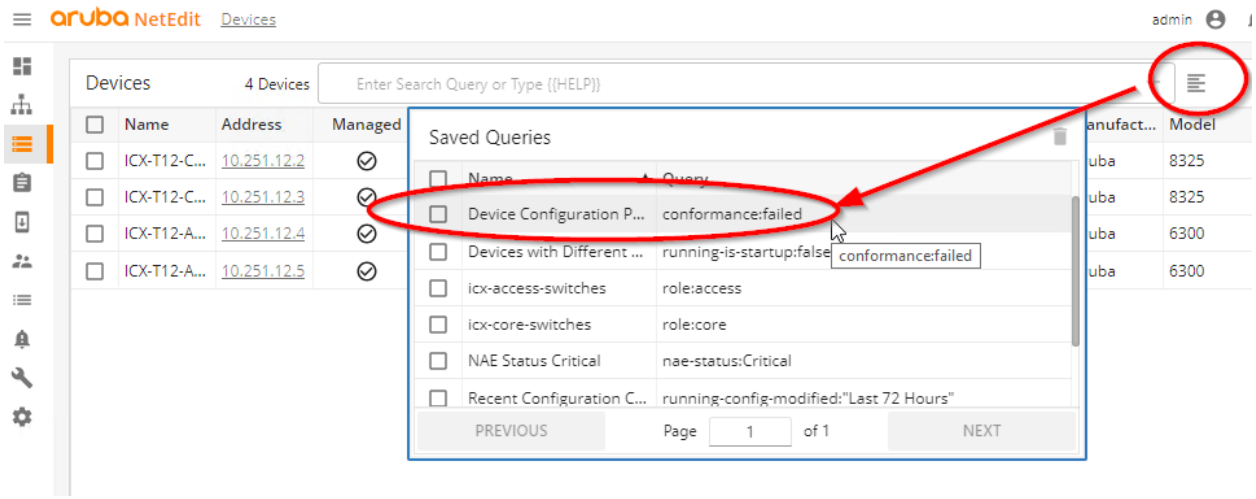

24. The filter should show:

- Core1 : configuration error via NetEdit
- Core2 : configuration error made directly on the device CLI

NOTE: If Core2 is not listed, 5 minutes may not have passed since the configuration change was made on the Core2 CLI. Wait a few minutes and retry the filter. You may also just continue without checking Core2.

25. Example output:

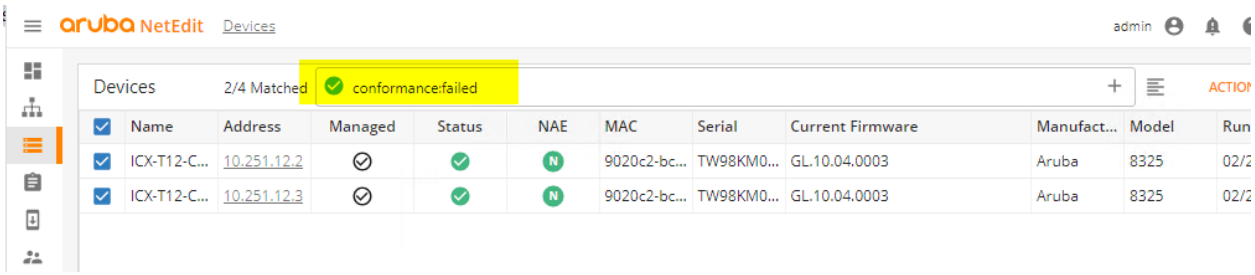

This shows that NetEdit conformance tests can also be used to verify an existing configuration or configuration changes that were made directly on the device CLI.

Correct the Configuration

Now the configuration will be corrected so that the location name conforms to the checked rules.

26. Select both the **Core1** and **Core2** devices. From the **Action** menu select **Edit Config**. Name the plan **'snmp location'** and click **Create** to define the plan.
27. Correct the SNMP location, set it to **'icx-lab'**, press <ENTER> to start the validation. The line should no longer show as 'red', since it conforms to the set rules.

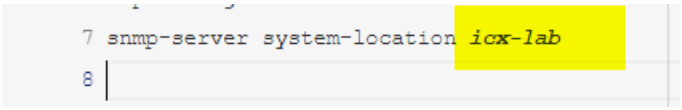

28. Click **Return to Plan**, select **Deploy**, and confirm using **Deploy**.
29. Wait a few moments for the plan to be deployed to both Core switches.
30. Once the plan has been deployed, navigate to **Devices** and apply the query filter **'conformance:failed'**. The list should be empty, since all devices are conforming the rules.

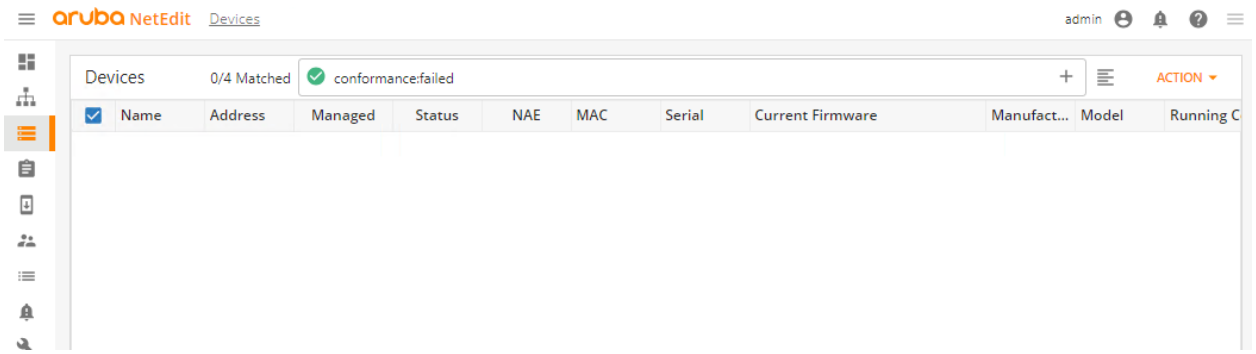

This demonstrates the conformance test feature.

The conformance test is not used in future labs, so the conformance test will be disabled.

31. Navigate to **Settings > Validation > Conformance** and select the **'icx-snmp'** conformance test.
32. Use the slider to disable the test and click **OK** to save the change.

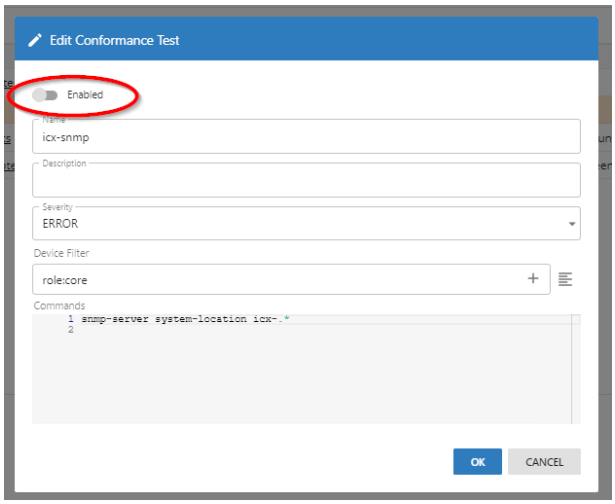

This completes the optional conformance check task.

You have completed Lab 2

Lab 03: Network Analytics Engine- NAE Configuration

Lab Diagram

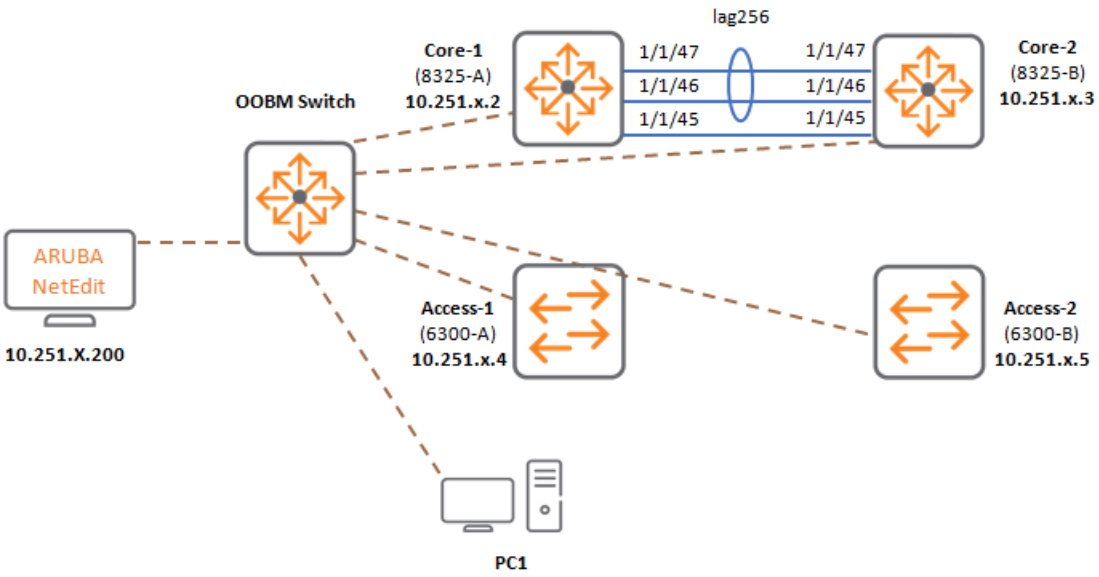

Overview

In this lab activity, the AOS-CX Network Analytics Engine feature will be explored.

First there is a review of the built-in NAE agent. Next a new script will be added and an agent will be configured.

There is also an optional task to add another NAE script that integrates with a ping monitor module of the switch.

The last task is to review how the NAE script and agent settings are stored in the switch's configuration.

Objectives

- Understand the difference between NAE scripts and agents
- Learn how to configure NAE agent parameters

- Understand NAE agent alerts and graphs
- Understand how NAE scripts and agents are stored in the switch configuration

Task 1: Prepare the setup

Objectives

- Verify devices are ready for the NAE lab activity
- Must have completed lab 02 - NetEdit configuration

Steps

1. Open a terminal connection to Core1.
2. Verify that time is currently synchronized using NTP.

```
ICX-Tx-Core1# show ntp status
NTP Status Information

NTP                          : Enabled
NTP Authentication           : Disabled
NTP Server Connections       : Using the mgmt VRF

System time                  : Thu Feb 13 09:18:46 EST 2020
NTP uptime                   : 2 days, 23 hours, 28 minutes, 30 seconds

NTP Synchronization Information

NTP Server                   : 10.253.1.15 at stratum 2
Poll interval                : 1024 seconds
Time accuracy                : Within -0.000247 seconds
Reference time               : Thu Feb 13 2020 8:59:02.459 as per US/Eastern
```

3. Examine the actual time.

```
ICX-Tx-Core1# show clock
Thu Feb 13 09:16:08 EST 2020
System is configured for timezone : US/Eastern
```

4. On PC1 (Management PC - OOBM), verify the time in a command prompt (cmd.exe). Enter the command **time** and just press <ENTER> when prompted to enter a new time.

```
C:\Users\student>time
The current time is:  9:16:13.51
Enter the new time:
```

The time difference should be within seconds. The NAE web UI is using the client browser time to get the time for the 'live' graphs, so therefore it is important to have correct time on the PC client and the switch.

IMPORTANT:
Notify your instructor if the time is still different!

Task 2: Review the built-in NAE script and agent

Diagram

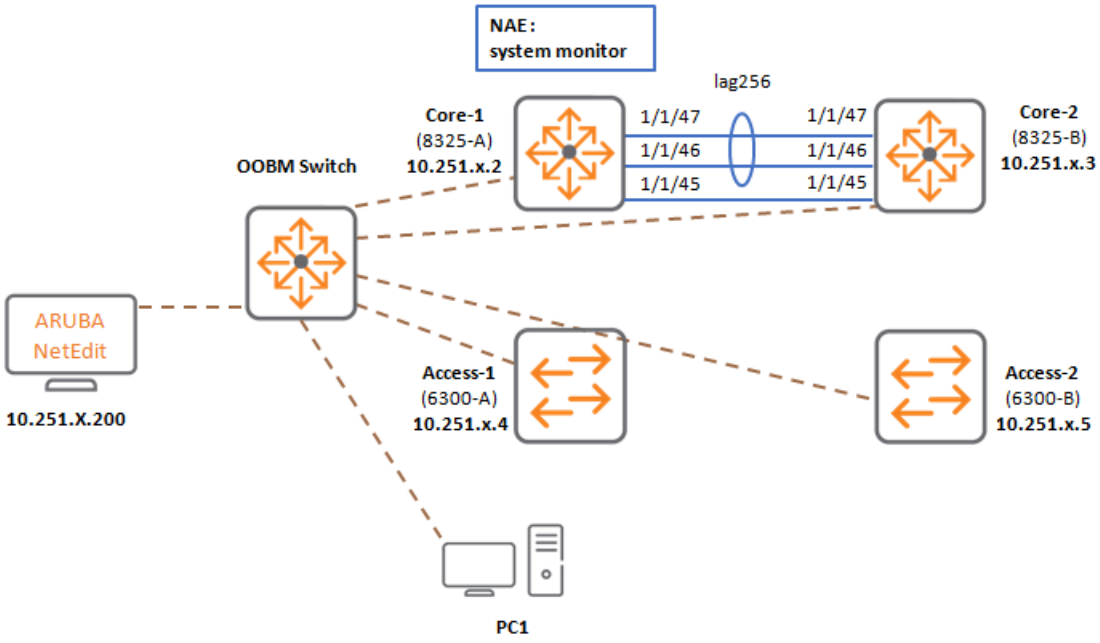

Objectives

- Get experience implementing an NAE script and agent
- Review the built-in System Monitor NAE agent
- Understand the meaning of parameters for an agent
- Review the alerts and alert actions generated by an agent

Steps

1. Open a session to PC1 (Management PC - OOBM).
2. On PC1, open a web browser to the Core1 OOBM IP: **https://10.251.x.2**, where x is your table number.
3. Log into the web interface with **admin/aruba123** credentials.

The Overview page will be displayed.

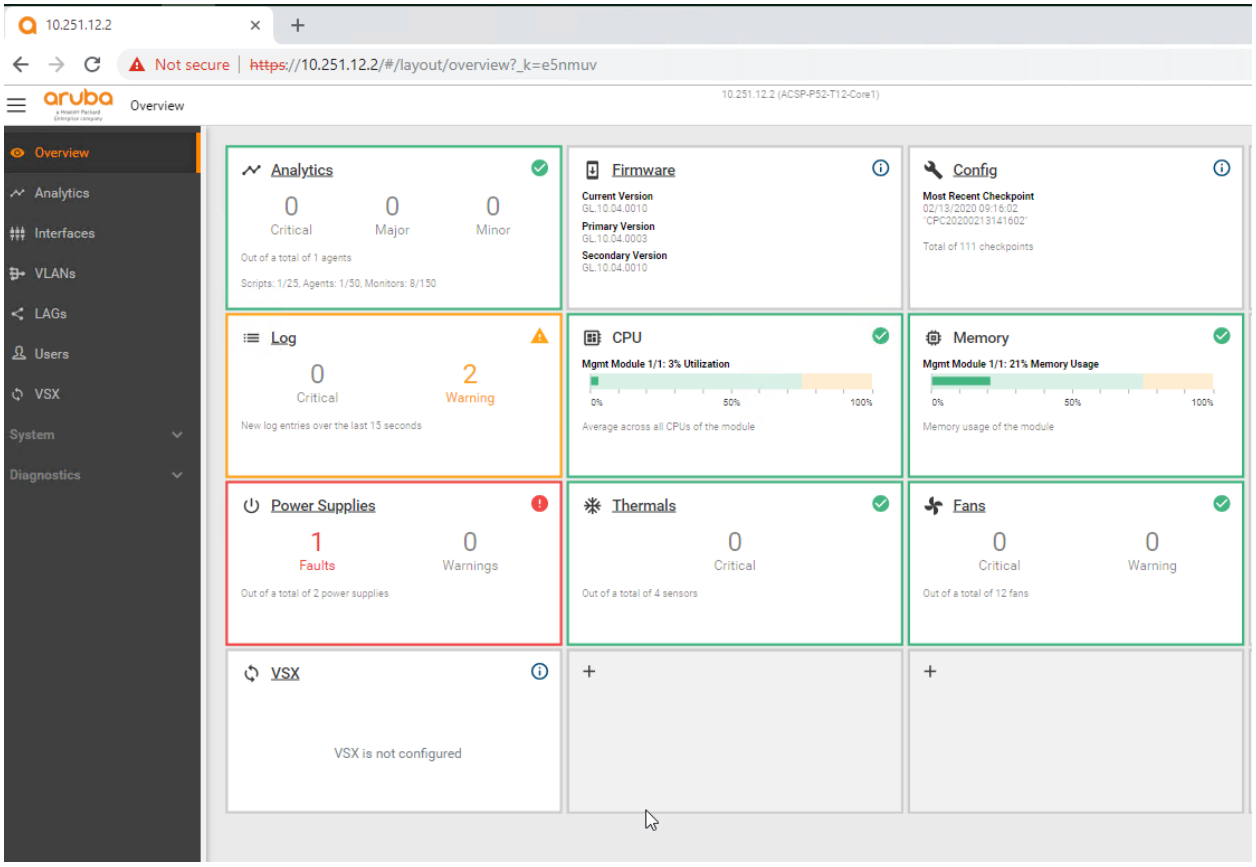

4. Navigate to the Analytics page by selecting **Analytics** in the left navigation pane. This will display the 'Agents' and the 'Scripts'.

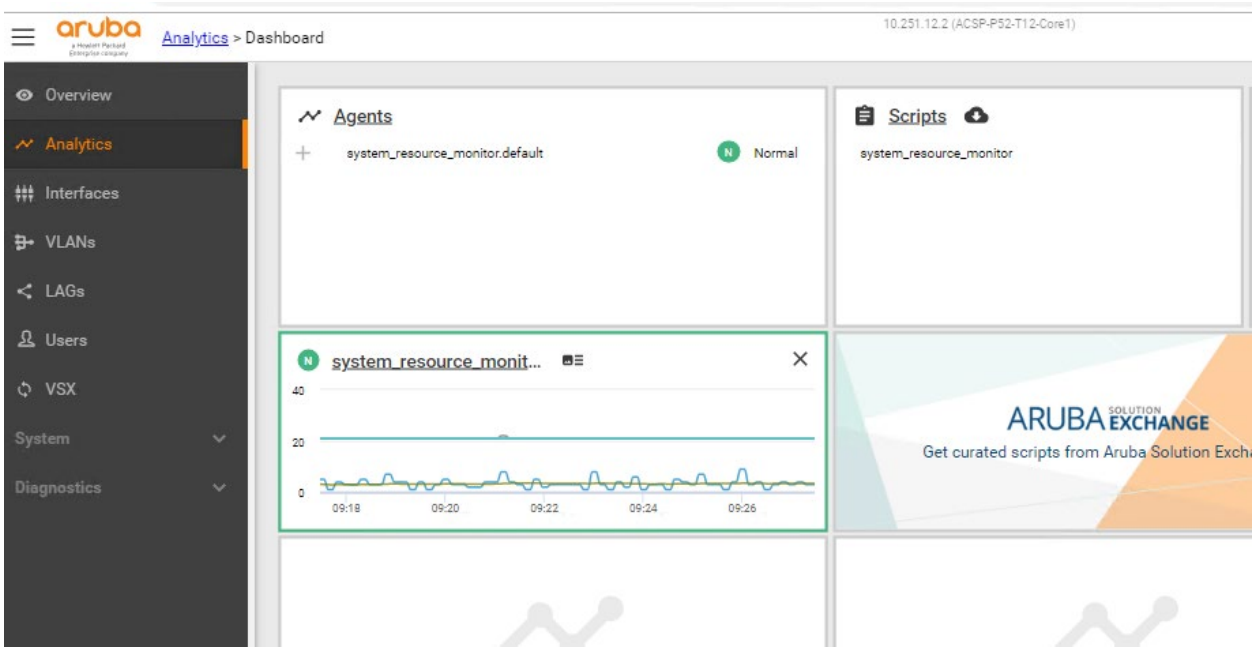

Scripts are the actual code that can be executed by NAE, while an Agent is an instance of the script. A script that is installed on the switch does not do anything by itself. An agent must be created based on the script to start the actual monitoring.

One script may be instantiated multiple times, that would depend on the script:

- Some scripts may monitor all interfaces of the switch, so only 1 instance (agent) would make sense
- Another script may only monitor 1 interface, where the interface that needs to be monitored would need to be entered by the administrator. This interface name would be a parameter on an agent. When the administrator wants to monitor 2 interfaces, 2 agents would be created based on this script. 1 agent with interface 1, the other agent with interface 2.

5. Click the **Scripts** link.

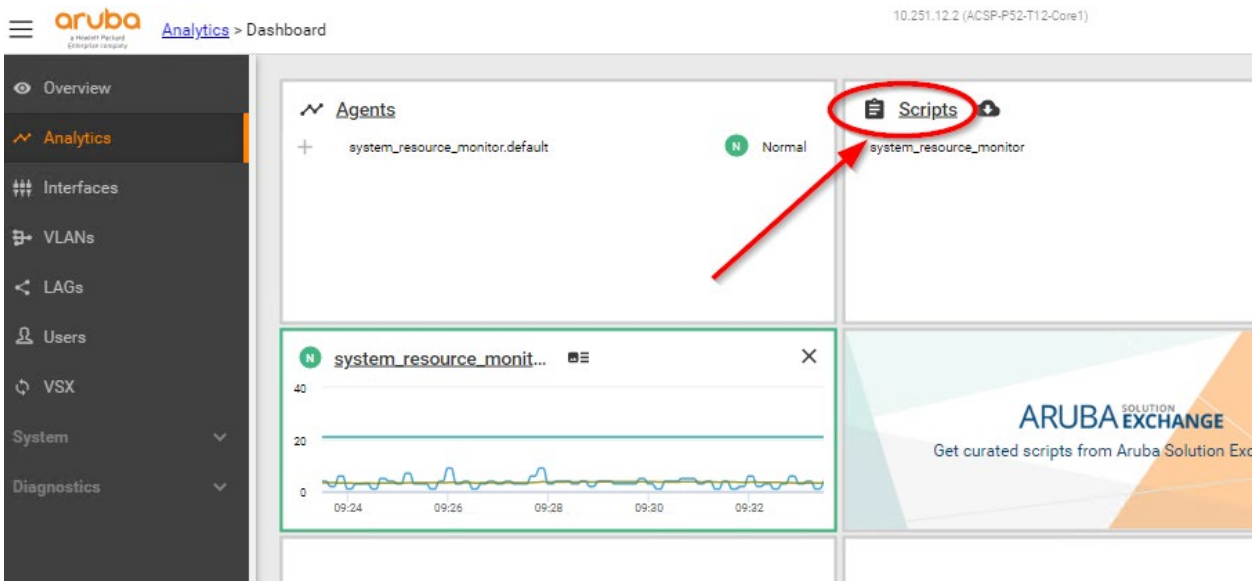

6. There is one built-in script, that is called the 'system_resource_monitor'. This is system-created and therefore cannot be deleted.

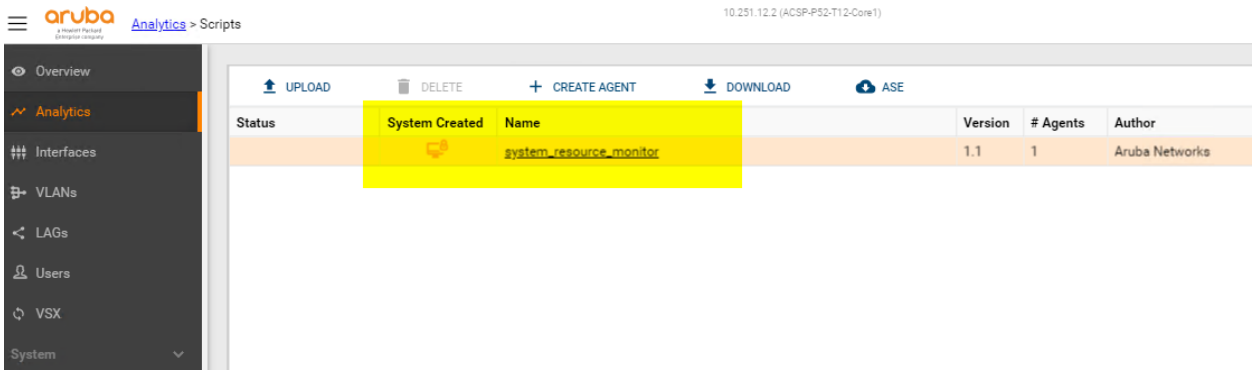

Example Script Details

7. Click the name of the script (**system_resource_monitor**) to see the details of the script. A script on the switch can also be downloaded to an administrator PC from this screen. The download is optional; it is not required in this lab activity.

> **NOTE:** The window layout may appear differently, based on your local computer's screen resolution.

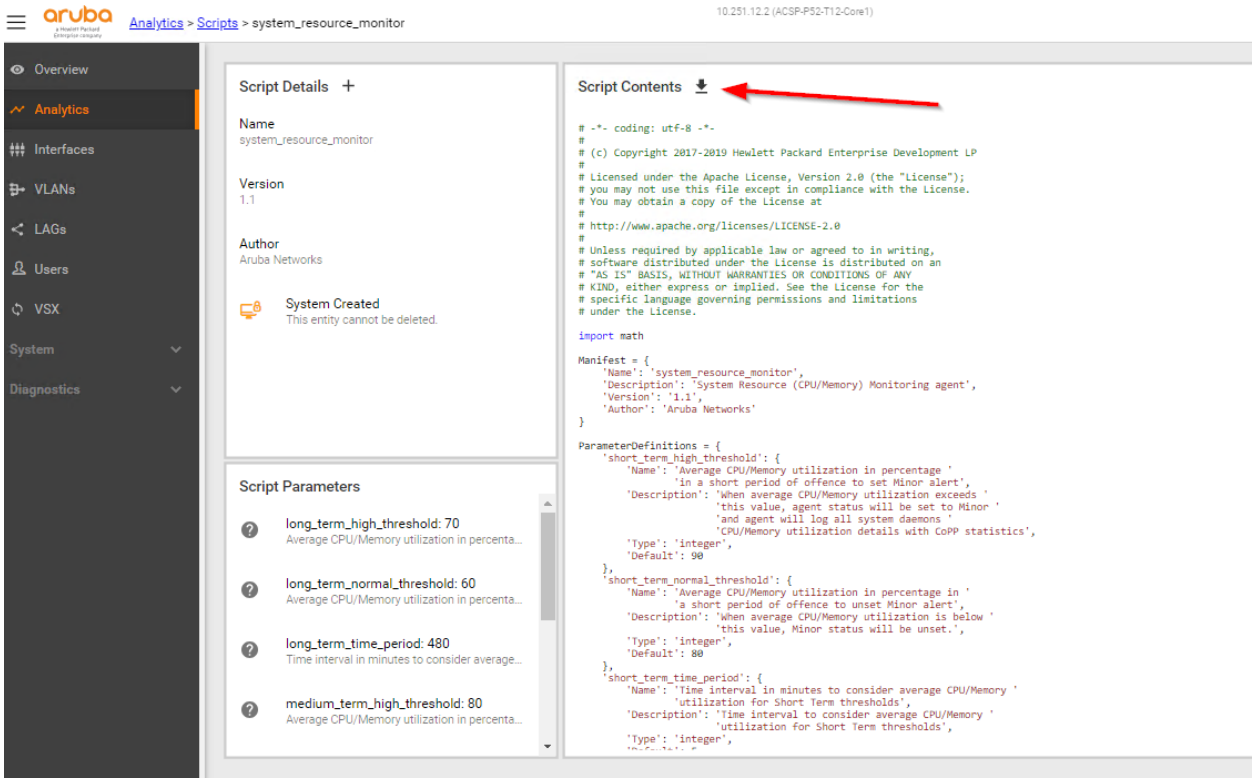

NOTE: The development of the scripts themselves is beyond the scope of this training. Please consult the product documentation if you want to learn more about the scripts and the syntax.

8. Navigate back to **Analytics** to see the Dashboard screen.

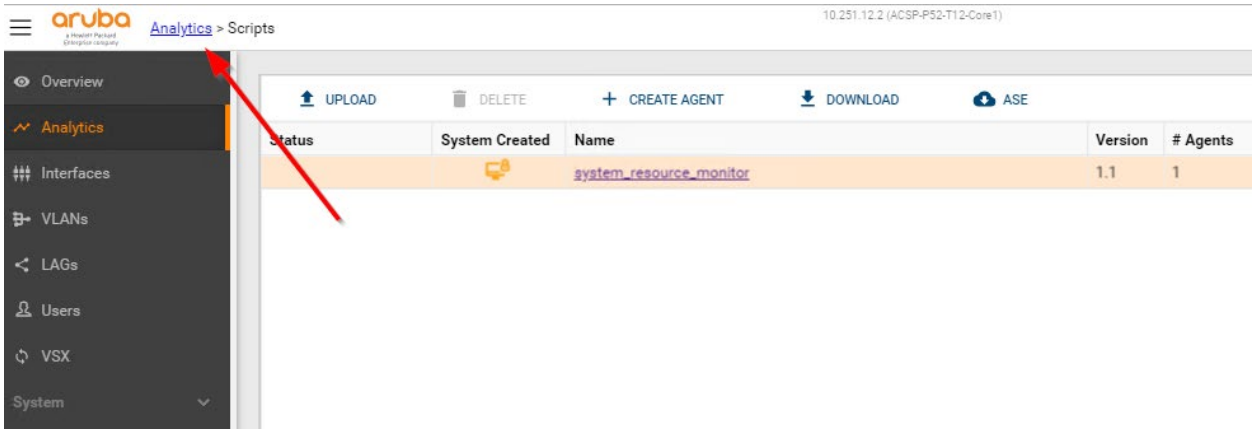

9. Click the **Agents** link to open the Agents window.

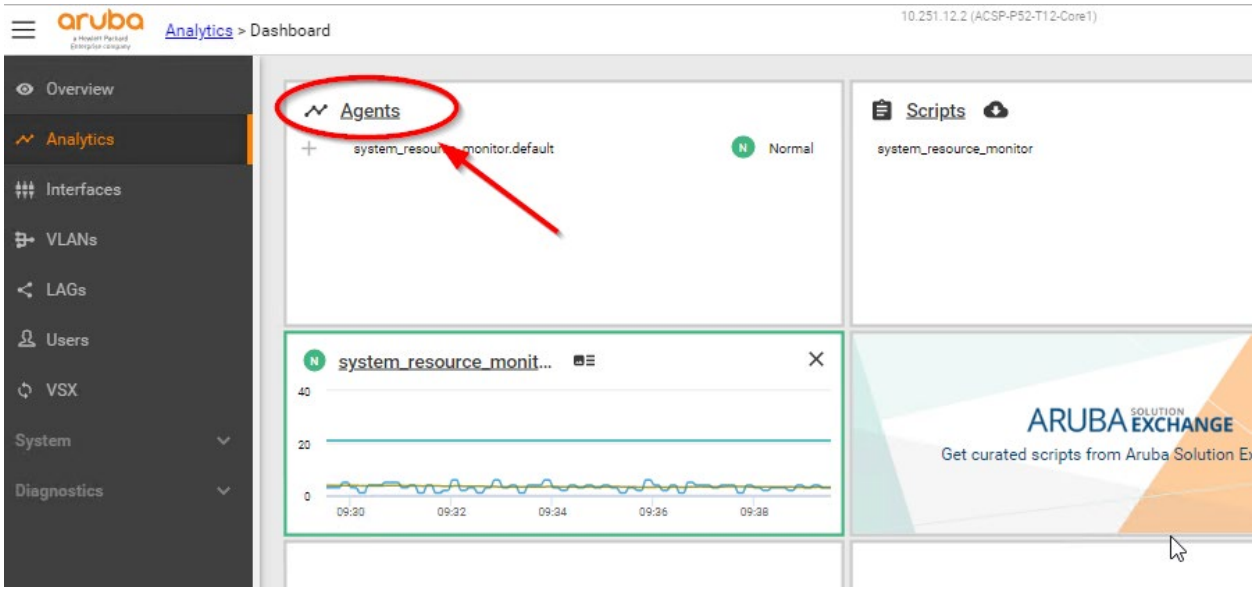

10. This will list the default agent that is running on the switch. The agents are the actual running instances of a script. In this case, there is a default agent that will be monitoring the system resources. This happens out-of-the-box.

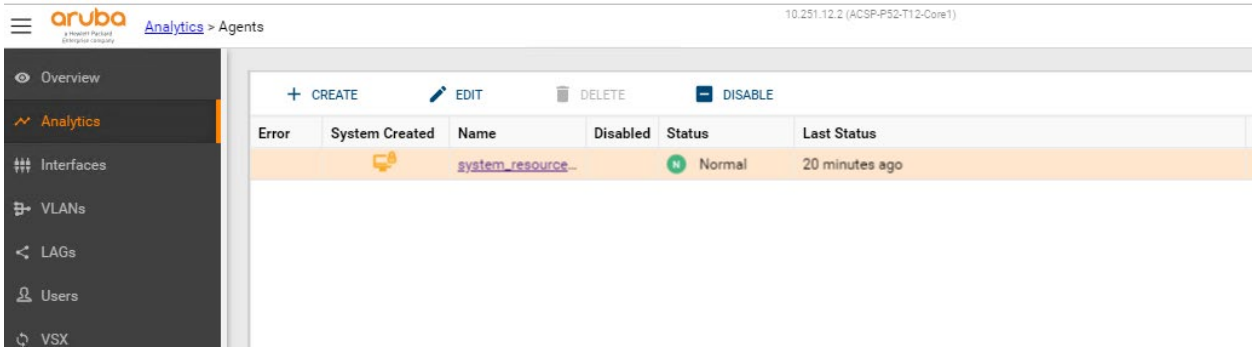

11. Click on the name of the Agent (**system_resource_monitor**) to see its details.
12. The default screen will show:
    - the agent details (script that is used for this agent),
    - the status
    - the parameters that have been set for this agent
    - the graph, if applicable
    - the alerts, if applicable

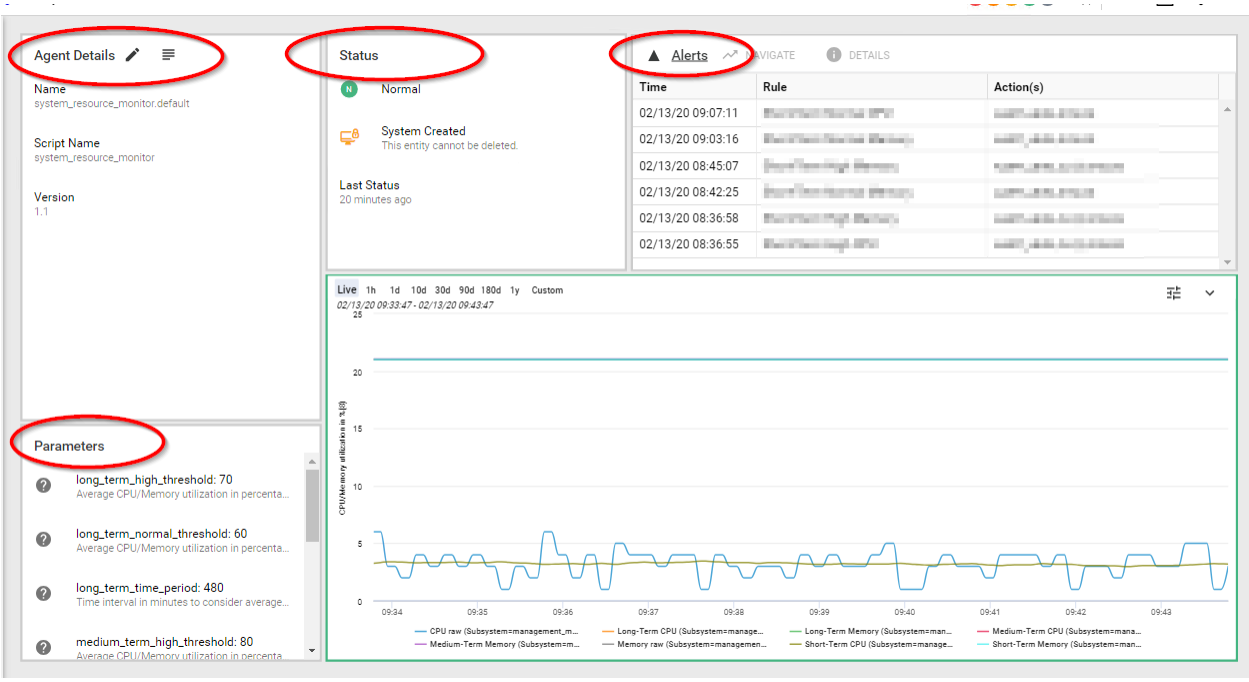

Graph Management

13. The switch has a local database that can store historical data for the agent. In the graph, the administrator can check the history of the agent. Click the '1h', '1d' and 'live' links to explore the differences.

> **NOTE:** Your actual graph will be different, the images in the lab guide were made on a system that has been running for some time. Some students report the switches begin updating after 10-12 hours.

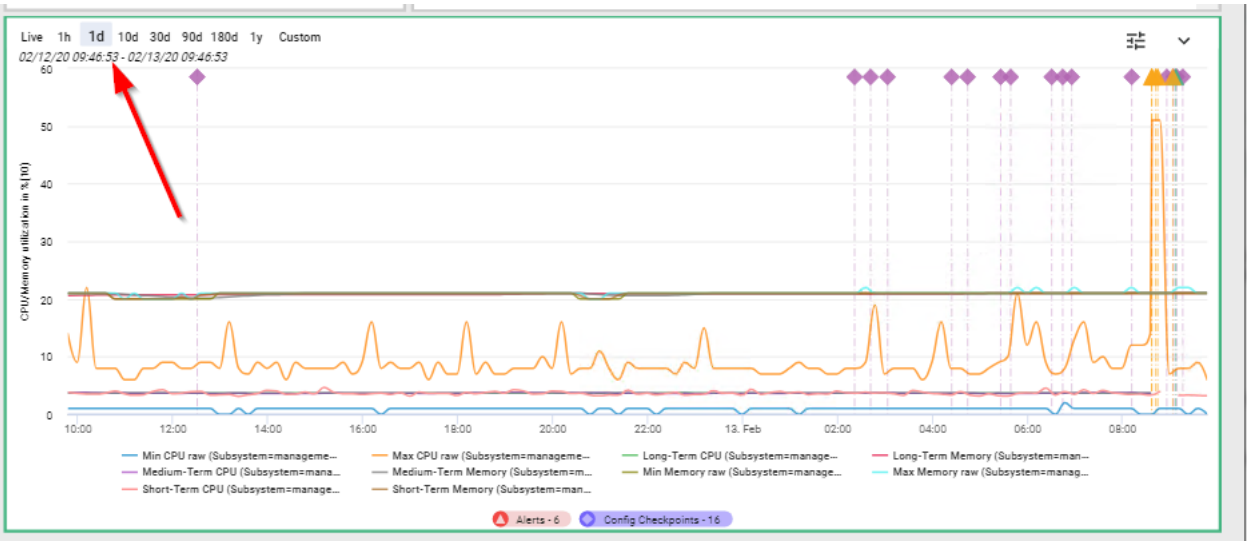

14. The admin can also apply a custom zoom by selecting a time range with the mouse (click and drag).

15. The administrator can choose to disable certain lines on the graph by clicking them at the bottom.

Alerts

Agents can monitor various options in the system. When the agent is configured to compare a value with some threshold, the agent can generate an alert when the threshold is exceeded.

The moment an alert is generated, the agent can simply display the alert, but it can also send out a syslog message, or it can run some CLI commands and capture the output. These actions are defined in the script.

Change Agent Parameters

To see an example alert, the default parameters that control the threshold of the built-in system monitor will be changed to extreme low values, so that will trigger an alert.

16. In the **Agent Details** pane, click the **pen** icon to edit the Agent settings.

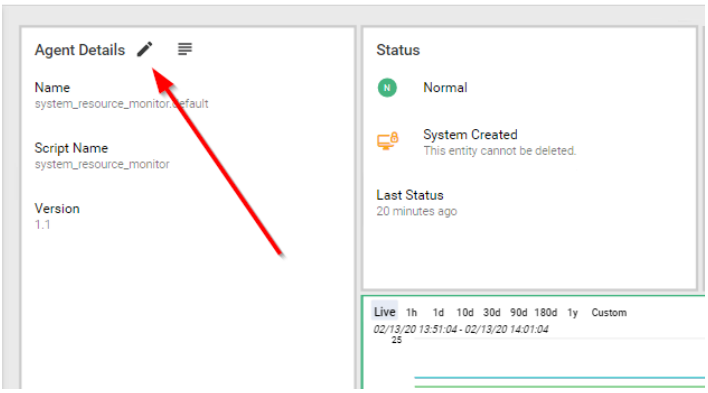

17. Under these settings, it is possible to enable or disable the agent. It is also possible to review or update the agent parameters.
18. Update these parameters for the Average CPU/Memory Utilization, expressed as a percentage:
    - short_term_high_threshold**:** **2**
    - short_term_normal_threshold**:** **1**

> **NOTE:** By setting the high_threshold for the average CPU/Memory to an extremely low number, the NAE system monitor agent will go into an Alert state.

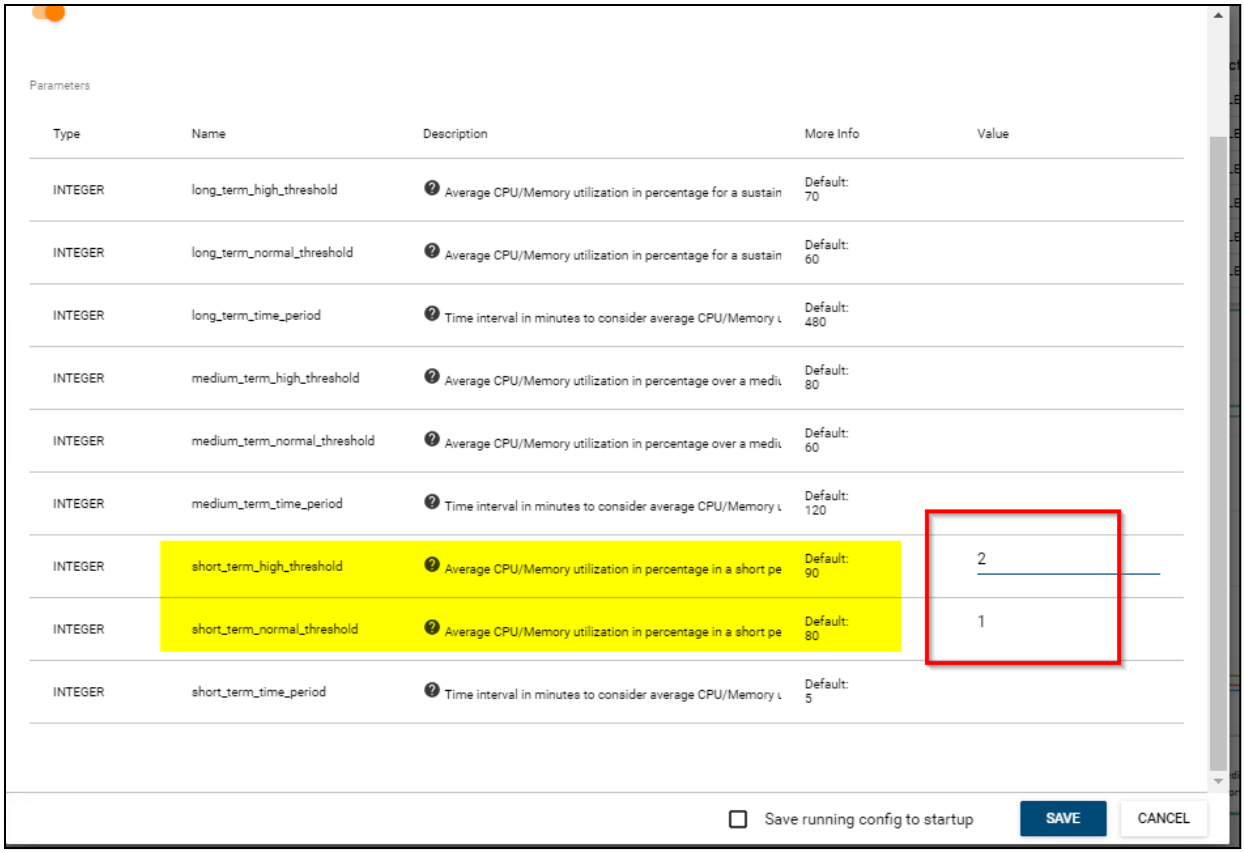

19. Click **Save** to save the settings. Next click **Close** to close the pop-up window.

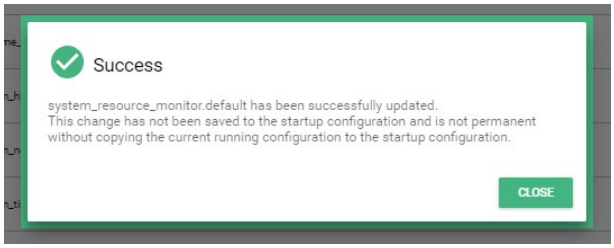

20. In about 60 seconds, a minor alert should be generated. The alert will be displayed at the top of the screen in the top right corner, in the agent details' Status pane to the right of the Agent Details pane, and in the graph bar (make sure **Live** is selected in the graph), displayed as an orange line with an orange triangle. (Note that the color of the alert in the graph corresponds to the alert level; in this example, since the alert categorized as *minor*, the color is orange).

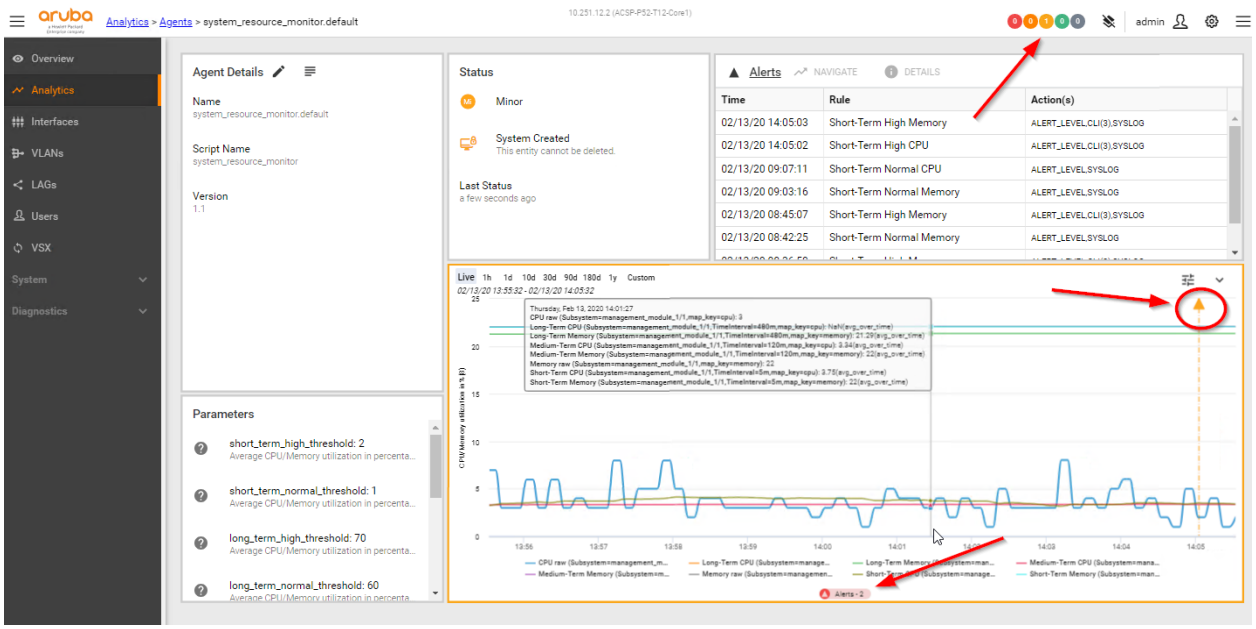

21. Open the details of the alert. This can be done either by clicking on the 'triangle' in the graph or by clicking on the alert in the alert list, and then clicking **DETAILS**.

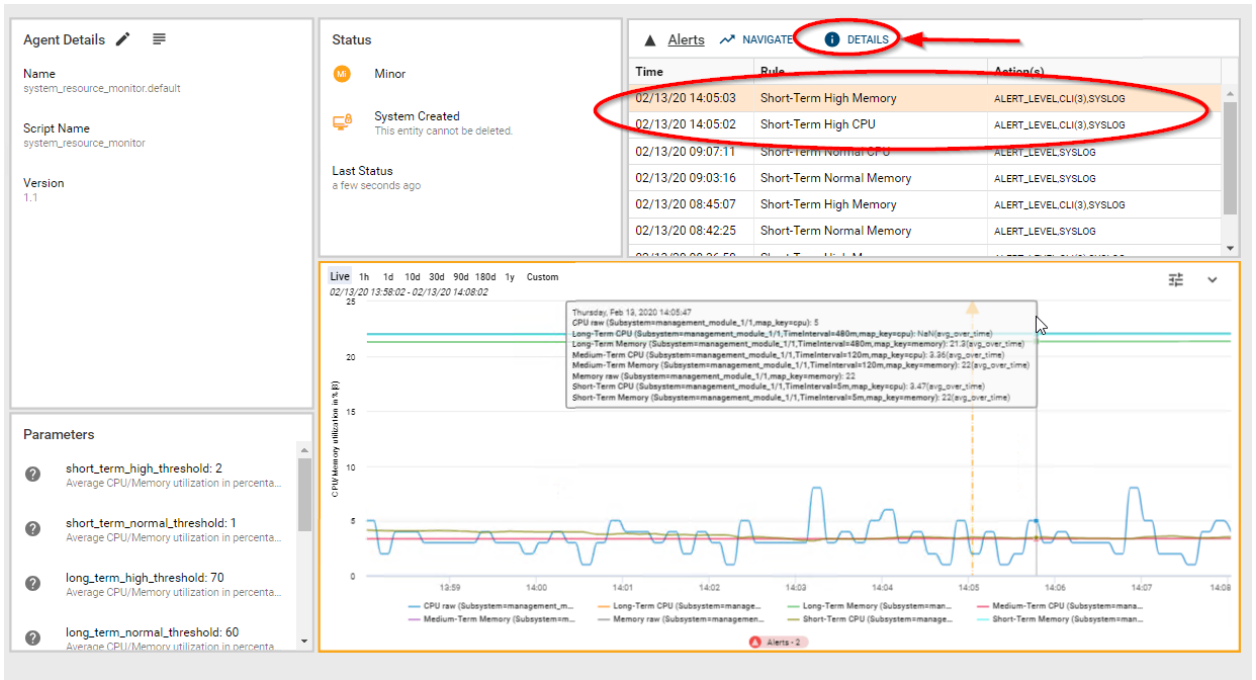

22. In the alert details, the administrator can see detailed information about the monitor that has triggered the alert: in this case, either high memory or high CPU utilization will be indicated. In this script, an alert will also capture CLI information from the switch, where this output is saved as part of the alert. This is very

convenient, since it allows the administrator to see additional information about the state of the switch at that point in time.

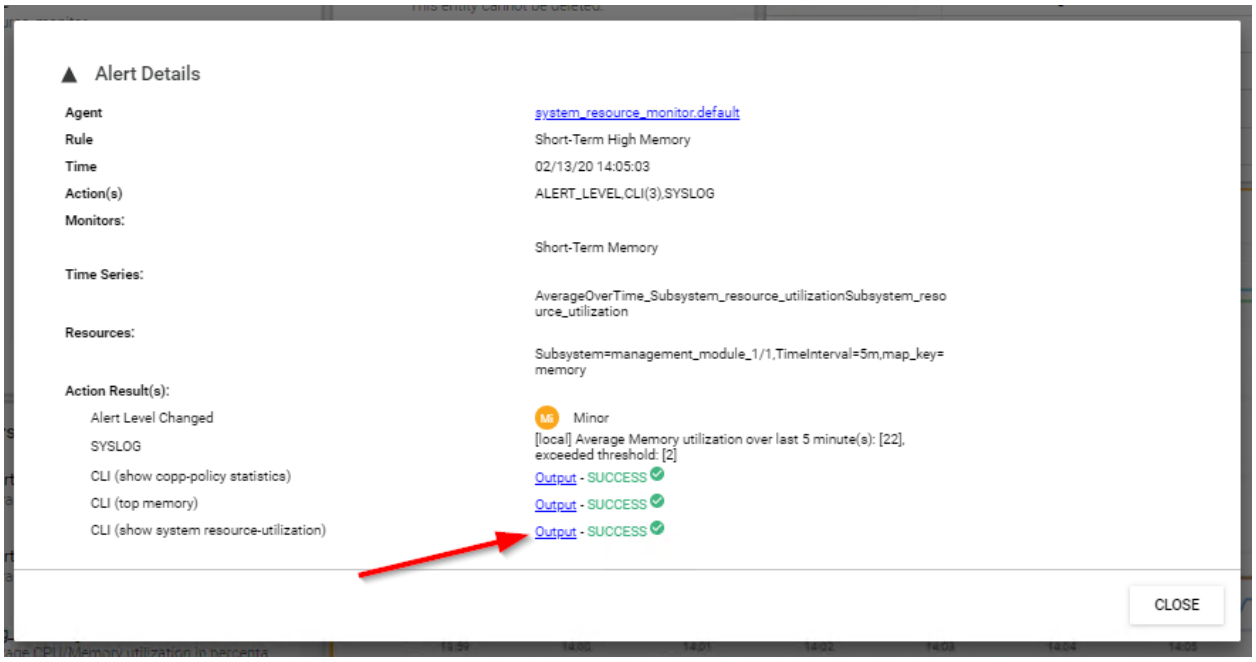

23. Click on the various 'output' links to see the details, this example shows the 'show system resource-utilization' output (click **Back** to return to the Alert Details window).

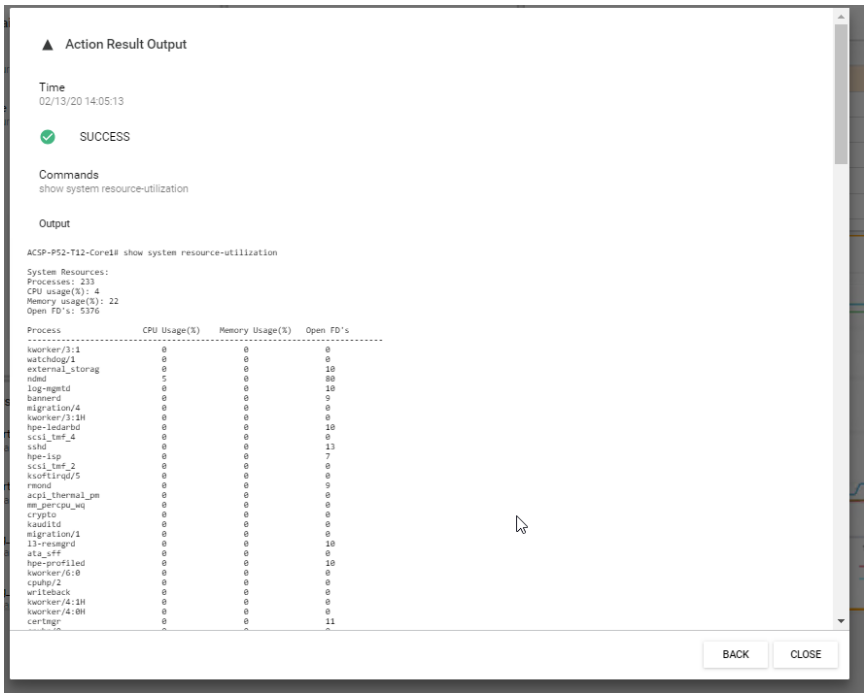

24. Close the CLI details and the alert.
25. Edit the **pen** icon for Agent Details again, and assign high custom values, so the alert can be cleared again. Then click **Save** and **Close** to confirm.
    - short_term_high_threshold: **91**
    - short_term_normal_threshold: **81**

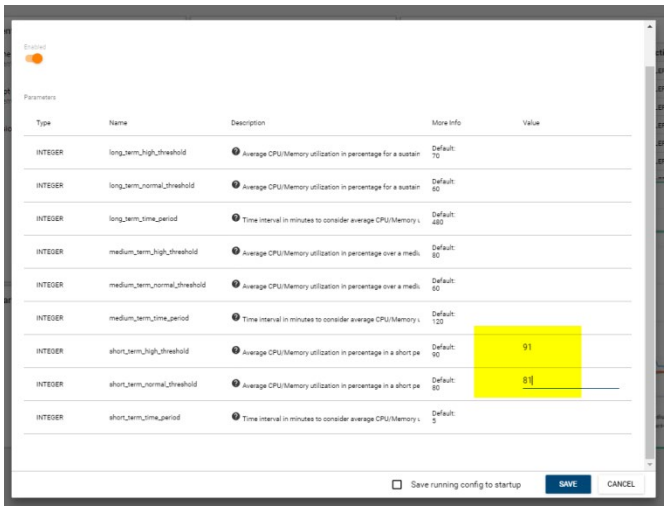

NOTE: It may take a few minutes for the alert to clear again (where the Status returns to a *normal* state). There is no need to wait for this.

Task 3: Add a New NAE Script and Agent

Diagram

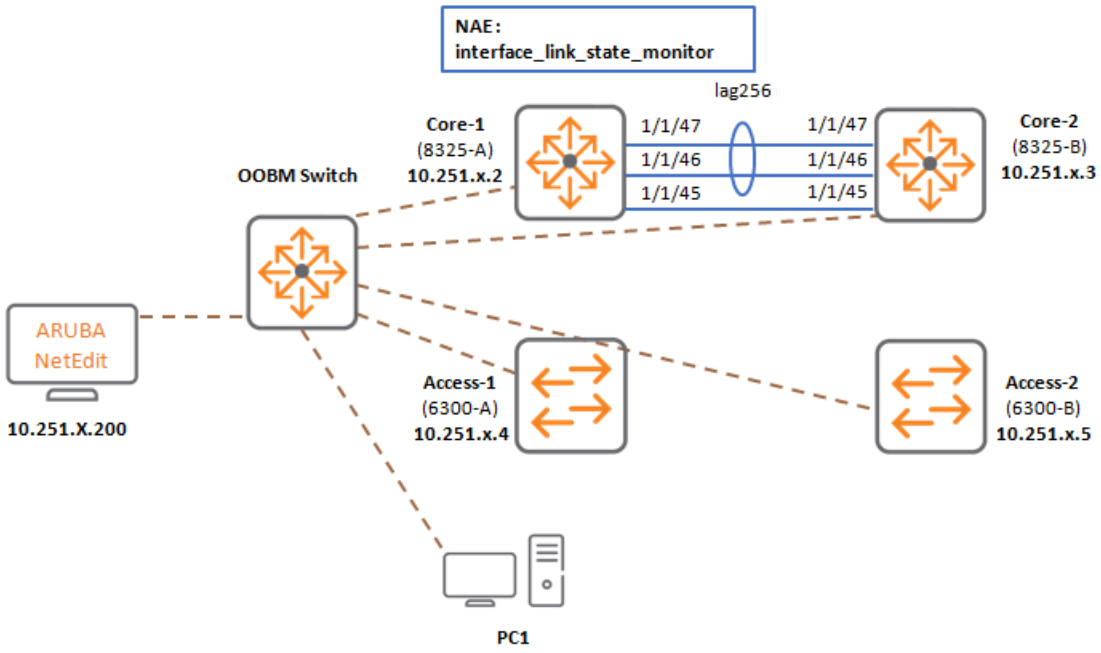

Objectives

In this task, a new script and agent will be added to Core1 to monitor link state changes on the switch.

Steps

Review Core1-to-Core2 Switch Link

In the NetEdit lab activity, a keep-alive IP link was previously defined between the Core1 and the Core2 switches using port 1/1/45. This link will be monitored using NAE.

First review or adjust this configuration between the Core1 and the Core2 switches.

Core1

1. Open a terminal on Core1 and enter Configuration mode.
2. Review the configuration of port 1/1/45.

```
ICX-Tx-Core1(config)# show run int 1/1/45
interface 1/1/45
    no shutdown
    vrf attach KA
    ip address 192.168.0.0/31
    exit
```

3. Verify reachability to Core2 using the KA VRF.

```
ICX-Tx-Core1(config)# do ping 192.168.0.1 vrf KA
PING 192.168.0.1 (192.168.0.1) 100(128) bytes of data.
108 bytes from 192.168.0.1: icmp_seq=1 ttl=64 time=0.133 ms
108 bytes from 192.168.0.1: icmp_seq=2 ttl=64 time=0.150 ms
108 bytes from 192.168.0.1: icmp_seq=3 ttl=64 time=0.184 ms
108 bytes from 192.168.0.1: icmp_seq=4 ttl=64 time=0.227 ms
108 bytes from 192.168.0.1: icmp_seq=5 ttl=64 time=0.242 ms

--- 192.168.0.1 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4117ms
rtt min/avg/max/mdev = 0.133/0.187/0.242/0.043 ms
```

Core2

4. **Optional step**: In case the ping did not work, review/update the Core2 configuration.

```
vrf KA
interface 1/1/45
    no shutdown
    vrf attach KA
    ip address 192.168.0.1/31
```

Install the new script

5. Using PC1, in the switch web UI, navigate to **Analytics > Scripts** .
6. Click the **Upload** link to install a new script.

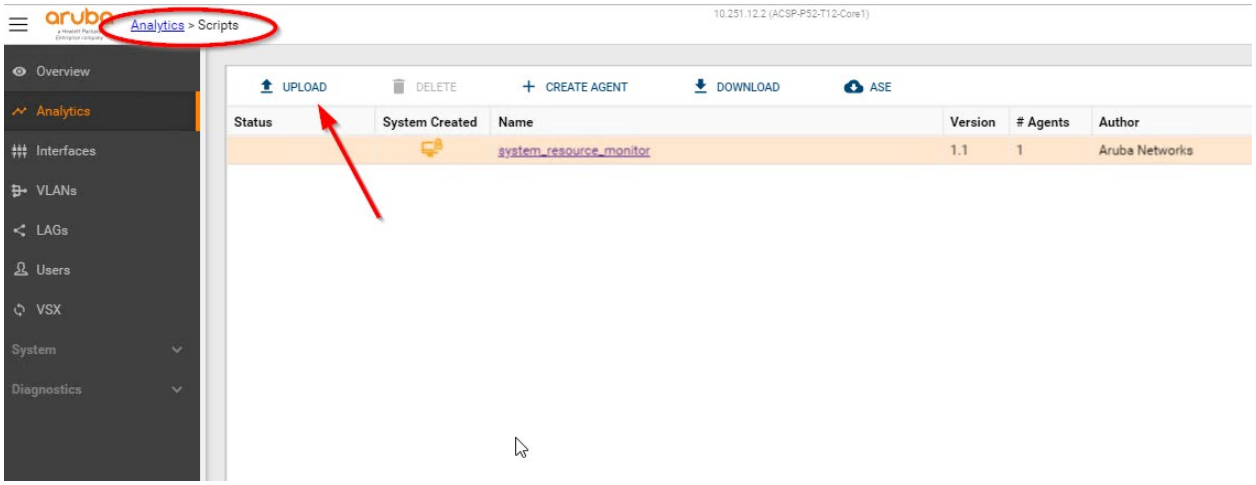

7. Click Browse and open the **ICX-Files** folder on the desktop, open the **nae** sub-folder and select the **'interface_link_state_monitor.1.0.py'** script, then click **Next.**

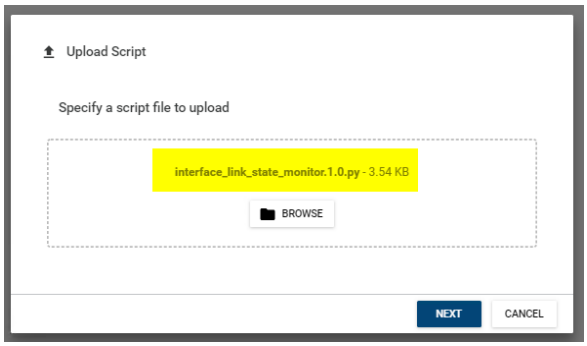

8. Review the script manifest data and click **Upload.**

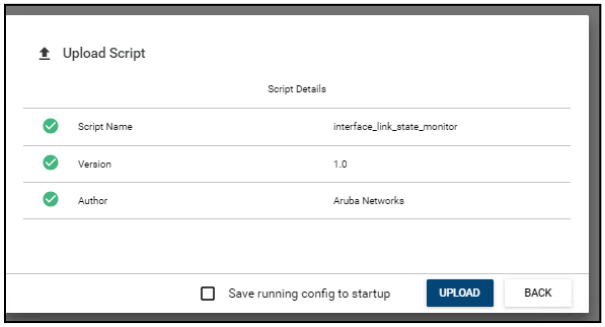

9. Click **Close** for the confirmation, the script will be shown in the list.

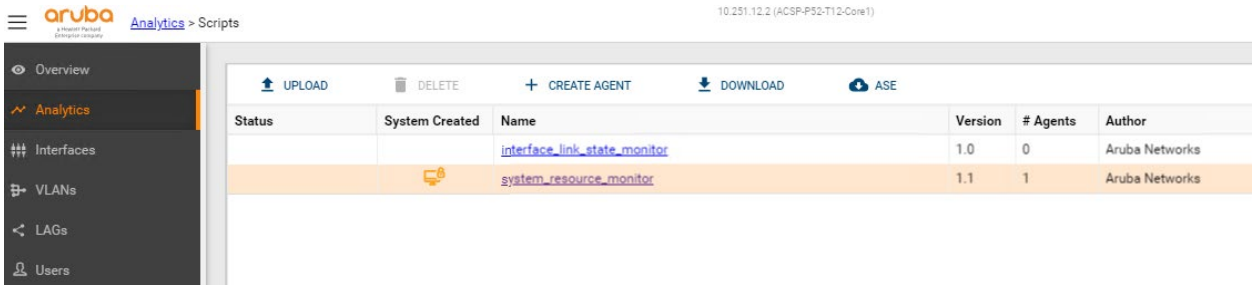

Create an Agent based on the new script

10. Navigate to **Analytics > Agents** and define a new agent by clicking **Create**.

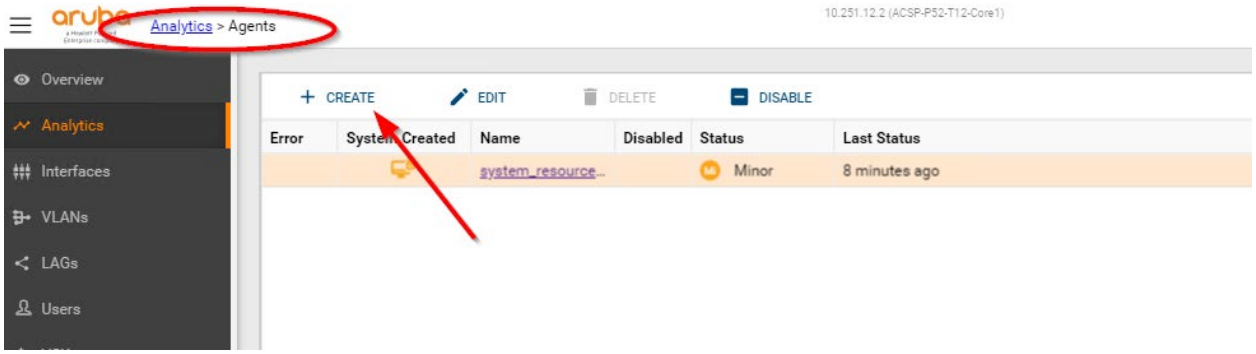

11. Select the **'interface_link_state_monitor'** from the list, and enter an agent name **iface-link-state-mon.** Then click **Create.**

12. Click **Close** in the confirmation window.
13. Navigate to **Analytics** to see the new Agent in the Dashboard.

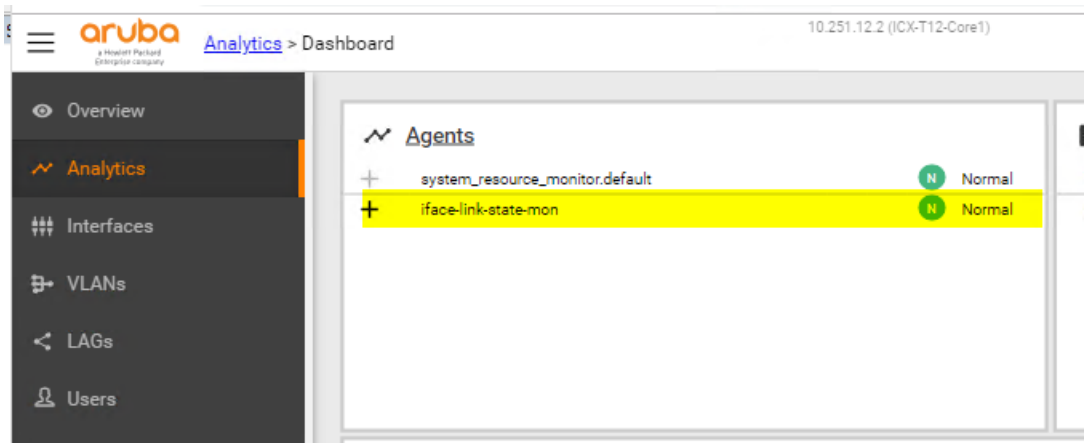

14. Select the name of the agent to see the details.

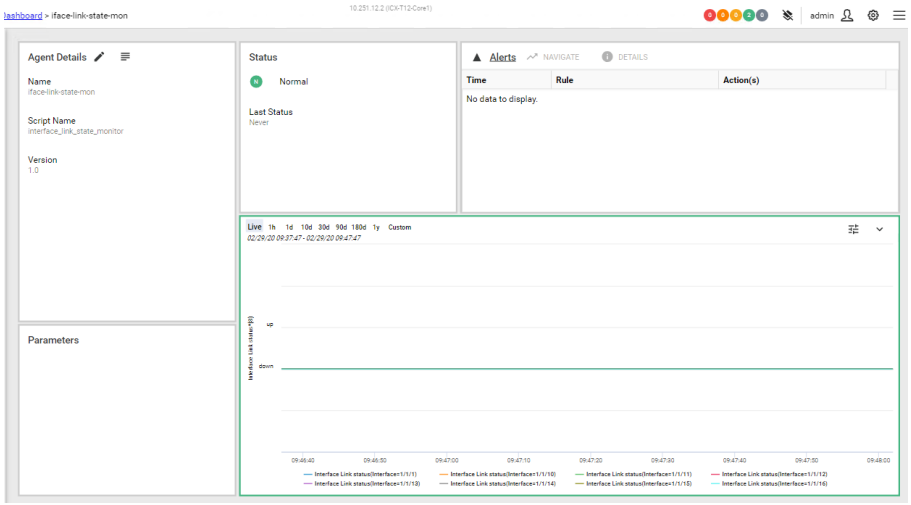

NOTE: This script does not have any parameters; it simply monitors all the interface transition from up to down.

Test the link state agent

15. On the Core1 switch, disable interface 1/1/45 and save the configuration.

```
ICX-Tx-Core1(config-if)# interface 1/1/45
ICX-Tx-Core1(config-if)# shutdown
```

16. A few moments later, save the configuration.

```
ICX-Tx-Core1(config-if)# write mem
Configuration changes will take time to process, please be patient.
```

17. Return to the switch Analytics Dashboard. The triangle will indicate the moment the alert was detected. The 'diamond' will indicate that a new configuration was saved.

> **NOTE:** The screenshot shows some previous alerts as well, these can be ignored.

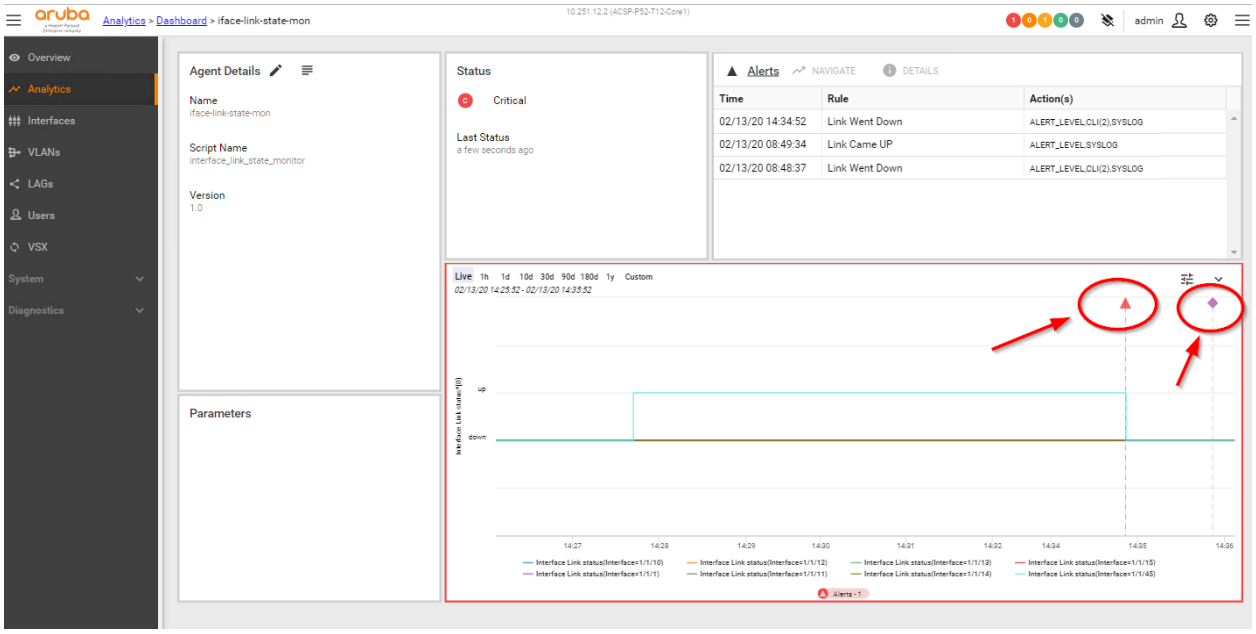

18. Click the **red triangle** in the graph to see the alert details. Review the captured CLI details by clicking the **Output** hyperlinks.

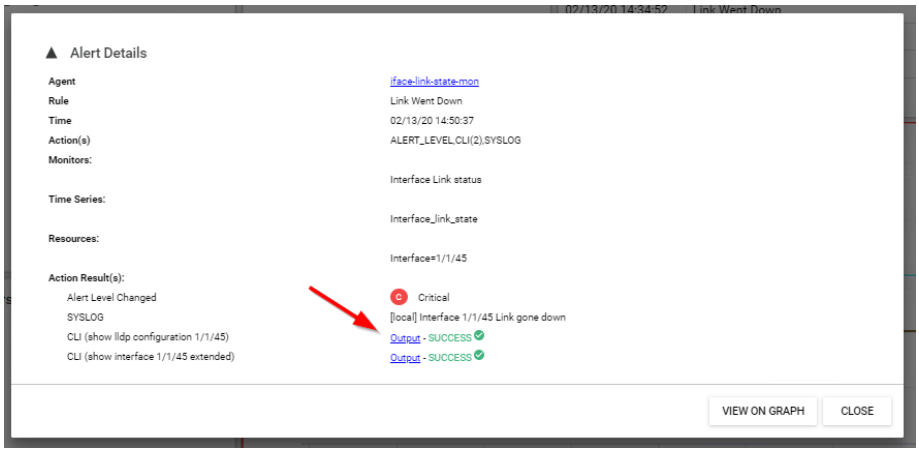

19. Example CLI details.

20. Click the **purple diamond** icon in the graph to see the details of the configuration change. This can make it easy for an administrator to correlate an alert to a configuration change.

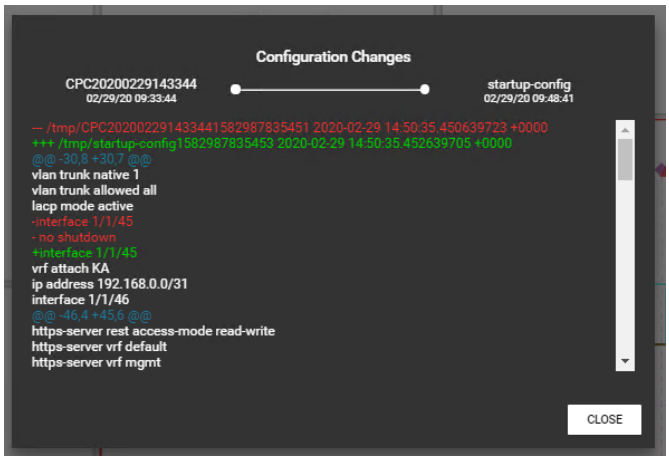

Restore the interface

21. On the Core1 switch, enable interface 1/1/45 and save the configuration.

```
ICX-Tx-Core1(config)# int 1/1/45
ICX-Tx-Core1(config-if)# no shutdown
ICX-Tx-Core1(config-if)# write mem
Configuration changes will take time to process, please be patient.
```

22. Verify that the agent state is back to normal in the Analytics Dashboard. The Status window pane should show a green icon and display “Normal”. In the Alerts section you should see an even listed that the link is now up (“Link Came Up”). In the graph, you should see a green arrow indicating the interface status change.

Also, the graph will display a bright blue line indicating that the status is up (examine the graph key for the interfaces at the bottom of the graph).

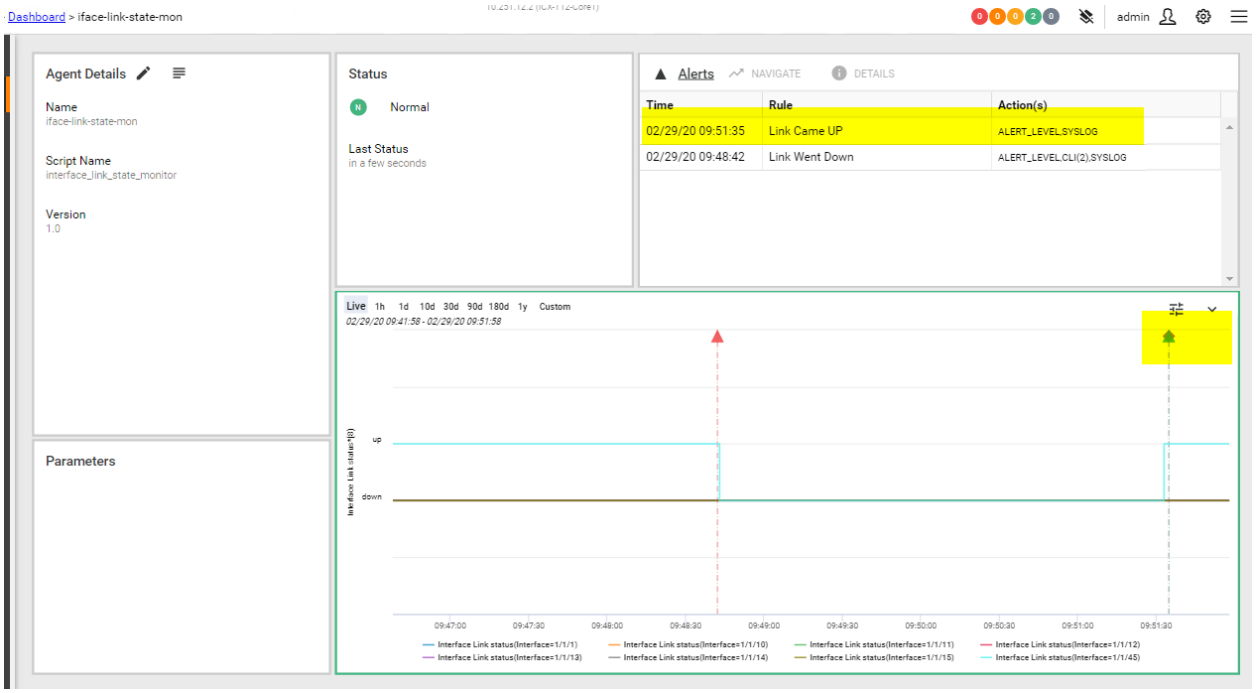

23. In the Agent Details pane, select the **pen** icon to edit the Agent. Move the slider to the off state to disable the agent. Click **Save** to save the settings; click **Close** for the confirmation.

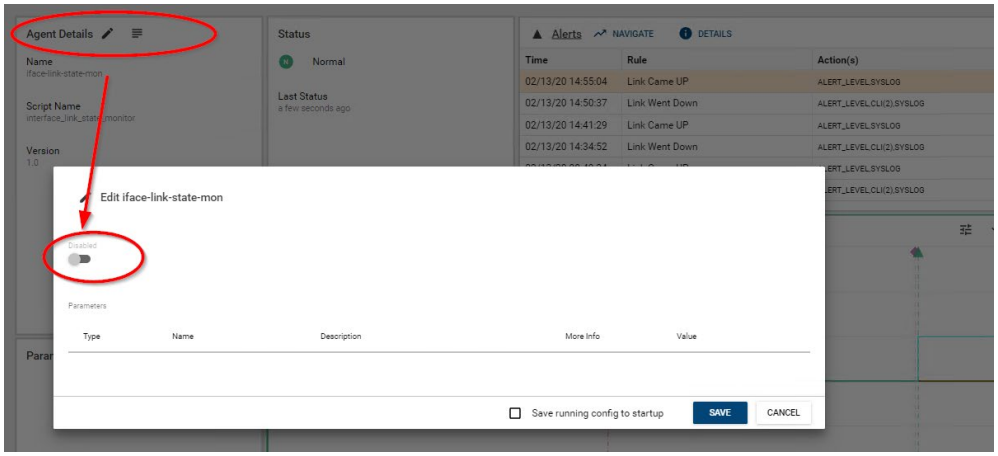

Optional Task 4: Connectivity Check

This task is **optional** and can be done if time permits. Check with your instructor.

If you skip this task, you can move to the next Task.

Diagram

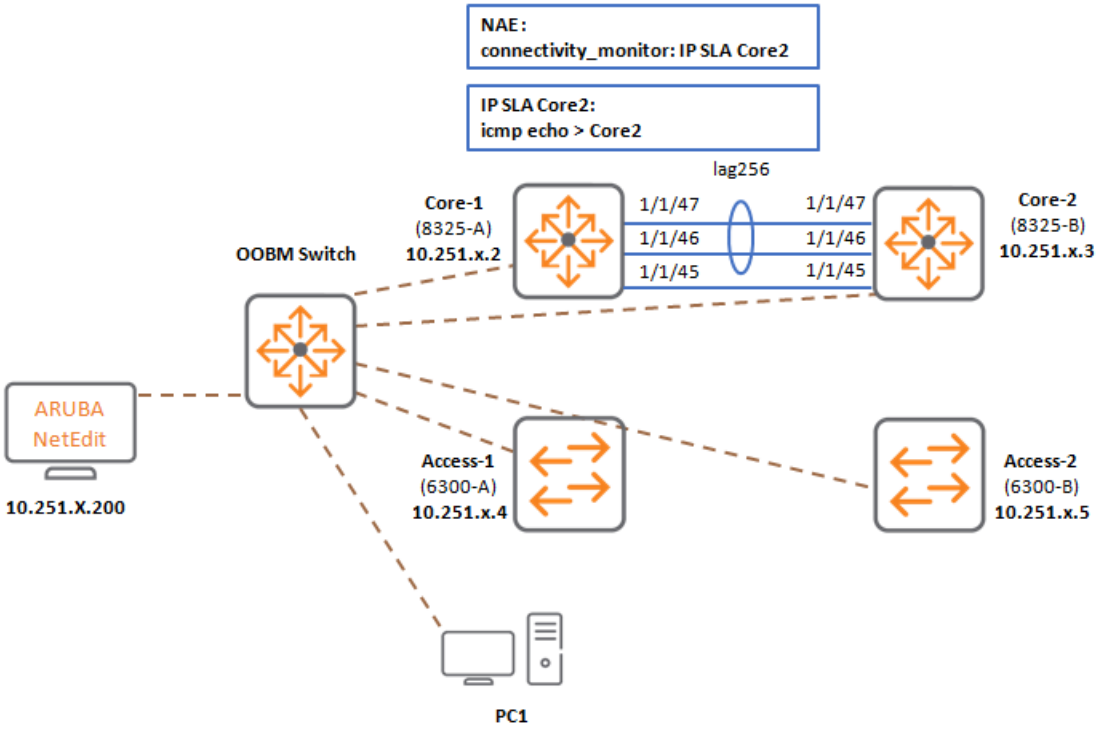

Objectives

In this task, a connectivity check script will be installed and configured. This script relies on the IP-SLA (Service Level Agreement) test feature set of the switch. An IP-SLA object in the switch is a 'test client' that can perform and repeat network tests, such as sending ICMP echo requests or HTTP get requests, and keep track of the results.

In this task, an IP-SLA test will be created that will verify ICMP connectivity with the peer switch Core2. With the IP-SLA in place, the NAE script will now keep track of the IP-SLA object and will be able to alert and monitor the results of the IP-SLA object.

Steps

Add IP-SLA on Core1

1. On the terminal of Core1, enter Configuration mode.
2. Configure a new IP-SLA to send an ICMP echo to the IP of Core2 every 5 seconds. Then start the test.

```
ICX-Tx-Core1(config)# ip-sla ping-core2
ICX-Tx-Core1(config-ip-sla-ping-core2)# vrf KA
ICX-Tx-Core1(config-ip-sla-ping-core2)# icmp-echo 192.168.0.1 probe-interval 5
ICX-Tx-Core1(config-ip-sla-ping-core2)# start-test
ICX-Tx-Core1(config-ip-sla-ping-core2)# exit
```

3. Wait about 10 seconds, then check the results of the IP-SLA test. The latest probe results should show that 1 packet was received, this is the ICMP response from Core2.

```
ICX-Tx-Core1(config)# show ip-sla all
SLA Name                : ping-core2
Status                  : running
...
IP-SLA session status
IP-SLA Name                          : ping-core2
IP-SLA Type                          : icmp-echo
Destination Host Name/IP Address     : 192.168.0.1
Source IP Address/IFName             :
Status                               : running

IP-SLA Session Cumulative Counters
Total Probes Transmitted             : 4
Probes Timed-out                     : 0
Bind Error                           : 0
Destination Address Unreachable      : 0
DNS Resolution Failures              : 0
Reception Error                      : 0
Transmission Error                   : 0

IP-SLA Latest Probe Results
Last Probe Time                      : 2020 Feb 14 01:24:32
Packets Sent                         : 1
Packets Received                     : 1
Packet Loss in Test                  : 0.0000%

Minimum RTT(ms)                      : 0
Maximum RTT(ms)                      : 0
Average RTT(ms)                      : 0
DNS RTT(ms)                          :
```

This IP-SLA is now ready to be tracked by NAE.

Add script and agent for the Connectivity reachability test

First add a new script that can monitor IP-SLA objects

4. On PC1, use a web browser to connect to Core1. Navigate to **Analytics > Scripts**.
5. Upload the script 'reachability.1.1.py': click **UPLOAD** in the top right corner. Click **Browse**. Find and select the **reachability.1.1.py** script in the **ICX-Files\nae**

folder. Click **Next**. Examine the script information, including the name, which is “connectivity_monitor”. Notice that the author is “Aruba Networks”. Click **UPLOAD**. Finally, click **CLOSE** to complete the upload. Verify that you see the “connectivity_monitor” script.

6. Navigate to the **Analytics > Agents** window.
7. Click **CREATE** to create a new Agent. Enter the following information:
   - Script **connectivity_monitor**
   - Agent Name **connectivity-ping-core2**

   For the parameters, enter the following:
   - Connectivity check rate **leave at the default (1min)**
   - IPSLA session name **ping-core2**

NOTE: The IPSLA session name **must** match with the IP-SLA object name that was created in the previous step, that is 'ping-core2' in this lab guide

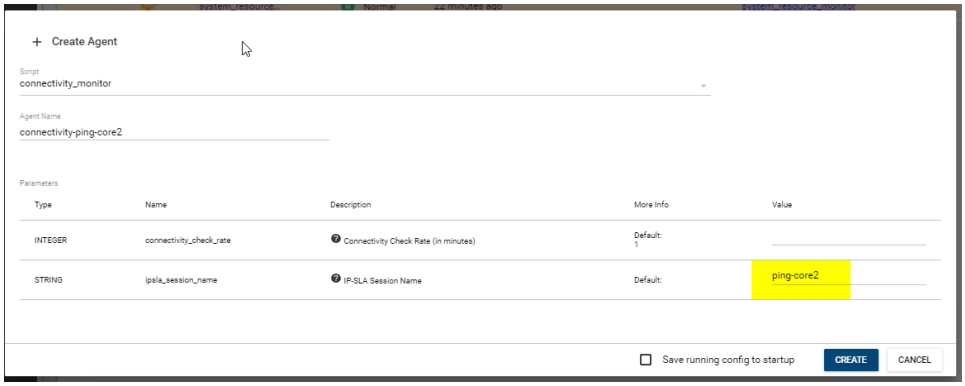

.

8. Click **CREATE** to submit the new agent settings, and click **CLOSE** for the confirmation.
9. Once the agent has been created, select the agent to see the details. It should appear 'normal' (green).

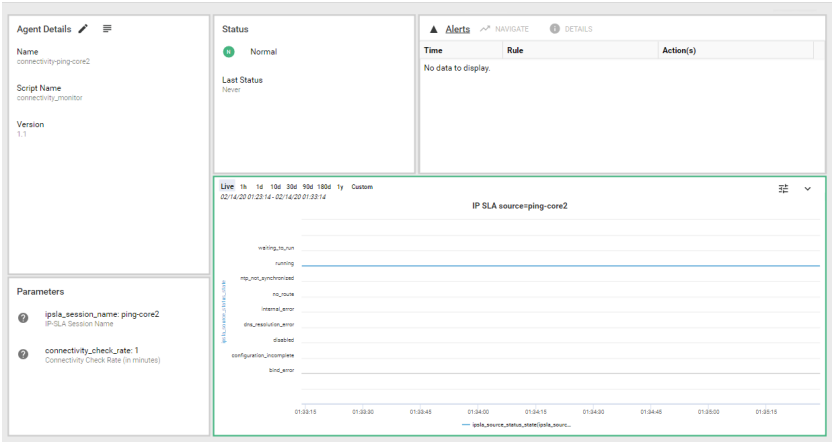

Testing the Connectivity Test Script

10. Open a terminal to Core2, and remove the IP address on 1/1/45. This will not cause a 'link down', as the first 'link status' script would be able to detect. However, the ping will no longer work, so this connectivity script should detect it.

```
ICX-Tx-Core2# configure terminal
ICX-Tx-Core2(config)# interface 1/1/45
ICX-Tx-Core2(config-if)# no ip address 192.168.0.1/31
```

11. Switch to the web UI of Core1. Access **Analytics** in the left navigation pane. After 1-2 minutes, the connectivity-ping-core2 agent under the Agents window pane should show a minor alarm and then a critical alarm. Click **connectivity-ping-core2** to access the agent details screen.

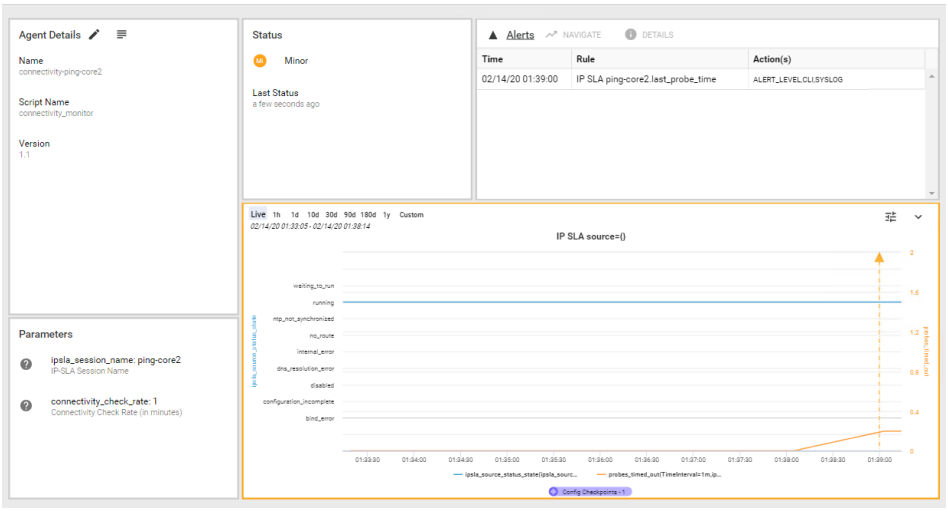

12. In the 'Alerts' window, select the most recent alert for the IP SLA event and then click **DETAILS**. Check the CLI output by clicking the Output hyperlink. The output should be familiar from when you executed the **show** command previously after configuring and verifying SLA from the CLI. NAE captures the IP-SLA state at the time of the alert.

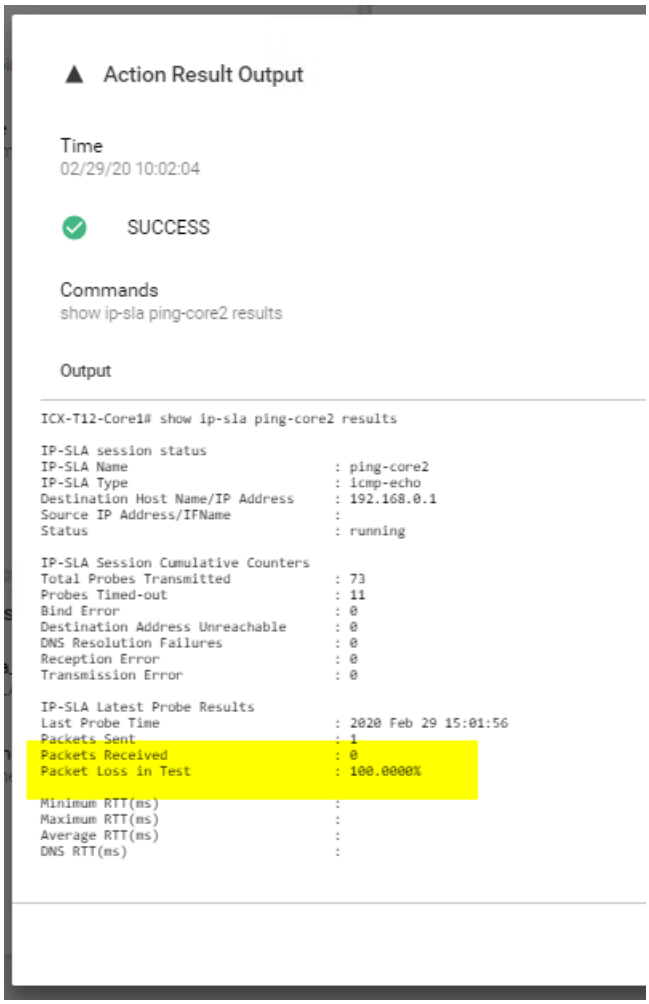

13. Click **CLOSE.**

Recover the Alert

14. Re-assign the IP address on Core2.

```
ICX-Tx-Core2(config-if)# ip address 192.168.0.1/31
```

15. Re-access the Analytics Dashboard in the web UI. After about 1 minute, the alert should be cleared. Notice that the Status window pane will show a green icon and “Normal”. The graph will indicate a green arrow for when the condition was recovered.

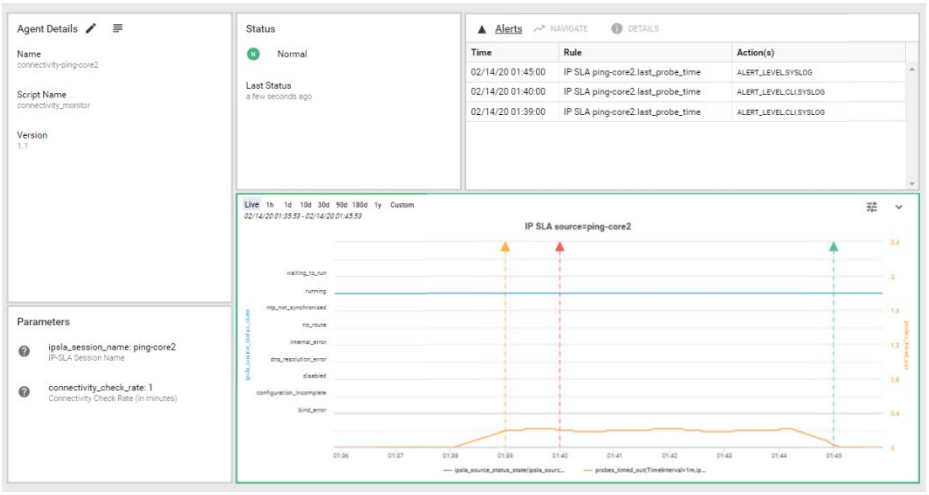

This concludes the optional Connectivity agent task.

Task 5: Review the NAE agent in the switch configuration file

Objectives

In this task, the NAE script and agent settings will be reviewed in the switch configuration.

Steps

Review the running-configuration

NAE script and agent information is stored in the switch's configuration file. This ensures that a backup/restore of the switch configuration will also maintain the configured monitoring solution.

1. Open a terminal connection to Core1 and review the installed scripts.

> **NOTE:** The listed scripts may be different in your output based on the completion of the previous, optional task. This can be ignored.

```
ICX-Tx-Core1# show nae-script
----------------------------------------------------------------
Script Name                     Version   Origin    Status
----------------------------------------------------------------
connectivity_monitor            1.1       user      VALIDATED
interface_link_state_monitor    1.0       user      VALIDATED
system_resource_monitor         1.1       system    VALIDATED
```

2. Review the NAE agents.

> **NOTE:** The listed agents may be different in your output based on the completion of the previous, optional task: this can be ignored.

```
ICX-Tx-Core1# show nae-agent
------------------------------------------------------------------------------------
-------------------------------------
Agent Name                          Script Name                     Version
Origin    Disabled  Status    Error
------------------------------------------------------------------------------------
-------------------------------------
connectivity-ping-core2             connectivity_monitor            1.1
user      false     NORMAL    NONE
iface-link-state-mon                interface_link_state_monitor    1.0
user      true      UNKNOWN   NONE
system_resource_monitor.default     system_resource_monitor         1.1
system    false     NORMAL    NONE
```

3. Review the running-configuration for lines that 'include' the string 'nae' . The script and the agent parameters are Base64 encoded in the configuration file.

```
ICX-Tx-Core1# show run | include nae
```

4. Example extract, the font size for the encode section has been reduced.

```
...
nae-script interface_link_state_monitor false
Iy0qLSBjb2Rpbmc6IHV0Zi04IC0qLQ0KIw0KI0NvcHlyaWdodCAoYykgMjAxNyBIZXdsZXR0IFBhY2thcmQgRW50ZXJwcmlzZSBEZXZlbG9wbWVudCBM
UA0KIw0KI0xpY2Vuc2VkIHVuZGVyIHRoZSBBcGFjaGUgTGljZW5zZSwgVmVyc2lvbiAyLjAgKHRoZSAiTGljZW5zZSIpOw0KI3lvdSBtYXkgbm90IHVz
ZSB0aGlzIGZpbGUgZXhjZXB0IGluIGNvbXBsaWFuY2Ugd2l0aCB0aGUgTGljZW5zZS4NCiNZb3UgbWF5IG9idGFpbiBhIGNvcHkgb2YgdGhlIExpY2Vu
c2UgYXQNCiMNCiNodHRwOi8vd3d3LmFwYWNoZS5vcmcvbGljZW5zZXMvTElDRU5TRS0yLjANCiMNCiNVbmxlc3MgcmVxdWlyZWQgYnkgYXBwbGljYWJs
ZSBsYXcgb3IgYWdyZWVkIHRvIGluIHdyaXRpbmcsDQojc29mdHdhcmUgZGlzdHJpYnV0ZWQgdW5kZXIgdGhlIExpY2Vuc2UgaXMgZGlzdHJpYnV0ZWQg
b24gYW4NCiMiQVMgSVMiIEJBU0lTLCBXSVRIT1VUIFdBUlJBTlRJRVMgT1IgQ09ORElUSU9OUyBPUiBBTlkNCiNLSU5ELCBlaXRoZXIgZXhwcmVzcyBv
ciBpbXBsaWVkLiBTZWUgdGhlIExpY2Vuc2UgZm9yIHRoZQ0KI3NwZWNpZmljIGxhbmd1YWdlIGdvdmVybmluZyBwZXJtaXNzaW9ucyBhbmQgbGltaXRh
dGlvbnMNCiN1bmRlciB0aGUgTGljZW5zZS4NCg0KTWFuaWZlc3QgPSB7DQogICAgJ05hbWUnOiAnaW50ZXJmYWNlX2xpbmtfc3RhdGVfbW9uaXRvcics
DQogICAgJ0Rlc2NyaXB0aW9uJzogJ0ludGVyZmFjZSBMaW5rIFN0YXR1cyBNb25pdG9yaW5nIHVzaW5nIEFnLA0K
ICAgICdBdXRob3InOiAnQXJ1YmEgTmV0d29ya3MnDQp9DQoNCmNsYXNzIFBvbGljeShOQUUpOg0KDQogICAgZGVmIF9faW5pdF9fKHNlbGYpOg0KI0lu
dGVyZmFjZSBzdGF0dXMNCiAgICAgICAgXJ1YmEgTmV0d29ya3MnDQp9DQoNCmNsYXNzIFBvbGljeShOQUUpOg0KDQogICAgZGVmIF9faW5pdF9fKHNlbGYpOg0KI0lu
ICAgICAgc2VsZi5tMSA9IE1vbml0b3IoDQogICAgICAgICAgICB1cmkxLA0KICAgICAgICAgICAgJ0ludGVyZmFjZSBMaW5rIHN0YXR1cycpDQogICAgICAg
ICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAg
YWN0aW9uX2ludGVyZmFjZV91cCkNCg0KI3ZhcmlhYmxlcw0KICAgICAgICBzZWxmLmxvZ2dlci5kZWJ1ZygiPT09PT09PT09PT09PSBEb3duIEQ0KICAg
ZWxmLmxvZ2dlci5kZWJ1ZygnaW50ZXJmYWNlX2lkIC0gJyArIGludGVyZmFjZV9pZCkNCg0KICAgICAgICBzZWxmLmxvZ2dlci5kZWJ1ZygiPT09PT09PT09
WydsaW5rc19kb3duJ10NCiAgICAgICAgICAgc2VsZi5sb2dnZXIuZGVidWcoJ2xpbmtzX2Rvd24gYmVmb3JlOiBbbJw0KICAgICAgICAgICAgICAgICAg
ICAgICArIGxpbmtzX2Rvd24gKyAnXScpDQogICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAg
ICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAg
J10gPSBsaW5rc19kb3duDQogICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAg
bicpDQogICAgICAgICAgICAgICBBY3Rpb25DTEkoInNob3cgbGxkcCBjb25maWd1cmF0aW9uICIgKyBpbnRlcmZhY2VfaWQpDQogICAgICAgICAgICAgICBBY3Rp
b25DTEkoInNob3cgaW50ZXJmYWNlICIgKyBpbnRlcmZhY2VfaWQgKyAiIGV4dGVuZGVkIikNCiAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAg
dmVsKCkgIT0gQWxlcnRMZXZlbC5DUklUSUNBTDoNCiAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAg
KQ0KICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAg
b3duICsgJ10nKQ0KICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAg
IGRlZiBhY3Rpb25faW50ZXJmYWNlX3VwKHNlbGYsIGV2ZW50KToNCiAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAg
PT09PT09PT09PT09PSIpDQogICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAg
WycgKyBsYWJlbCArICddJykNCiAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAg
bGYubG9nZ2VyLmRlYnVnKCdpbnRlcmZhY2VfaWQgLSAnICsgaW50ZXJmYWNlX2lkKQ0KICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAg
J2xpbmtzX2Rvd24nXQ0KICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAg
ICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAg
bmtzX2Rvd24gPSBsaW5rc19kb3duLnJlcGxhY2UoKGludGVyZmFjZV9pZCArICc6KQ0KICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAg
bmtzX2Rvd24nXSA9IGxpbmtzX2Rvd24NCiAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAg
ICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAg
ZmFjZV9pZCArICcgTGluayBjYW1lIHVwJykNCiAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAg
X2FsZXJ0X2xldmVsKCkgaXMgbm90IE5vbmU6DQogICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAgICAg
Zi5sb2dnZXIuZGVidWcoIj09PT09PT09PT09PT09PT0gL1VwID09PT09PT09PT09PT09PT0iKQ0K
...
```

5. The administrator can simply review the original script text with a Base64 decode of the script. Below is an online example **https://www.base64decode.org/** .

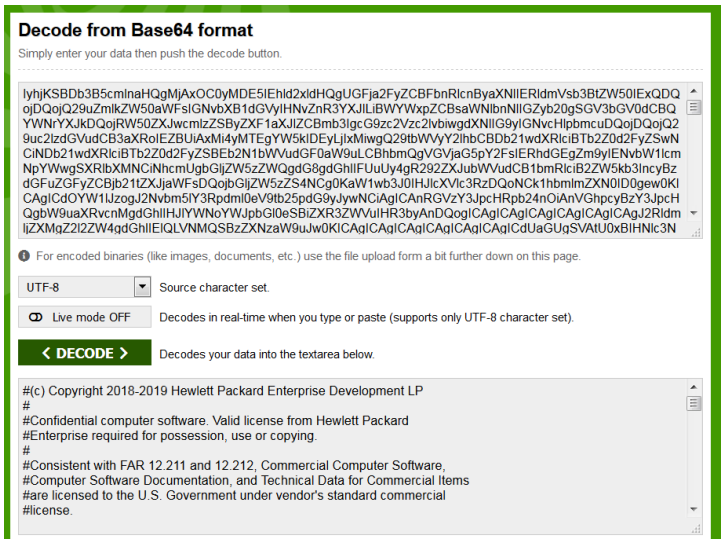

NOTE: This is informational only: there is no need to actually decode your script in this lab.

6. Each of the created agents based on the script, the agent state (false in the first line), and the parameters (base64 encoded) are also stored in the configuration file.

```
nae-agent connectivity_monitor connectivity-ping-core2 false
ipsla_session_name:cGluZy1jb3JlMg==

nae-agent interface_link_state_monitor iface-link-state-mon true

nae-agent system_resource_monitor system_resource_monitor.default false
short_term_high_threshold:OTE= short_term_normal_threshold:ODE=
```

Review the switch NAE capacities

7. The switch can handle several scripts and agents. To review the currently active amount of scripts and agents versus the platform limits, use the **show capacities-status nae** command.

```
ICX-Tx-Core1# show capacities-status nae

System Capacities Status: Filter NAE
Capacities Status Name                                                    Value Maximum
------------------------------------------------------------------------------------
Number of configured NAE agents currently active in the system              3      50
Number of configured NAE monitors currently active in the system           12     150
Number of configured NAE scripts currently active in the system             3      25
```

You have completed Lab 3!

Lab 04: Configuring VSX- AOS-CX Virtual Switching Extension

Lab Diagram

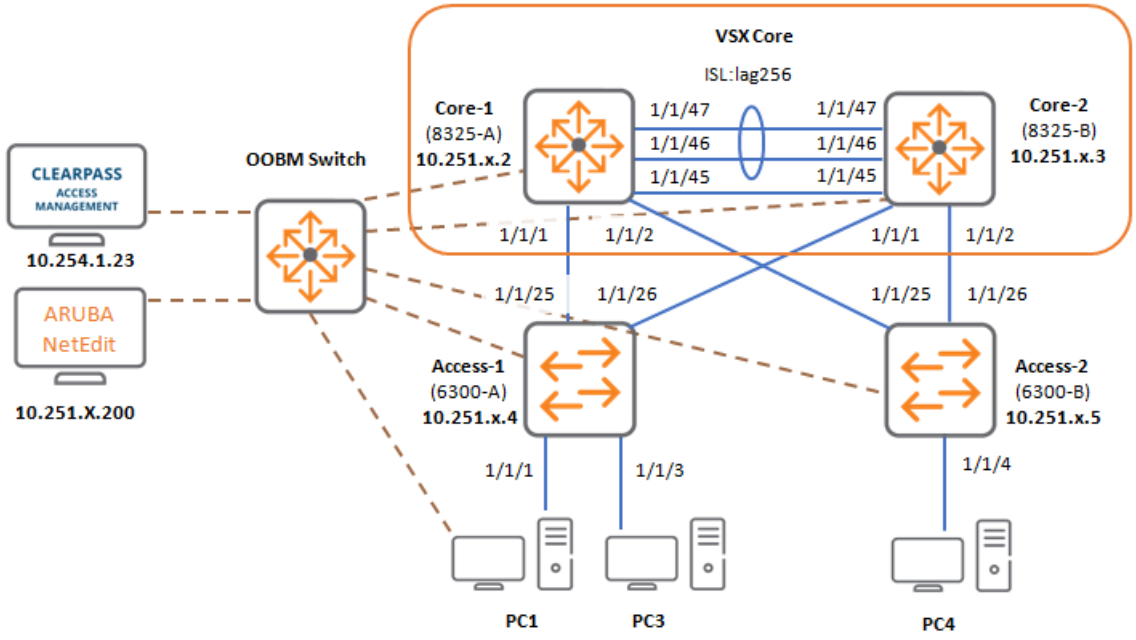

Overview

In this lab activity, VSX will be configured between the two core switches in the lab setup.

VSX has similar benefits as VSF; however, VSX also offers better high availability required in core and data center environments. VSX binds two ArubaOS-CX switches of the same model type to operate as one device for layer 2. VSX also operates as independent nodes for layer 3.

Objectives

- Configure VSX
- Understand the VSX configuration synchronization feature
- Configure VSX to peer devices using VSX LAG
- Understand and configure VSX Active Gateway (first hop redundancy)
- Configure DHCP relay
- Configure and understand split-system protection

> **IMPORTANT:** The checkpoint of this lab activity will be used as the base configuration for many other lab activities. Make sure to verify your steps and pay attention on which device you are working on.

Task 1: Verify Lab Start Configuration

Objectives

– Load the checkpoint of Lab 02 - NetEdit configuration

Steps

1. Open a console connection to the 6300A. Login using **admin**/**aruba123**

```
ICX-Tx-Access1# copy checkpoint icx-lab02-netedit running-config
Configuration changes will take time to process, please be patient.
ICX-Tx-Access1#
```

2. Open a console connection to the 6300B. Login using **admin**/**aruba123**

```
ICX-Tx-Access2# copy checkpoint icx-lab02-netedit running-config
Configuration changes will take time to process, please be patient.
ICX-Tx-Access2#
```

3. Open a console connection to the 8325A. Login using **admin**/**aruba123**

```
ICX-Tx-Core1# copy checkpoint icx-lab02-netedit running-config
Configuration changes will take time to process, please be patient.
ICX-Tx-Core1#
```

4. Open a console connection to the 8325B. Login using **admin**/**aruba123**

```
ICX-Tx-Core2# copy checkpoint icx-lab02-netedit running-config
Configuration changes will take time to process, please be patient.
ICX-Tx-Core2#
```

Task 2: VSX Basic Setup

Diagram

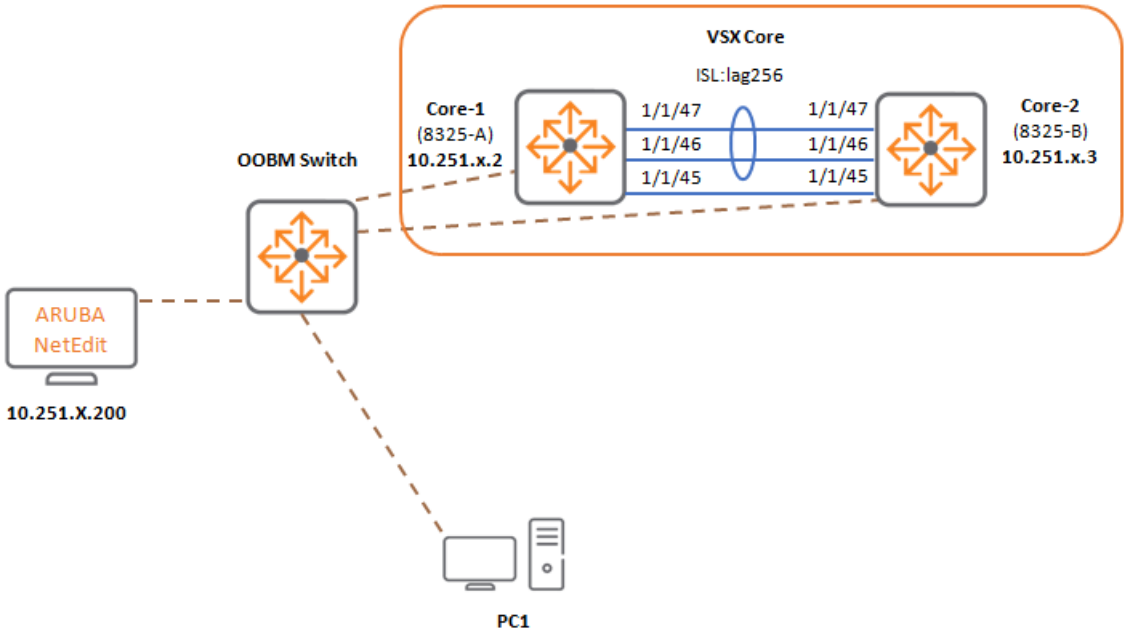

Objectives

- Configure a link aggregation that will be used as the Inter Switch Link (ISL)
- Configure the keepalive link
- Setup the ISL link
- Configure the VSX roles
- Verify the status of the VSX cluster

Steps

Core1

1. Open a terminal session to Core1 and enter Configuration mode.
2. Take note of the current version.

```
ICX-Tx-Core1(config)# show version
```

3. Review the LAG256 to Core2. This LAG was defined in the NetEdit lab activity.

Some best practice guidelines:

- Use the last LAG number supported by the switch. In this case, number 256.
- Allow all the VLANs on the ISL.
- Enable LACP, use the standard LACP timers (30 seconds).

```
ICX-Tx-Core1# show running-config interface lag256
```

```
interface lag 256
    no shutdown
    no routing
    vlan trunk native 1
    vlan trunk allowed all
    lacp mode active
    exit
```

4. Enable support for jumbo frames on the ISL ports. This will be required when using dynamic segmentation.

```
ICX-Tx-Core1# configure terminal
ICX-Tx-Core1(config)# interface 1/1/46-1/1/47
ICX-Tx-Core1(config-if-<1/1/46-1/1/47>)# mtu 9198
ICX-Tx-Core1(config-if-<1/1/46-1/1/47>)# exit
```

5. Review the physical ports of the LAG256 and verify jumbo frame support is enabled.

```
ICX-Tx-Core1(config)# show run int 1/1/46
interface 1/1/46
    no shutdown
    mtu 9198
    lag 256
    exit
```

```
ICX-Tx-Core1(config)# show run int 1/1/47
interface 1/1/47
    no shutdown
    mtu 9198
    lag 256
    exit
```

Core2

6. Open a console to Core2 and enter Configuration mode.
7. Check the version, make sure this is the same version as Core1. Contact your instructor if it would be different.

```
ICX-Tx-Core2(config)# show version
```

8. Enable jumbo frame support on the ISL ports 1/1/46 and 1/1/47.

```
ICX-Tx-Core2# configure terminal
ICX-Tx-Core2(config)# interface 1/1/46-1/1/47
ICX-Tx-Core2(config-if-<1/1/46-1/1/47>)# mtu 9198
ICX-Tx-Core2(config-if-<1/1/46-1/1/47>)# exit
```

Core1 - Review status

9. On Core1, review the status of the LAG between the Core1 and Core2 switches.

```
ICX-Tx-Core1(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  90:20:c2:bc:17:00 65534  256  up
1/1/47  lag256     48    1     ALFNCD  90:20:c2:bc:17:00 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  90:20:c2:bc:97:00 65534  256
1/1/47  lag256     48    1     ALFNCD  90:20:c2:bc:97:00 65534  256
```

Review the VSX Keepalive Link

The keepalive link is used to handle a split-brain scenario. This means that all the links between the VSX members are down, while both switches are still up.

VSX will make an IP connectivity check with its peer. Since this IP check has nothing to do with the rest of the IP network, it is recommended to put this IP link in a dedicated VRF, a separate routing table.

NOTE: The split-brain scenario will be tested in a later task.

Core1

10. On Core1, verify the connectivity to Core2 for this VRF. This VRF and IP have been defined in the NetEdit lab activity.

```
ICX-Tx-Core1(config)# do ping 192.168.0.1 vrf KA
PING 192.168.0.1 (192.168.0.1) 100(128) bytes of data.
108 bytes from 192.168.0.1: icmp_seq=1 ttl=64 time=0.202 ms
108 bytes from 192.168.0.1: icmp_seq=2 ttl=64 time=0.171 ms
108 bytes from 192.168.0.1: icmp_seq=3 ttl=64 time=0.175 ms
108 bytes from 192.168.0.1: icmp_seq=4 ttl=64 time=0.157 ms
108 bytes from 192.168.0.1: icmp_seq=5 ttl=64 time=0.226 ms

--- 192.168.0.1 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4099ms
rtt min/avg/max/mdev = 0.157/0.186/0.226/0.026 ms
```

```
ICX-Tx-Core1(config)#
```

NOTE: If the ping fails, here is the configuration that is required. You can use this to correct your configuration if needed:

Core1

```
vrf KA
interface 1/1/45
   no shutdown
   vrf attach KA
   ip address 192.168.0.0/31
   exit
```

Core2

```
vrf KA
interface 1/1/45
   no shutdown
   vrf attach KA
   ip address 192.168.0.0/31
   exit
```

NOTE: The /31 subnet mask can be used on the AOS-CX devices on point to point connections that only have 2 hosts. Thanks to the /31 support, there is no waste of the network and broadcast addresses of the subnet.

Configure the VSX roles

Some best practice guidelines:

- On the VSX primary switch, set the system-mac manually. This will ensure that in case this switch needs to be replaced due to hardware failure, the new switch can be configured with the same system-mac as the original switch. By default, the hardware system MAC is used, which would result in a different system MAC address after a hardware change.
- Use the default values for the ISL link timers.
- Make sure the role is meaningful based on the hostname. For example, assign the primary role on core1 and the secondary role on core2.
- Enable the VSX-global sync. This will keep many parts of the switch configuration automatically synchronized between the switches.

Core1

11. On Core1, configure the following. Replace xx with your table # (Table 1 use '01', Table 12 use '12').

```
ICX-Tx-Core1(config)# vsx
ICX-Tx-Core1(config-vsx)# system-mac 02:01:00:00:xx:00
ICX-Tx-Core1(config-vsx)# inter-switch-link lag 256
ICX-Tx-Core1(config-vsx)# role primary
ICX-Tx-Core1(config-vsx)# vsx-sync vsx-global
ICX-Tx-Core1(config-vsx)# exit
```

Core2

12. On Core2, configure the following:

```
ICX-Tx-Core2(config)# vsx
ICX-Tx-Core2(config-vsx)# inter-switch-link lag 256
ICX-Tx-Core2(config-vsx)# role secondary
ICX-Tx-Core2(config-vsx)# exit
```

Core1

13. On Core1, now review the VSX status.

```
ICX-Tx-Core1(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:xx:00   02:01:00:00:xx:00
Platform            8325                8325
Software Version    GL.10.04.0030       GL.10.04.0030
Device Role         primary             secondary

ICX-Tx-Core1(config)#
```

NOTE: In the above output, the switch software version may be different from your actual lab environment. However, both switches must have the **same** version.

14. Compare the running configuration between Core1 and Core2.

Core1

```
ICX-Tx-Core1(config)# show running-config | begin 5 vsx
vsx
    system-mac 02:01:00:00:xx:00
    inter-switch-link lag 256
    role primary
    vsx-sync vsx-global
ip dns server-address 10.254.1.21 vrf mgmt
https-server rest access-mode read-write
```

```
https-server vrf mgmt
ICX-Tx-Core1(config)#
```

Core2

```
ICX-Tx-Core2(config)# show running-config | begin 5 vsx
vsx
    system-mac 02:01:00:00:xx:00
    inter-switch-link lag 256
    role secondary
    vsx-sync vsx-global
ip dns server-address 10.254.1.21 vrf mgmt
https-server rest access-mode read-write
https-server vrf mgmt
ICX-Tx-Core2(config)#
```

Notice that Core2 has received the 'vsx-sync' command automatically, as well as the 'system-mac' command.

Core1

15. On Core1, review the VSX section of the running configuration.

```
ICX-Tx-Core1(config)# show running-config vsx
vsx
    system-mac 02:01:00:00:xx:00
    inter-switch-link lag 256
    role primary
interface lag 256
   no shutdown
   no routing
    vlan trunk native 1 tag
    vlan trunk allowed all
   lacp mode active
interface 1/1/46
   no shutdown
   lag 256
interface 1/1/47
   no shutdown
   lag 256
ICX-Tx-Core1(config)#
```

NOTE: Notice in the above output that the interface that is used for the ISL will automatically set the native VLAN to 'native tagged' state. This means that there is no untagged traffic passing on the ISL.
One of the benefits of having the native VLAN tagged is that any 802.1p CoS value would be preserved on the ISL.

16. Review the running configuration on Core2.

```
ICX-Tx-Core2(config)# show running-config vsx
vsx
```

```
    system-mac 02:01:00:00:xx:00
    inter-switch-link lag 256
    role secondary
interface lag 256
   no shutdown
   no routing
    vlan trunk native 1 tag
    vlan trunk allowed all
   lacp mode active
interface 1/1/47
   no shutdown
   lag 256
interface 1/1/46
   no shutdown
   lag 256
ICX-Tx-Core2(config)#
```

17. Next, review the configuration that is synchronized between the two core switches.

Core1

```
ICX-Tx-Core1(config)# show running-config vsx-sync
Current vsx-sync configuration:
!
!Version ArubaOS-CX GL.10.04.0030
!export-password: default
vsx
    system-mac 02:01:00:00:xx:00
    vsx-sync vsx-global
```

Core2

```
ICX-Tx-Core2(config)# show running-config vsx-sync
Current vsx-sync configuration:
!
!Version ArubaOS-CX GL.10.04.0030
!export-password: default
vsx
    system-mac 02:01:00:00:xx:00
    vsx-sync vsx-global
ICX-Tx-Core2(config)#
```

Enable the Keepalive link

Enable VSX with the keepalive link. This is the detection of split brain, in case all the links of the ISL would be down. The split brain will be tested in a later Task.

Core1

18. Configure the keepalive link, make sure to check the source and destination IP.

```
ICX-Tx-Core1(config)# vsx
```

```
ICX-Tx-Core1(config-vsx)# keepalive peer 192.168.0.1 source 192.168.0.0 vrf KA
ICX-Tx-Core1(config-vsx)# exit
```

Core2

19. Set the keepalive link, make sure to inverse the source and destination IP.

```
ICX-Tx-Core2(config)# vsx
ICX-Tx-Core2(config-vsx)# keepalive peer 192.168.0.0 source 192.168.0.1 vrf KA
ICX-Tx-Core2(config-vsx)# exit
```

Verify the VSX status

Core1

20. On the core1, verify the VSX status.

```
ICX-Tx-Core1(config)# show vsx brief
ISL State                              : In-Sync
Device State                           : Peer-Established
Keepalive State                        : Keepalive-Init
Device Role                            : primary
Number of Multi-chassis LAG interfaces : 0
ICX-Tx-Core1(config)#
```

NOTE: The above output was taken right after enabling the Keep-alive link, showing the Keepalive State as 'Keepalive-init'. A few seconds later, the state will transition to 'Keepalive-Established', so you may see that state already in your output.

Core2

21. Review the same output on the console of Core2.

```
ICX-Tx-Core2(config)# show vsx brief
ISL State                              : In-Sync
Device State                           : Peer-Established
Keepalive State                        : Keepalive-Established
Device Role                            : secondary
Number of Multi-chassis LAG interfaces : 0
ICX-Tx-Core2(config)#
```

Core1

When an administrator is connected to Core1, they may need to review some command output on the VSX peer device (Core2). Instead of having to connect to the peer device, many commands have the 'vsx-peer' command option.

This indicates to the CLI that the command should be executed on the peer device, so the output of the command comes effectively from the peer device, but there is no need to log into the peer device.

22. On Core1, run the same command, but add the 'vsx-peer' option. Compare this output with the previous output that was shown on the console of the Core2 device. The command output should be the same.

```
ICX-Tx-Core1(config)# show vsx brief vsx-peer
ISL State                              : In-Sync
Device State                           : Peer-Established
Keepalive State                        : Keepalive-Established
Device Role                            : secondary
Number of Multi-chassis LAG interfaces : 0
```

23. On Core1, test this feature again with the 'show lldp neighbor 1/1/46' command. When the command is executed using vsx-peer, the output will come from Core2, so Core1 will be listed as an LLDP peer.

```
show lldp neighbor-info 1/1/46
show lldp neighbor-info 1/1/46 vsx-peer
```

24. On Core1, verify that the keepalive link is properly connected.

```
ICX-Tx-Core1(config)# show vsx status keepalive
Keepalive State           : Keepalive-Established
Last Established          : Mon Dec 23 16:15:34 2019
Last Failed               : Mon Dec 23 16:13:27 2019
Peer System Id            : 02:01:00:00:xx:00
Peer Device Role          : secondary

Keepalive Counters
Keepalive Packets Tx      : 385
Keepalive Packets Rx      : 255
Keepalive Timeouts        : 0
Keepalive Packets Dropped : 0
ICX-Tx-Core1(config)#
```

25. On Core1, execute the same command for the VSX peer (Core2). Note the Peer device role and the counters in the 2 command outputs.

```
ICX-Tx-Core1(config)# show vsx status keepalive vsx-peer
Keepalive State           : Keepalive-Established
Last Established          : Mon Dec 23 16:16:16 2019
Last Failed               :
Peer System Id            : 02:01:00:00:xx:00
Peer Device Role          : primary

Keepalive Counters
Keepalive Packets Tx      : 360
Keepalive Packets Rx      : 360
Keepalive Timeouts        : 0
Keepalive Packets Dropped : 0
```

Task 3: VSX Configuration Synchronization

Diagram

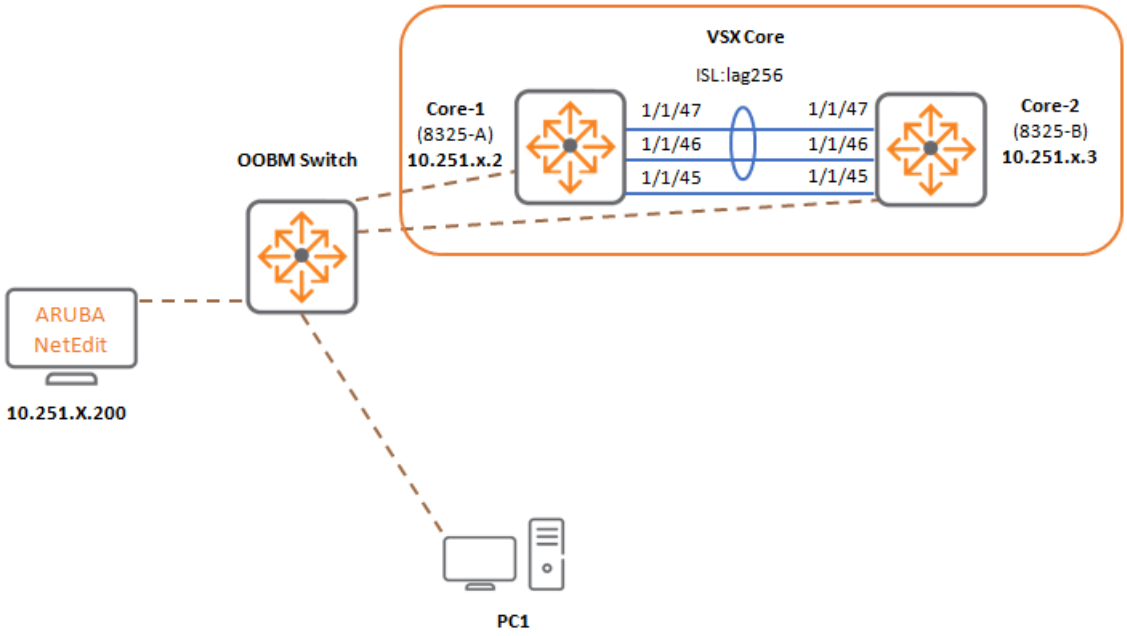

Objectives

- Review default configuration synchronization behavior on both VSX switches
- Enable configuration synchronization
- Verify VSX operation
- Review out of sync status

Steps

Review the default configuration synchronization behavior on both VSX switches

In these steps, the current status of the configuration sync will be reviewed. Only a limited set of configurations ('vsx-global sync') are synchronized between the VSX peers at this point. This introduces potential configuration risks, since an administrator should remember to define a VLAN (or most other settings) on **both** VSX peers. In the next steps, this default behavior will be demonstrated. Afterwards, the VSX configuration synchronization feature will be used to automatically synchronize the configuration between the VSX peers.

In the configuration, there are:

- **Global** configuration settings: global ACL, time configuration etc.

  Synchronization of these settings can be enabled under the VSX global context

- **Context** configuration settings: VLANs

  Synchronization of these settings must be enabled per context.

The lab will introduce a global and a context configuration change.

Core1

1. On the Core1, verify that the config sync process is active.

```
ICX-Tx-Core1(config)# show vsx status config-sync
Admin state             : Enabled
Operational State       : Operational
Error State             : None
Recommended remediation : N/A
Current time            : Mon Dec 23 16:28:42 2019
Last sync time          : Sun Dec 22 21:21:10 2019
ICX-Tx-Core1(config)#
```

Global configuration change

As an example, a new admin user called 'demo' will be defined on the Core1. Since there is no configuration synchronization at this point, the user will not exist on Core2.

The admin users are part of the 'aaa' vsx-sync feature, that will be enabled in the next section.

2. On Core1, define a new admin user.

```
ICX-Tx-Core1(config)# user demo group administrators password plaintext demo
```

3. Review the running configuration.

```
ICX-Tx-Core1(config)# show running-config | include user
user admin group administrators password ciphertext
AQBapV8rLxVdvc31eF2hwggG6NdJ/p/bqXrRlSsdibXBcJ5SYgAAAI9vpWqHUNb5frqUOqkMdrazRAVY0
tNdl36lkRQVKggG5ldaN9iEzCWFKQz1b7fMpNMt8xX0nal2et4T8N1aXuXnvJFKH+X1BPs7BSiKboKXGi
EW4/sy/Sv6EGz5yshSNuoB
user demo group administrators password ciphertext
AQBapazfG0p4r497EJ3fcdfaSHIeRNhPeHHvMVYX9r8YpjulYgAAAHGn29l4MqabnQyJ0tQpanbi+eao6
jzb/tfgRDXefU5s+4U5qDInY5fTjoGUNSxQjTHNeUpgWQ9cbCCXxirTU9eKfbWbU3jS45B9GGwlM5nySf
AopiFDLebNg+S3yEW8o2fy
ICX-Tx-Core1(config)#
```

4. Review the running configuration of the vsx-peer (**Core2**), the user demo should not exist.

```
ICX-Tx-Core1(config)# show running-config vsx-peer | include user
user admin group administrators password ciphertext
ICX-Tx-Core1(config)#
```

This was an example of a global configuration change.

Context configuration change

Next you will define a new VLAN on Core1 and verify the status on Core2.

Core1

5. On Core1, define VLAN 11.

```
ICX-Tx-Core1(config)# vlan 11
ICX-Tx-Core1(config-vlan-11)# exit
```

6. Verify the VLAN list.

```
ICX-Tx-Core1(config)# show vlan

--------------------------------------------------------------------------------
VLAN  Name                                Status  Reason  Type      Interfaces
--------------------------------------------------------------------------------
1     DEFAULT_VLAN_1                      up      ok      default   lag256
11    VLAN11                              up      ok      static    lag256
```

7. Next, verify the VLAN list on the vsx-peer Core2.

```
ICX-Tx-Core1(config)# show vlan vsx-peer

--------------------------------------------------------------------------------
VLAN  Name                                Status  Reason  Type      Interfaces
--------------------------------------------------------------------------------
1     DEFAULT_VLAN_1                      up      ok      default   lag256
```

This demonstrates that each switch has a local configuration, and, by default, a configuration change on Core1 is not automatically pushed to Core2.

Enable Global configuration synchronization for other features.

The administrator must 'opt-in' for the global features that should be synchronized by the VSX members.

8. On the Core1, review the current VSX configuration.

```
ICX-Tx-Core1(config)# show running-config vsx-sync
Current vsx-sync configuration:
!
!Version ArubaOS-CX GL.10.04.0030
!export-password: default
vsx
    system-mac 02:01:00:00:xx:00
```

```
    vsx-sync vsx-global
```

9. Next enable some VSX sync features. For this lab exercise, the 'aaa' option is required, since that includes the admin users.

```
ICX-Tx-Core1(config)# vsx
ICX-Tx-Core1(config-vsx)# vsx-sync aaa bfd-global bgp dhcp-relay mclag-interfaces
ICX-Tx-Core1(config-vsx)# vsx-sync ospf qos-global route-map
ICX-Tx-Core1(config-vsx)# exit
```

Verify synchronization operation for Global features

10. Review the updated output of the VSX sync command.

```
ICX-Tx-Core1(config)# show run vsx-sync
Current vsx-sync configuration:
!
!Version ArubaOS-CX GL.10.04.0030
!export-password: default
user admin group administrators password ciphertext
AQBapXRaY6Hk0KGjVzqZjewoYg94PjoFJCPdL8qqusr+/7b/YgAAADi/osVJXNfAMvTortdokcXfSSG1/
9RzFDtehOzaNeea8iMUWOtNNv1SVTufQwLeVHXnk2Crpat4r29ENMWQJ76/Z1cngFVibQV4M7bELYq8+G
JCgZfKHoTgjlNOKjqCb5EO
user demo group administrators password ciphertext
AQBapWf0qG1No/HM5fQb3+D0teB4f9bvCrhVtqExP0xflB1wYgAAAGmR1x0axs0A5LVMSvTMmJAllPr49
4ilg4ZcSzOCZGXVRDvfXlIgo9+kVA9rAWc42ONmXLElQ4+0wF3vGiW1Tj1I8ht4VUrCF4iKIhf+BVc+wQ
l/lcAcgN+qdHDQZ006CFvb
!
!
!
vsx
    system-mac 02:01:00:00:xx:00
    vsx-sync aaa bfd-global bgp dhcp-relay mclag-interfaces ospf qos-global
route-map vsx-global
```

11. Run the command again with the 'vsx-peer' option. The config of the local Core1 and remote Core2 should be the same now, so the user 'demo' exists in both switch configurations.

```
ICX-Tx-Core1(config)# show running-config vsx-sync vsx-peer
Current vsx-sync configuration:
!
!Version ArubaOS-CX GL.10.04.0030
!export-password: default
user admin group administrators password ciphertext
user demo group administrators password ciphertext
!
!
```

```
!
vsx
    system-mac 02:01:00:00:xx:00
    vsx-sync aaa bfd-global bgp dhcp-relay mclag-interfaces ospf qos-global
route-map vsx-global
```

12. Cleanup: On Core1, remove the ***demo*** user account.

```
ICX-Tx-Core1(config)# no user demo
User demo's home directory and active sessions will be deleted.
Do you want to continue (y/n)? y

ICX-Tx-Core1(config)#
```

This demonstrates global feature synchronization.

Synchronization for context features

For a context feature, such as a VLAN, the synchronization must be enabled within the context.

13. On the Core1, enter the VLAN 11 context and enable the VSX sync option for this VLAN.

```
ICX-Tx-Core1(config)# vlan 11
ICX-Tx-Core1(config-vlan-11)# vsx-sync
ICX-Tx-Core1(config-vlan-11)# exit
ICX-Tx-Core1(config)#
```

14. Review the VLAN list of the VSX Peer device.

```
ICX-Tx-Core1(config)# show vlan vsx-peer

--------------------------------------------------------------------------------
VLAN  Name                                Status  Reason Type      Interfaces
--------------------------------------------------------------------------------
1     DEFAULT_VLAN_1                      up      ok     default   lag256
11    VLAN11                              up      ok     static    lag256
ICX-Tx-Core1(config)#
```

15. Enable VLAN 12 and VLAN 13, these will be used in later labs.

```
ICX-Tx-Core1(config)# vlan 12,13
ICX-Tx-Core1(config-vlan-<12,13>)# vsx-sync
ICX-Tx-Core1(config-vlan-<12,13>)# exit
```

16. Verify the VLANs exist on the VSX peer.

```
ICX-Tx-Core1(config)# show vlan vsx-peer

--------------------------------------------------------------------------------
```

```
VLAN  Name                                Status  Reason   Type      Interfaces
------------------------------------------------------------------------------
1     DEFAULT_VLAN_1                      up      ok       default   lag256
11    VLAN11                              up      ok       static    lag256
12    VLAN12                              up      ok       static    lag256
13    VLAN13                              up      ok       static    lag256
```

Review the sync process rules

At this point, the lab has demonstrated that the VSX configuration synchronization feature can help the administrator to perform an action only once on the primary switch, and the synchronization will handle the configuration change on the secondary switch.

There are some rules for the synchronization process:

- For any feature that IS ENABLED for synchronization:
  - Configuration changes MUST be made on the primary switch.
  - If any change is made on the secondary switch, the change will be overwritten by the VSX sync with the primary switch configuration. Therefore, the secondary switch SHOULD NOT be used for those configuration changes.
- For any feature that IS NOT ENABLED for synchronization:
  - Changes can be made on either switch and will only exist in the local configuration of each switch.
- Synchronization status changes:
  - When synchronization is turned off for a feature, each switch maintains the current configuration for that feature in the local configuration. Disabling the configuration sync will not remove that configuration on the secondary node.
  - When both switches have local configuration changes for a feature, and the feature is enabled for synchronization, the version of the primary switch will be leading. The secondary switch will lose any local configuration changes related to this feature.
  - If the sync process is 'out of sync', due to an ISL failure or a member reboot, the system will automatically update all the settings for the enabled features the first time they come back online. Since this is a complete feature sync, any new, modified or removed configurations are automatically synchronized as well. Any changes applied locally on the secondary node would be lost at this point.

Task 4: VSX Layer-2 - VSX Link Aggregation (VSX LAG)

Diagram

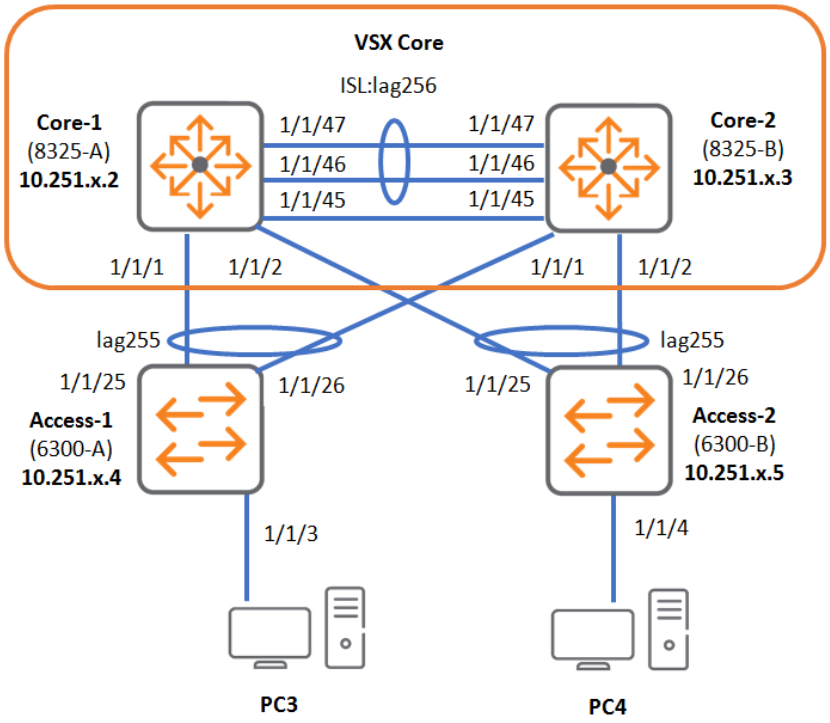

Objectives

- Setup a VSX LAG with config sync
- Configure peer switches (Access1 and Access2)
- Verify connectivity
- Configure a VSX LAG to the Aruba Mobility Controller

This task will cover the redundant Layer-2 connectivity to a VSX system. Since both VSX members have their own management and control planes, the default behavior would use each system's own unique MAC address for Layer-2 protocols.

Since VSX wants to present itself as a single Layer-2 device to the peer devices, the VSX LAG feature can be used so both VSX members will present themselves with the same MAC address for Layer-2 protocols such as LACP and STP.

A VSX member can still have a local LAG, using only local ports.

Any LAG that spans ports of both members of the VSX system, the administrator must define this LAG as a VSX LAG.

In this task, two VSX LAGs will be defined: one to Access1 and one to Access2.

Steps

Core1: Link-Aggregation to Access1

1. Open a console connection to Core1.
2. Enter configuration mode and define a new LAG of type multi-chassis. This is the VSX LAG.

```
ICX-Tx-Core1(config)# interface lag 1 multi-chassis
```

3. By default, an VSX LAG is:

   - Shutdown

   - LACP enabled

   - Switched port (a VSX LAG cannot be a routed port)

NOTE: LACP is the recommended method to connect with the peer switches. When it is not possible to use LACP with the peer device, a static (no protocol) VSX LAG can be created using 'interface lag <id> multi-chassis static'.

NOTE: In some scenarios, the peer device may not always support LACP. This may happen with servers during the boot process (when the server admin wants to use a PXE boot for the initial server provisioning). This may also occur when the peer switch must be deployed with ZTP (zero touch provisioning), since the peer switch does not have an LACP configuration at factory default.

For these use cases, it is possible to configure AOS-CX with the LACP fall-back feature. With this feature enabled, the switch will keep one interface in a forwarding state when no LACP packets are received from the peer device. In the default LACP mode, both interfaces are blocked, so the peer device is not be able to reach the network.

Review this default state by reviewing the current configuration for the current context.

```
ICX-Tx-Core1(config-lag-if)# show run current-context
interface lag 1 multi-chassis
    no routing
    vlan access 1
    lacp mode active
ICX-Tx-Core1(config-lag-if)#
```

4. Make the VSX LAG a VLAN trunk and permit VLANs 1 and 11-13.

```
ICX-Tx-Core1(config-lag-if)# vlan trunk allowed 1,11-13
```

NOTE: In the above command, 11-13 is a range, so VLAN 12 will be included as

well.

5. Set a description on the LAG.

```
ICX-Tx-Core1(config-lag-if)# description access1
```

6. Enable the port.

```
ICX-Tx-Core1(config-lag-if)# no shutdown
```

7. Review the configuration changes.

```
ICX-Tx-Core1(config-lag-if)# show run current-context
interface lag 1 multi-chassis
    description access1
    no shutdown
    no routing
    vlan trunk native 1
    vlan trunk allowed 1,11-13
    lacp mode active
ICX-Tx-Core1(config-lag-if)# exit
```

Core2: Link-Aggregation to Access1

8. Open a console connection to Core2.
9. Define a new VSX LAG with LAG id 1 of type multi-chassis.

```
ICX-Tx-Core2(config)# interface lag 1 multi-chassis
```

10. Review the current configuration.

```
ICX-Tx-Core2(config-lag-if)# show run current-context
interface lag 1 multi-chassis
    description access1
    no routing
    vlan trunk native 1
    vlan trunk allowed 1,11-13
    lacp mode active
ICX-Tx-Core2(config-lag-if)#
```

Q: Why is the VSX LAG already a VLAN trunk and why are the VLANs allowed?

A: Due to the global vsx config-sync 'mclag-interfaces'

11. Even though the VLAN configuration has been synchronized, the interface must still be enabled per member. Enable the interface.

```
ICX-Tx-Core2(config-lag-if)# no shutdown
```

12. Review the configuration.

```
ICX-Tx-Core2(config-lag-if)# show run current-context
interface lag 1 multi-chassis
    no shutdown
    description access1
    no routing
    vlan trunk native 1
    vlan trunk allowed 1,11-13
    lacp mode active
ICX-Tx-Core2(config-lag-if)# exit
```

Core1: Assign local physical port to the VSX LAG

13. On the Core1, assign interface 1/1/1 to the LAG 1: this is the LAG that was created in the previous steps.

Description	Access1
MTU	9100
LAG	1
Status	enabled (no shutdown)

```
ICX-Tx-Core1(config)# int 1/1/1
ICX-Tx-Core1(config-if)# description access1
ICX-Tx-Core1(config-if)# lag 1
ICX-Tx-Core1(config-if)# mtu 9100
ICX-Tx-Core1(config-if)# no shutdown
ICX-Tx-Core1(config-if)# exit
```

Core2

14. On the Core2, repeat the commands of the previous step for interface 1/1/1.

```
ICX-Tx-Core2(config)# interface 1/1/1
ICX-Tx-Core2(config-if)# description access1
ICX-Tx-Core2(config-if)# lag 1
ICX-Tx-Core2(config-if)# mtu 9100
ICX-Tx-Core2(config-if)# no shutdown
ICX-Tx-Core2(config-if)# exit
```

Configure the Link-Aggregation to Access2 using NetEdit

In these steps, the VSX LAG to Access2 will be defined. These steps will be done using NetEdit. Since many steps are the same on both Core switches, NetEdit is very convenient to perform these changes.

15. **Optional:** Below are the manual steps for both Core1 and Core2 in case you do not want to use NetEdit:

```
interface lag 2 multi-chassis
    description access2
    no shutdown
```

```
    no routing
    vlan trunk native 1
    vlan trunk allowed 1,11-13
    lacp mode active

interface 1/1/2
    description access2
    no shutdown
    mtu 9100
    lag 2
```

16. Access PC1 and connect to NetEdit. Login using **admin**/**aruba123**.
17. Navigate to **Devices**. Select both **Core1** and **Core2**. From the **Action** menu select **Edit Config**.
18. Enter **'lag2'** as the plan name, click **Create.**
19. Define a new lag interface after the previously defined lag 1 configuration.

NOTE: Remember to use a <SPACE> to enter the interface context after entering the first line!

```
interface lag 2 multi-chassis
    description access2
    no shutdown
    no routing
    vlan trunk native 1
    vlan trunk allowed 1,11-13
    lacp mode active
```

20. This should be the result for the LAG.

```
20 system interface-group 1 speed 10g
21 interface lag 1 multi-chassis
22     description access1
23     no shutdown
24     no routing
25     vlan trunk native 1
26     vlan trunk allowed 1,11-13
27     lacp mode active
28 interface lag 2 multi-chassis
29     description access2
30     no shutdown
31     no routing
32     vlan trunk native 1
33     vlan trunk allowed 1,11-13
34     lacp mode active
35
```

21. Next define the interface 1/1/2 context, link it to the lag2.

```
interface 1/1/2
    description access2
    no shutdown
    mtu 9100
    lag 2
```

22. This should be the result.

```
41 interface 1/1/1
42     description access1
43     no shutdown
44     mtu 9100
45     lag 1
46 interface 1/1/2
47     description access2
48     no shutdown
49     mtu 9100
50     lag 2
51
```

23. Click **Return to Plan**. Click **Deploy** and confirm by clicking **Deploy.**
24. Once the deployment has completed, use **Commit** to save the configuration and confirm with **Commit**.

Access1 and Access2 LAG to the VSX Core using NetEdit

Both Access1 and Access2 need to get the uplink LAG defined. Since both access switches will be using the same uplink ports, NetEdit will be used to push the same configuration to both Access switches.

Both access switches will be using the same LAG ID as the upstream LAG to connect to the Core switches. This make it easy to identify the LAG. In these labs, the LAG ID 255 will be used for this LAG.

NOTE: The lab instructions are based on NetEdit, if you cannot use NetEdit or need to push the commands directly to Access1 and Access2, these are the commands:

```
vlan 11-13
interface lag 255
   description core
   no routing
   vlan trunk allowed all
   lacp mode active
   no shutdown
   exit

interface 1/1/25,1/1/26
   lag 255
   mtu 9100
   no shutdown
   exit
```

The allowed VLAN list is controlled on the Core side of the link, to simplify future labs, the Access side of the VLAN trunk is set to 'allow all vlans'.

25. Use the PC1 to connect to NetEdit. Login using **admin**/**aruba123**.
26. Navigate to **Devices**. Select both **Access1** and **Access2**. From the **Action** menu select **Edit Config**.
27. Enter **'lag255 to core'** as the plan name, click **Create.**
28. Define the VLANs 11-13.

```
vlan 11-13
```

29. The result in NetEdit:

```
 7 ssh server vrf mgmt
 8 vlan 1
 9 vlan 11-13
10 spanning-tree
```

30. Define LAG 255.

```
interface lag 255
   description core
   no shutdown
   no routing
   vlan trunk native 1
   vlan trunk allowed all
   lacp mode active
```

31. The result in NetEdit.

```
15 vsf member 1
16     type jl668a
17 interface lag 255
18     description core
19     no shutdown
20     no routing
21     vlan trunk native 1
22     vlan trunk allowed all
23     lacp mode active
24 interface 1/1/1
25     shutdown
```

32. Map the uplink interfaces 1/1/25 and 1/1/26 to lag255, set the MTU and enable the ports. Make sure the 'shutdown' command does not exist on a line by itself.

```
interface 1/1/25
   no shutdown
   lag 255
   mtu 9100
```

```
interface 1/1/26
   no shutdown
   lag 255
   mtu 9100
```

33. The result in NetEdit:

34. Click **Return to Plan**, click **Deploy** and confirm by clicking **Deploy.**
35. Wait for the deployment to complete, next click **Commit** and confirm by clicking **Commit.**

Optional Steps: Verification of the VSX LAG using the Switch CLI

These steps are optional. The next section will perform the same steps using NetEdit.

Now that the VSX LAG configuration has been done between the VSX core and the 2 Access switches, several configuration and status verification commands can be used.

The next section will use:

- LLDP neighbors to verify the physical link and neighbor switches.
- Access switch LACP status to verify the local and neighbor LAG link states.
- Core switch VSX LAG status to verify the local and neighbor LAG link states.

36. **Optional:** Review the LLDP neighbors on Core1, Core2, Access1 and Access2.

```
show lldp neighbor
```

37. **Optional:** On Access1, review the LAG status. Repeat this on Access2.

```
ICX-Tx-Access1(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  lag255     26    1     ALFNCD  88:3a:30:98:30:c0 65534  256  up
1/1/26  lag255     27    1     ALFNCD  88:3a:30:98:30:c0 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  lag255     1     1     ALFNCD  02:01:00:00:01:00 65534  1
1/1/26  lag255     1001  1     ALFNCD  02:01:00:00:01:00 65534  1
ICX-Tx-Access1(config)#
```

38. **Optional:** On Core1, review the LAG status.

```
ICX-Tx-Core1(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     ALFNCD  02:01:00:00:01:00 65534  1    up
1/1/2   lag2(mc)   2     1     ALFNCD  02:01:00:00:01:00 65534  2    up
1/1/46  lag256     47    1     ALFNCD  90:20:c2:bc:17:00 65534  256  up
1/1/47  lag256     48    1     ALFNCD  90:20:c2:bc:17:00 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)   26    1     ALFNCD  88:3a:30:98:30:c0 65534  256
1/1/2   lag2(mc)   26    1     ALFNCD  88:3a:30:97:b6:00 65534  256
1/1/46  lag256     47    1     ALFNCD  90:20:c2:bc:97:00 65534  256
1/1/47  lag256     48    1     ALFNCD  90:20:c2:bc:97:00 65534  256
ICX-Tx-Core1(config)#
```

Q: Why is the actor system-id for interfaces 1/1/1 and 1/1/46 different?

A: 1/1/1 is part of a VSX LAG, so it uses the VSX system mac. 1/1/46 is part of a local LAG, so it can use the base system MAC instead of the VSX configured system mac.

39. **Optional:** On Core1, review the LAG status with the peer VSX status.

```
ICX-Tx-Core1(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        ALFNCD  02:01:00:00:01:00 65534    1
1/1/2   lag2(mc)   2       1        ALFNCD  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)   26      1        ALFNCD  88:3a:30:98:30:c0 65534    256
1/1/2   lag2(mc)   26      1        ALFNCD  88:3a:30:97:b6:00 65534    256


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        ALFNCD  02:01:00:00:01:00 65534    1
1/1/2  lag2(mc)   1002    1        ALFNCD  02:01:00:00:01:00 65534    2


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)   27      1        ALFNCD  88:3a:30:98:30:c0 65534    256
1/1/2  lag2(mc)   27      1        ALFNCD  88:3a:30:97:b6:00 65534    256
ICX-Tx-Core1(config)#
```

End of optional CLI verification steps.

Verification of the VSX LAG using NetEdit

These steps can only be used if the deployment was pushed using NetEdit.

40. Using the PC1, access NetEdit.
41. Navigate to **Plans**, open the latest plan, it was named **'lag255 to core'** if you followed the naming suggestion of the lab guide. This was the plan that was used to push the LAG255 to the Access1 and Access2 switches.

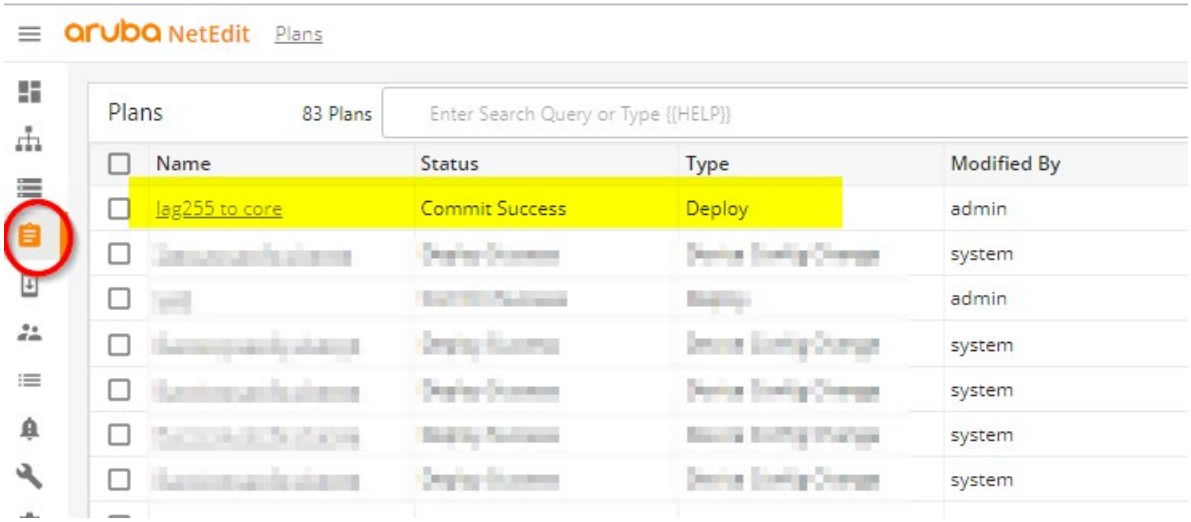

42. Click **Change Validation**.

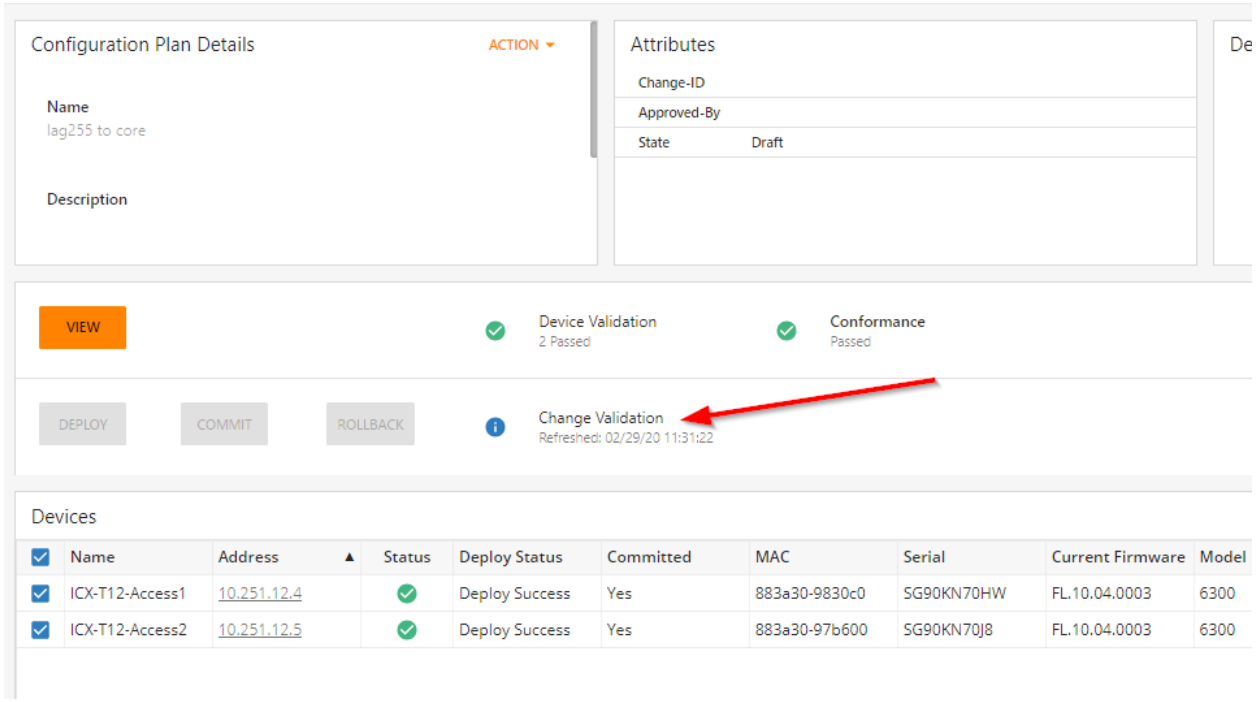

43. Expand the **ICX-TX-Access1 > show lacp interfaces'** validation test.

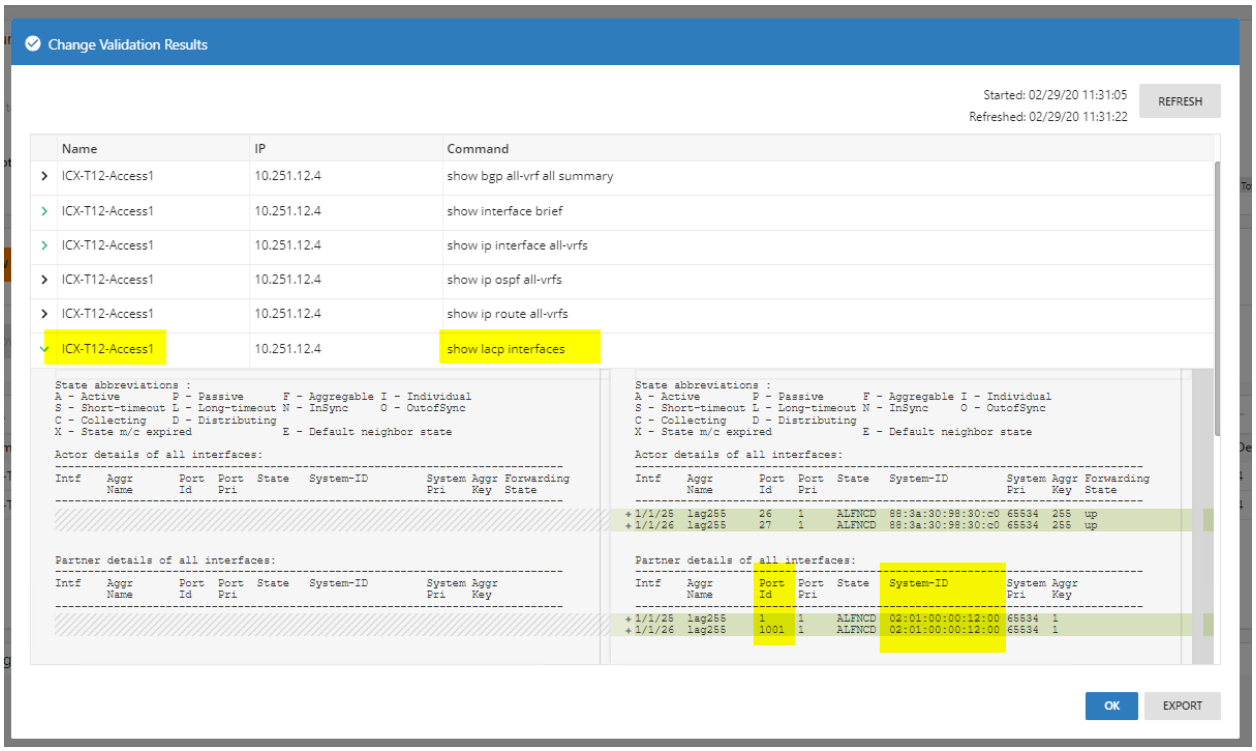

Q: What is the Forwarding State of the ports in the LAG255?

A: Both ports 1/1/25 and 1/1/26 should have state 'up'

Q: Under the 'Partner details', what is the Partner System-Id?

A: This is the LACP system ID, which is based on the Switch system ID. This was configured statically to 02:01:00:00:xx:00.

Q: Is the Partner System-Id the same for the ports 1/1/25 and 1/1/26?

A: Yes, VSX ensures that both members of the VSX are communicating with the same System-Id to peer devices, so they appear as 1 device.

44. Expand 'ICX-Tx-Access1' > 'show lldp neighbor-info' to see the validation details.

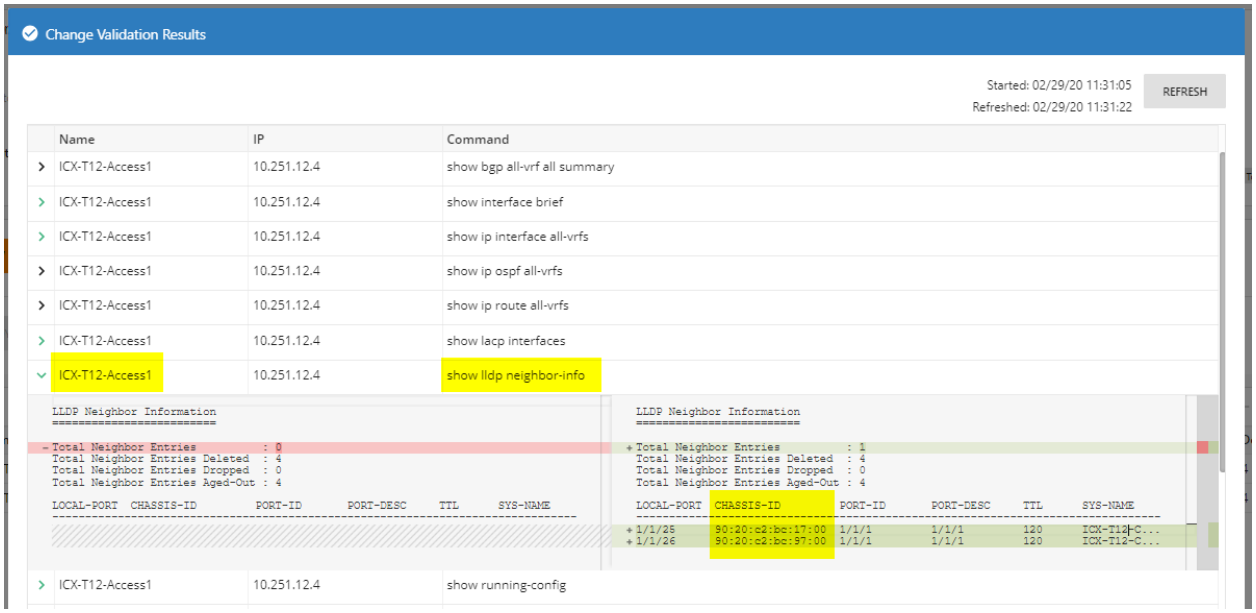

Q: Why is the Chassis-Id different from the LACP System-id?

A: The VSX core consists of 2 independent control planes, with their own IP and MAC address. They only share a common System-Id for Layer-2 protocols such as LACP to make their partners believe they are 1 system. LLDP is a discovery protocol and does not require the use of the common System-Id, so each VSX switch uses its own System MAC address for LLDP.

45. Scroll down to find the 'ICX-Tx-Access2' > 'show lacp interfaces' and verify the LAG255 is up for both interfaces 1/1/25 and 1/1/26.

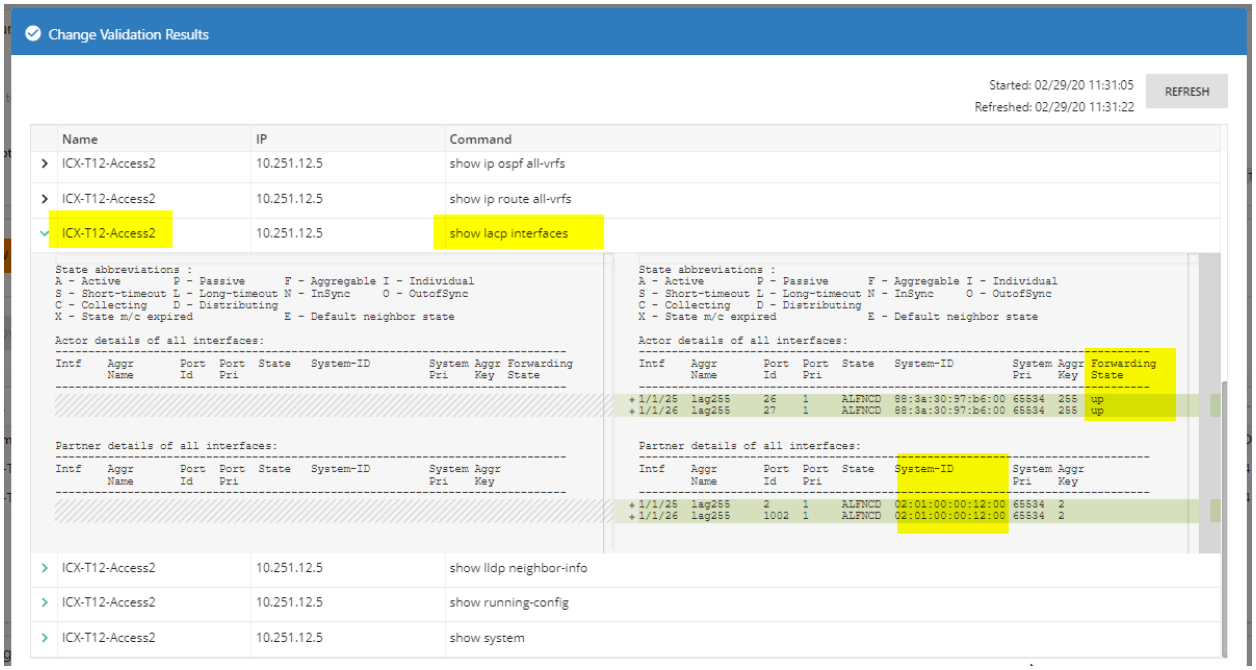
Change Validation Results
Started: 02/29/20 11:31:05
Refreshed: 02/29/20 11:31:22
REFRESH
Name
IP
Command
ICX-T12-Access2
10.251.12.5
show ip ospf all-vrfs
ICX-T12-Access2
10.251.12.5
show ip route all-vrfs
ICX-T12-Access2
10.251.12.5
show lacp interfaces
State abbreviations :
A - Active P - Passive 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 Name Port Id Port Pri State System-ID System Pri Aggr Key Forwarding State
+ 1/1/25 lag255 26 1 ALFNCD 88:3a:30:97:b6:00 65534 255 up
+ 1/1/26 lag255 27 1 ALFNCD 88:3a:30:97:b6:00 65534 255 up
Partner details of all interfaces:
Intf Aggr Name Port Id Port Pri State System-ID System Pri Aggr Key
+ 1/1/25 lag255 2 1 ALFNCD 02:01:00:00:12:00 65534 2
+ 1/1/26 lag255 1002 1 ALFNCD 02:01:00:00:12:00 65534 2
ICX-T12-Access2
10.251.12.5
show lldp neighbor-info
ICX-T12-Access2
10.251.12.5
show running-config
ICX-T12-Access2
10.251.12.5
show system

VSX LAG to Aruba Mobility Controller > NetEdit

Diagram

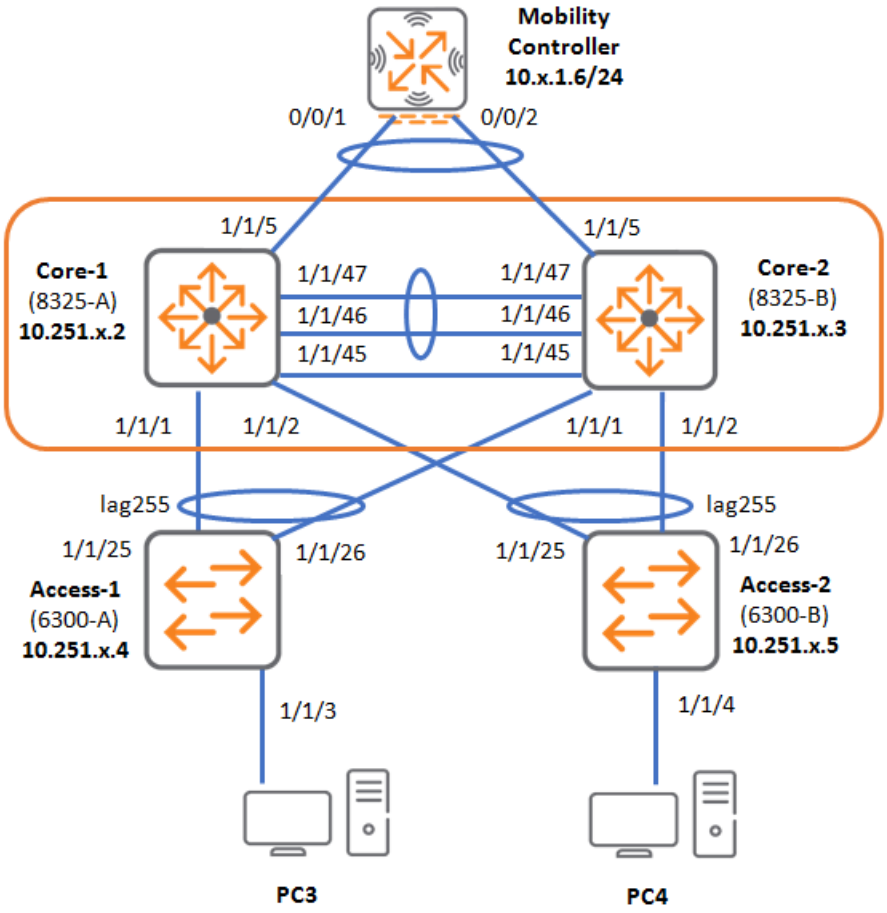

The Aruba Mobility Controller has been pre-configured in the lab with a LAG on its ports to the Core switches.

As an exercise, try configuring the VSX LAG to the Mobility Controller on your own.

You may use either NetEdit (use plan name **'vsx lag mc'**) or the CLI for this exercise.

Use these settings for the VSX LAG:

- Interface 1/1/5 on both Core1 and Core2
- MTU of 9100 (to support GRE tunnels in the dynamic segmentation lab)
- VLAN trunk, allow all VLANs

In case you are unsure, these are the commands:

46. CLI commands to be applied to **<u>both</u> Core1** and **Core2**, either in terminal or NetEdit:

```
configure terminal
interface lag 5 multi-chassis
    vlan trunk allowed all
    no shutdown

interface 1/1/5
    mtu 9100
    lag 5
    no shutdown
```

47. Verify connectivity, LACP state is UP for 1/1/5 and check VLANs are allowed on the Core1 and its vsx-peer.

```
show lacp interfaces
show lacp interfaces vsx-peer
```

> **NOTE:** Make sure to **'Commit'** the configuration when using NetEdit, or use 'write memory' when using the terminal session.

48. For reference only: example validation result for the LACP interfaces on Core1 when using NetEdit.

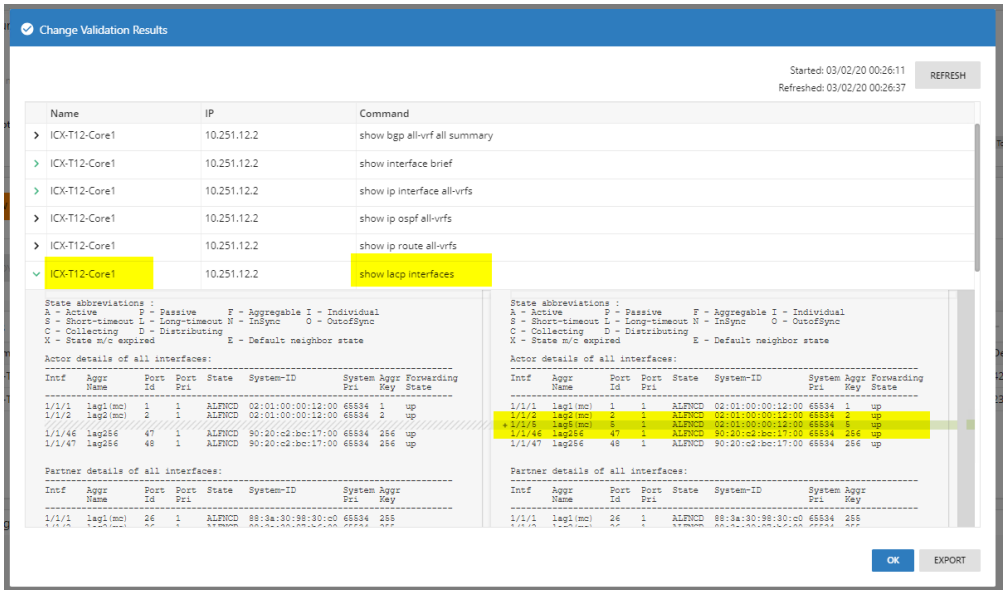

49. For reference only: example validation result for the LACP interfaces on Core2 when using NetEdit.

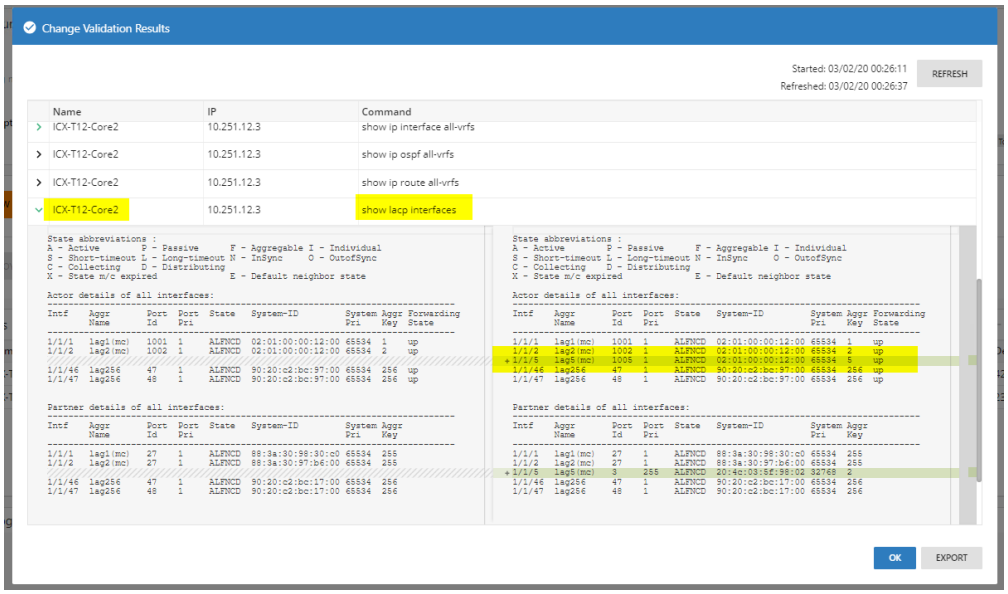

Task 5: VSX Layer3 - Active Gateway

Diagram

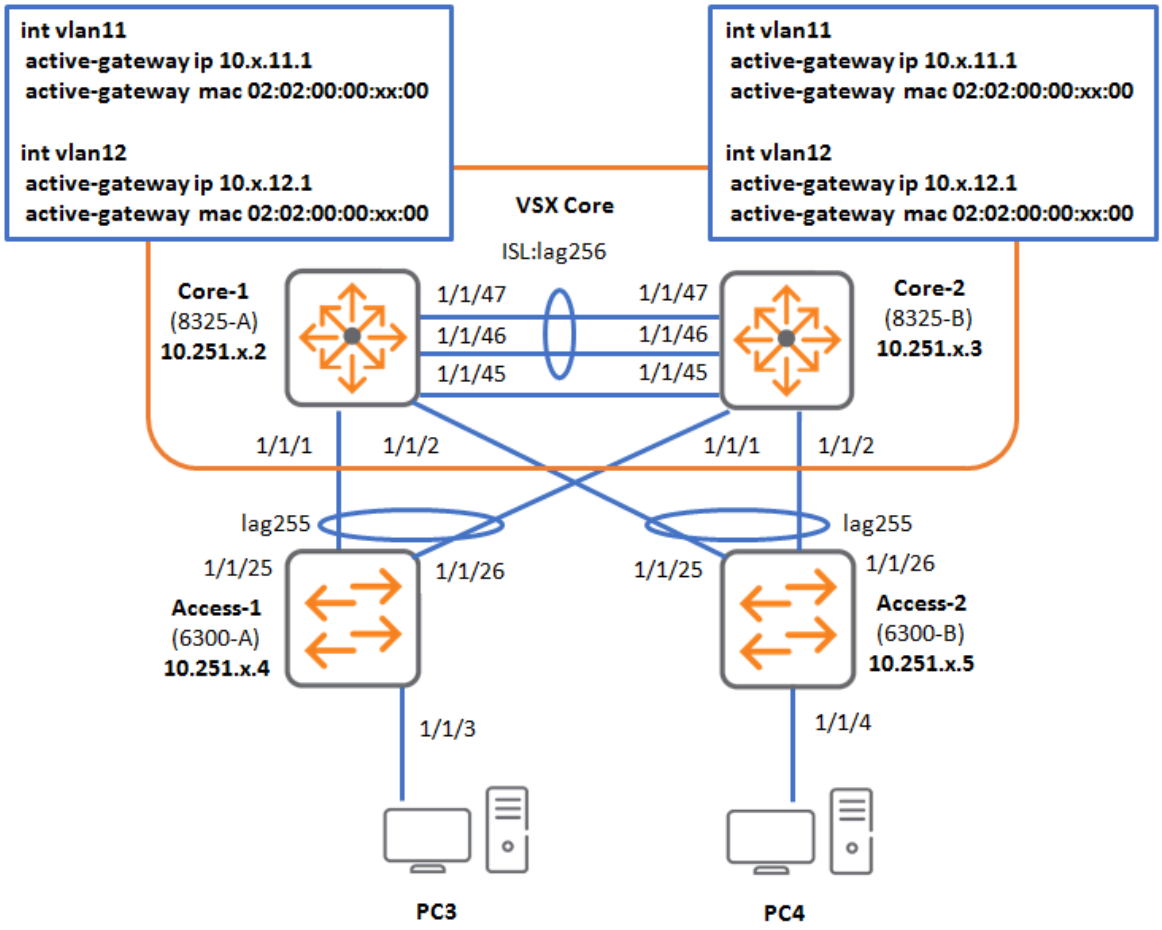

Objectives

- Understand the Active Gateway principle
- Configuration of the Active Gateway
- Verify the Active Gateway configuration
- Verify the Active Gateway operation on a client
- Setup Active Gateway in VLAN 1 to the Mobility Controller
- Configure IP helper on the VSX Core

Understand the Active Gateway principle

In a VSX system, active gateway provides redundant default gateway functionality for the end-hosts. The default gateway of the end-host is automatically handled by both of the VSX systems.

This functionality is similar to VRRP, but VRRP only operates in an active/standby mode, using a keep-alive between the active and the standby system to detect the state. To keep the ARP entry of the default gateway valid in case of a VRRP failover, VRRP uses a common MAC address between the VRRP hosts. In the case of VRRP, this is based on the VRRP VRID.

The active gateway feature also uses a virtual MAC address to ensure the ARP entry on the end-hosts is stable. However, the virtual MAC is set by the administrator, it is not controlled by some other ID. The virtual MAC can be re-used over multiple VLANs, since the MAC address is only used within the VLAN.

In the VSX setup, the virtual MAC and active gateway IP is programmed in the hardware tables of both VSX switches, so this results in an active/active setup.

Since the peer switches are typically connected with a Multi-Chassis Link Aggregation (VSX LAG), the peer switch load-distribution decides if traffic is sent to VSX member 1 or member 2. Whatever switch receives the traffic first, will be the switch that handles the traffic. Therefore, it does not matter whether VSX member 1 or member 2 receives the traffic first, either of them can route the traffic.

In case of a link-failure, the peer switch will switch the traffic over the other link, so the other VSX member will route all the traffic.

Each VSX member still has a local IP address, next to the active gateway IP address. The same active gateway IP address and active gateway MAC address must be configured on both VSX members.

Steps

Configuration

For each VLAN connected to end-hosts, the administrator must set an active gateway IP and virtual MAC address on both VSX members.

This will be the IP address that must be set on the end-hosts as the default gateway IP.

Core1: VLAN11

1. Open a console connection to Core1, enter the configuration mode.
2. Enter the interface VLAN 11 context, configure the local switch IP address. Replace x in the IP address with your table number (e.g. For table 1, use '1', for table 12 use '12'). Enable L3 counters, this will provide IP statistics on the interface.

```
ICX-Tx-Core1(config)# interface vlan 11
ICX-Tx-Core1(config-if-vlan)# ip address 10.x.11.2/24
ICX-Tx-Core1(config-if-vlan)# l3-counters
```

3. Enable the VSX-sync feature for this L3 VLAN interface.

```
ICX-Tx-Core1(config-if-vlan)# vsx-sync active-gateways policies
```

4. Configure the active gateway IP and MAC address. Replace x in the IP address and the MAC address with your table number.

For example, table 1 should use '01', for table 12 use '12'.

```
ICX-Tx-Core1(config-if-vlan)# active-gateway ip 10.x.11.1 mac 02:02:00:00:xx:00
ICX-Tx-Core1(config-if-vlan)# exit
```

5. Verify the configuration. Output example:

```
ICX-Tx-Core1(config)# show interface vlan11

Interface vlan11 is up
 Admin state is up
 Link transitions: 0
 Description:
 Hardware: Ethernet, MAC Address: 90:20:c2:bc:17:00
 IPv4 address 10.x.11.2/24
    active-gateway ip mac 02:02:00:00:xx:00
    active-gateway ip 10.x.11.1
 Rx
       L3:
            0 packets, 0 bytes
 Tx
       L3:
            0 packets, 0 bytes

ICX-Tx-Core1(config)#
```

Core2: VLAN11

6. Open a console connection to Core2, enter the configuration mode.
7. Enter the interface VLAN 11 context, configure the Core2 switch IP address, enable L3 counters.

```
ICX-Tx-Core2(config)# interface vlan 11
ICX-Tx-Core2(config-if-vlan)# ip address 10.x.11.3/24
ICX-Tx-Core2(config-if-vlan)# l3-counters
```

8. Review the current configuration of the VLAN 11 interface. Verify that VSX-sync has completed the active gateway configuration synchronization on Core2.

```
ICX-Tx-Core2(config-if-vlan)# show run current-context
interface vlan11
    vsx-sync active-gateways policies
    ip address 10.12.11.3/24
    active-gateway ip mac 02:02:00:00:12:00
    active-gateway ip 10.12.11.1
    l3-counters
```

9. Verify the operational status. Example output from table 12:

```
ICX-Tx-Core2(config-if-vlan)# show interface vlan11

Interface vlan11 is up
 Admin state is up
 Link transitions: 0
 Description:
 Hardware: Ethernet, MAC Address: 90:20:c2:bc:97:00
 IPv4 address 10.x.11.3/24
    active-gateway ip mac 02:02:00:00:xx:00
    active-gateway ip 10.x.11.1
 Rx
       L3:
            0 packets, 0 bytes
 Tx
       L3:
            0 packets, 0 bytes

ICX-Tx-Core2(config-if-vlan)# exit
```

VLAN 12 Active Gateway

Core1

10. Configure VLAN12 on Core1. Replace x in the IP address and the MAC address with your table number.

For example, table 1 should use '01', for table 12 use '12' for xx in the MAC address.

```
ICX-Tx-Core1(config)# interface vlan 12
ICX-Tx-Core1(config-if-vlan)# ip address 10.x.12.2/24
ICX-Tx-Core1(config-if-vlan)# l3-counters
ICX-Tx-Core1(config-if-vlan)# vsx-sync active-gateways policies
ICX-Tx-Core1(config-if-vlan)# active-gateway ip 10.x.12.1 mac 02:02:00:00:xx:00
ICX-Tx-Core1(config-if-vlan)# exit
```

Core2

11. Configure VLAN12 on Core2.

```
ICX-Tx-Core2(config)# interface vlan 12
ICX-Tx-Core2(config-if-vlan)# ip address 10.x.12.3/24
ICX-Tx-Core2(config-if-vlan)# l3-counters
ICX-Tx-Core2(config-if-vlan)# exit
```

VLAN 1 Active Gateway

Setup VLAN 1 active gateway. This will be used for communication to the Aruba mobility controller and ensure the mobility controller has a redundant default gateway.

Core1

12. On Core1, configure the VLAN1 active gateway.

```
ICX-Tx-Core1(config)# interface vlan1
ICX-Tx-Core1(config-if-vlan)# vsx-sync active-gateways policies
ICX-Tx-Core1(config-if-vlan)# active-gateway ip mac 02:02:00:00:xx:00
ICX-Tx-Core1(config-if-vlan)# active-gateway ip 10.x.1.1
ICX-Tx-Core1(config-if-vlan)# exit
```

NOTE: Core2 has been configured with an IP address already, the vsx-sync will ensure the active gateway configuration is set on Core2 as well.

Configure DHCP Relay on VSX Core

Diagram

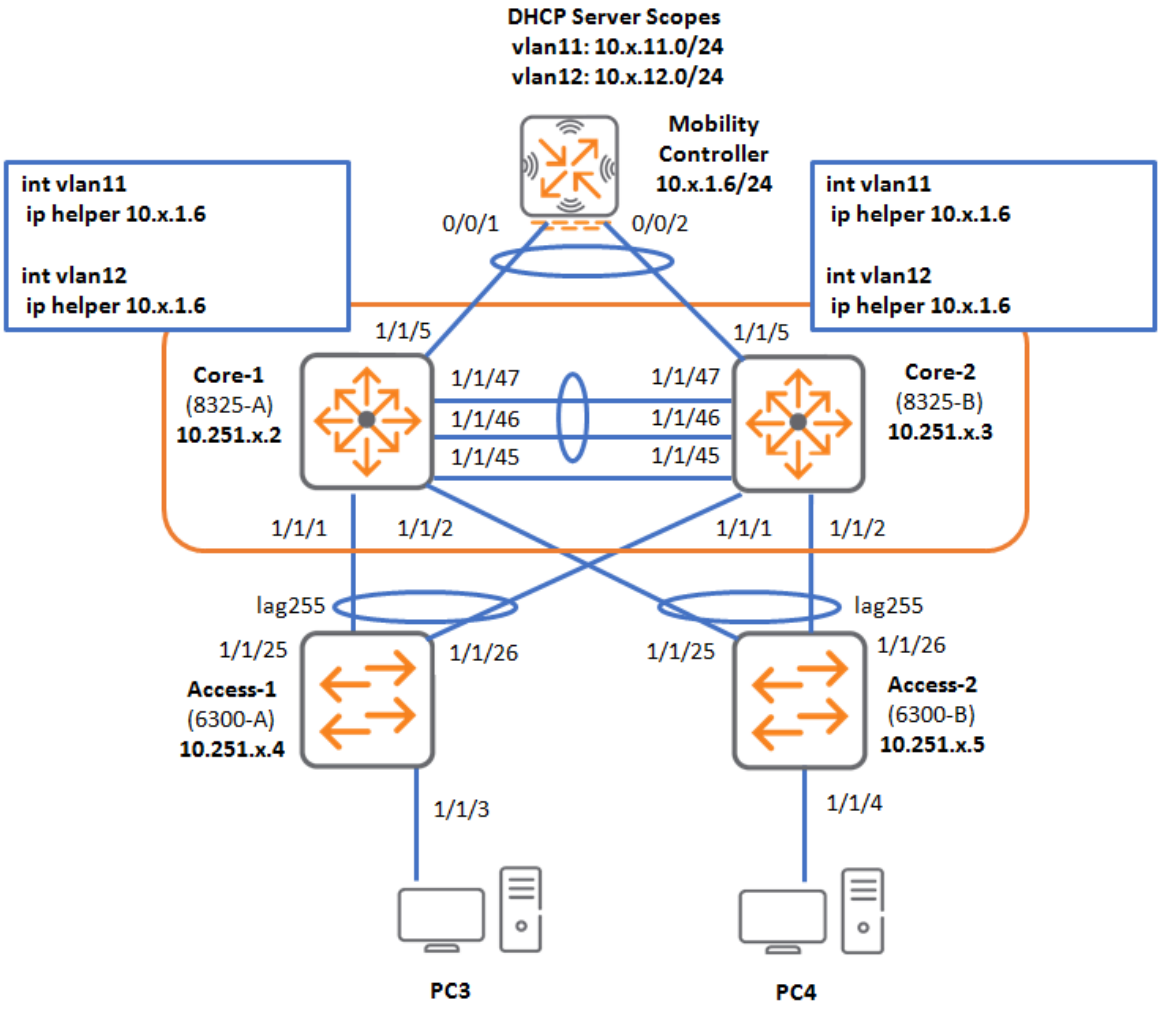

In this lab environment, DHCP scopes have been configured on the MC. On the VSX Core, the DHCP relay, also known as **'ip helper-address**' will be configured to point to the MC as the DHCP server.

Core1

13. On the Core1, apply the 'ip helper' command to VLAN 11. The MC has IP address 10.x.1.6.

```
ICX-Tx-Core1(config)# interface vlan11
ICX-Tx-Core1(config-if-vlan)# ip helper-address 10.x.1.6
ICX-Tx-Core1(config-if-vlan)# exit
```

14. Repeat the configuration for VLAN12.

```
ICX-Tx-Core1(config)# interface vlan12
```

```
ICX-Tx-Core1(config-if-vlan)# ip helper-address 10.x.1.6
ICX-Tx-Core1(config-if-vlan)# exit
```

15. Verify the 'ip helper' configuration was synchronized to Core2.

```
ICX-Tx-Core1(config)# show running-config interface vlan11 vsx-peer
interface vlan11
    vsx-sync active-gateways policies
    ip address 10.12.11.3/24
    active-gateway ip mac 02:02:00:00:12:00
    active-gateway ip 10.12.11.1
    l3-counters
   ip helper-address 10.x.1.6
    exit
```

Verify Active gateway operation on the client PC

Access1

16. Open a terminal connection to Access1, enter the configuration context.
17. Assign interface 1/1/3 (connected to the client PC) to VLAN 11.

```
ICX-Tx-Access1(config)# interface 1/1/3
ICX-Tx-Access1(config-if)# vlan access 11
ICX-Tx-Access1(config-if)# no shutdown
ICX-Tx-Access1(config-if)# exit
```

18. Open a connection to PC3 (connected to Access1 port 1/1/3).
19. On the client PC3, renew your IP address. Open a command prompt (cmd.exe) with administrator privileges, release and renew the IP address using "**ipconfig /release**" and "**ipconfig /renew**".

TIP: You may also choose to disable and enable the 'Lab NIC' interface in the network connections list.
Click **Start** > **Settings** > **Network & Internet** > **Ethernet** > **Change adapter options**.
Right-click on Lab NIC and select **Disable**.

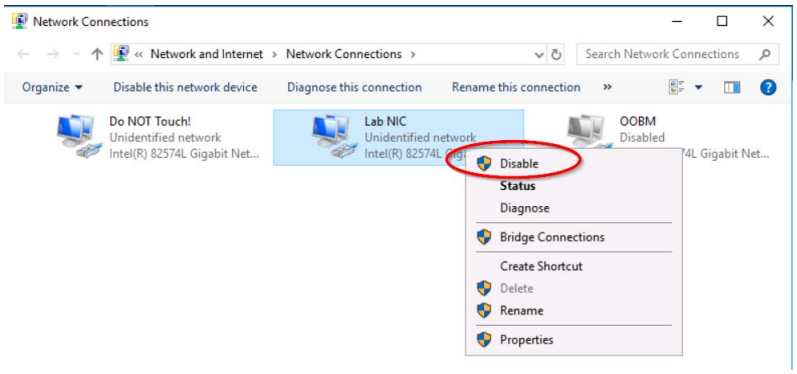

Wait a few moments, then right-click the Lab NIC and select **Enable**

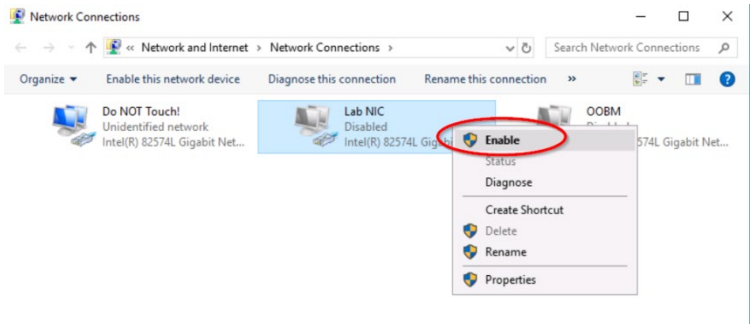

In the rest of the course labs, this procedure is referred to as **bounce the Lab NIC** (disable and enable it)

The IP address can be shown by right-clicking the **Lab NIC** > **Status** > **Details**.

20. On the Access1 console, verify the mac-address table of VLAN 11. The mac-address of the client should now be listed on interface 1/1/3.

Q: Can you explain the 3 mac addresses that have been learned on lag255?

A: For both core switches, the system MAC address, and the active gateway MAC.

```
ICX-Tx-Access1(config)# show mac-address-table vlan 11
MAC age-time             : 300 seconds
Number of MAC addresses : 4

MAC Address          VLAN     Type                        Port
-------------------------------------------------------------
90:20:c2:bc:97:00    11       dynamic                     lag255
90:20:c2:bc:17:00    11       dynamic                     lag255
00:50:56:b1:7a:37    11       dynamic                     1/1/3
02:02:00:00:12:00    11       dynamic                     lag255
```

21. On the client PC3, verify the system has received an IP address from the VLAN 11 ip range (10.x.11.0/24). In a command prompt (cmd.exe), run **'ipconfig'**.

```
C:\Users\student> ipconfig

Windows IP Configuration


Ethernet adapter Do NOT Touch!:

...

Ethernet adapter Lab NIC:

   Connection-specific DNS Suffix  . :
   IPv4 Address. . . . . . . . . . . : 10.x.11.51
```

```
   Subnet Mask . . . . . . . . . . . : 255.255.255.0
   Default Gateway . . . . . . . . . : 10.x.11.1

Ethernet adapter OOBM:

...

C:\Users\student>
```

22. Verify the IP path with a traceroute to the MC (10.x.1.6).

```
C:\Users\student>tracert -d 10.x.1.6

Tracing route to 10.12.1.6 over a maximum of 30 hops

  1    <1 ms    <1 ms    <1 ms  10.x.11.2
  2    <1 ms    <1 ms    <1 ms  10.x.1.6

Trace complete.

C:\Users\student>
```

Q: The client PC default gateway is 10.x.11.1, why does the traceroute show 10.x.11.2 or 10.x.11.3?

A: The core switch that receives the request (either Core1 or Core2) will respond with its local interface IP address.

23. **Optional step**: Verify the ARP entry on the PC3. In a command prompt (cmd.exe), enter **'arp -a'**. The IP address of the default gateway (10.x.11.1) should be listed with the active gateway MAC address. This confirms that the configured VSX active gateway MAC is learned by the client's systems.

```
...
Interface: 10.12.11.51 --- 0x10
  Internet Address      Physical Address      Type
  10.x.11.1             02-02-00-00-12-00     dynamic
  10.x.11.2             90-20-c2-bc-17-00     dynamic
  10.x.11.3             90-20-c2-bc-97-00     dynamic
  10.x.11.255           ff-ff-ff-ff-ff-ff     static
  224.0.0.22             01-00-5e-00-00-16     static
  224.0.0.251            01-00-5e-00-00-fb     static
  224.0.0.252            01-00-5e-00-00-fc     static
  239.255.255.250        01-00-5e-7f-ff-fa     static
  255.255.255.255        ff-ff-ff-ff-ff-ff     static
...
```

Verify DHCP relay for PC4 - The client connected to Access2

Access2

24. Open a terminal connection to Access2, enter the configuration context.
25. Assign interface 1/1/4 (connected to the client PC4) to VLAN 12.

```
ICX-Tx-Access2(config)# interface 1/1/4
ICX-Tx-Access2(config-if)# vlan access 12
ICX-Tx-Access2(config-if)# no shutdown
ICX-Tx-Access2(config-if)# exit
```

26. On PC4 (connected to Access2 port 1/1/4), renew your IP address. Open a command prompt (cmd.exe) with administrator privileges, release and renew the IP address using "**ipconfig /release**" and "**ipconfig /renew**" or by **bouncing** the **Lab NIC** network connection (disable/enable).

NOTE: If PC4 did not get an IP address in the 10.x.12.0/24 range, review your IP helper and VLAN configuration.

Optional Task 6: VSX Failover Tests

This task is **optional** and can be done if time permits. Check with your instructor.

Diagram

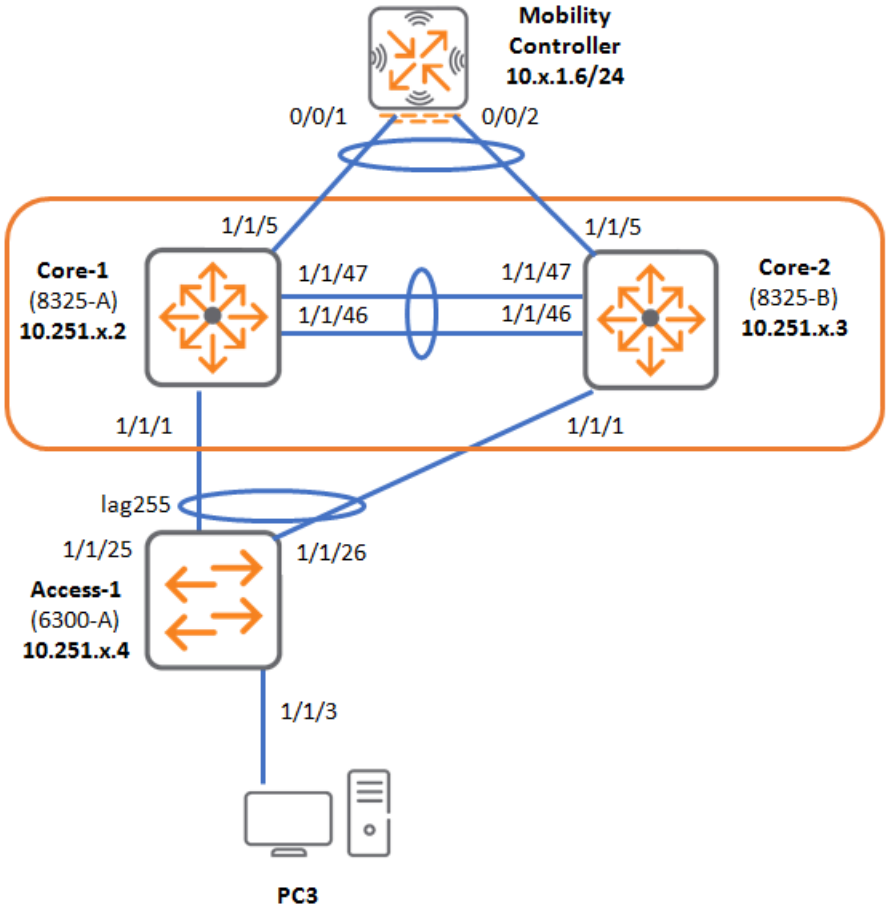

Objectives

- Link failure, verify L3 inter-VLAN traffic failover
- Link failure of a single ISL port, without causing a split-brain

Steps

Link failure, verify L3 inter-VLAN traffic failover

In this section, a ping will be done from PC3 (connected to Access1 1/1/3 - VLAN11) to the MC (connected to VLAN1).

While the ping is running, Access 1 uplinks will be disabled and re-enabled, one by one.

The result should be that the ping has only a minor interruption.

1. Open a session to PC3. Open a command prompt and start a continuous ping to the MC (10.x.1.6) .

```
C:\Users\student> ping 10.x.1.6 -t

Pinging 10.x.1.6 with 32 bytes of data:
Reply from 10.x.1.6: bytes=32 time<1ms TTL=63
Reply from 10.x.1.6: bytes=32 time<1ms TTL=63
...
```

2. Open a terminal connection to Access1, clear the interface statistics and then check the interface statistics for the uplink LAG.

The ping from PC3 sends an echo request and the MC will send an echo reply back. This is why the output will show both RX and TX statistics. The distribution of traffic depends on the hashing algorithm.

```
ICX-Tx-Access1# clear interface statistics

ICX-Tx-Access1# show interface lag255 statistics
```

Example output for the statistics:

```
-------------------------------------------------------------------------------------------------------------------
Interface       RX Bytes  RX Packets   TX Bytes  TX Packets  RX Broadcast  RX Multicast  TX Broadcast   TX Multicast
-------------------------------------------------------------------------------------------------------------------
1/1/25 - lag255  2279     17           2811      20          3             1             0                7
1/1/26 - lag255  1755     18           1148      14          3             1             0                0
lag255           4034     35           3959      34          6             2             0                7
```

3. On Access1, enter the configuration mode and disable uplink 1.

```
ICX-Tx-Access1# configure terminal
ICX-Tx-Access1(config)# interface 1/1/25
ICX-Tx-Access1(config-if)# shutdown
```

4. On PC3, check the ping operation is still successful. It is normal to either have no packet loss or 1 missed ping, since the ping is only performed once per second.
5. On Access1, enable the uplink1, verify the LACP state is operational again.

```
ICX-Tx-Access1(config-if)# no shutdown
```

Example of link that has **not yet completed** the LACP negotiation state (this output would only be seen a few seconds after enabling the interface).

```
ICX-Tx-Access1(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  lag255     26    1     ALFO    88:3a:30:97:b6:00 65534  256  lacp-block
1/1/26  lag255     27    1     ALFNCD  88:3a:30:97:b6:00 65534  256  up


Partner details of all interfaces:
--------------------------------------------------------------------------------
Intf    Aggr       Port  Port  State   System-ID         System Aggr
        Name       Id    Pri                             Pri    Key
--------------------------------------------------------------------------------
1/1/25  lag255     2     1     ALFO    02:01:00:00:01:00 65534  2
1/1/26  lag255     1002  1     ALFNCD  02:01:00:00:01:00 65534  2
```

Example of an operational link:

```
ICX-Tx-Access1(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  lag255     26    1     ALFNCD  88:3a:30:97:b6:00 65534  256  up
1/1/26  lag255     27    1     ALFNCD  88:3a:30:97:b6:00 65534  256  up


Partner details of all interfaces:
--------------------------------------------------------------------------------
Intf    Aggr       Port  Port  State   System-ID         System Aggr
        Name       Id    Pri                             Pri    Key
--------------------------------------------------------------------------------
1/1/25  lag255     2     1     ALFNCD  02:01:00:00:12:00 65534  2
1/1/26  lag255     1002  1     ALFNCD  02:01:00:00:12:00 65534  2
ICX-Tx-Access1(config-if)#
```

6. On Access1, disable uplink2.

```
ICX-Tx-Access1(config-if)# interface 1/1/26
ICX-Tx-Access1(config-if)# shutdown
```

7. On PC3, check the ping operation is still successful. It is normal to either have no packet loss or 1 missed ping, since the ping is only performed once per second.

8. On Access1, enable the uplink2, verify it is operational again (check for the 'up' state in the LACP output).

```
ICX-Tx-Access1(config-if)# no shutdown
```

```
ICX-Tx-Access1(config-if)# show lacp interfaces
```

NOTE: You may need to wait a few seconds and repeat the 'show lacp interfaces' command to see the 'up' state.

NOTE: You may leave the continuous ping active on PC3, it will be used in the next activity.

This demonstrates a link failure and traffic failover.

Link failure of a single ISL port

Diagram

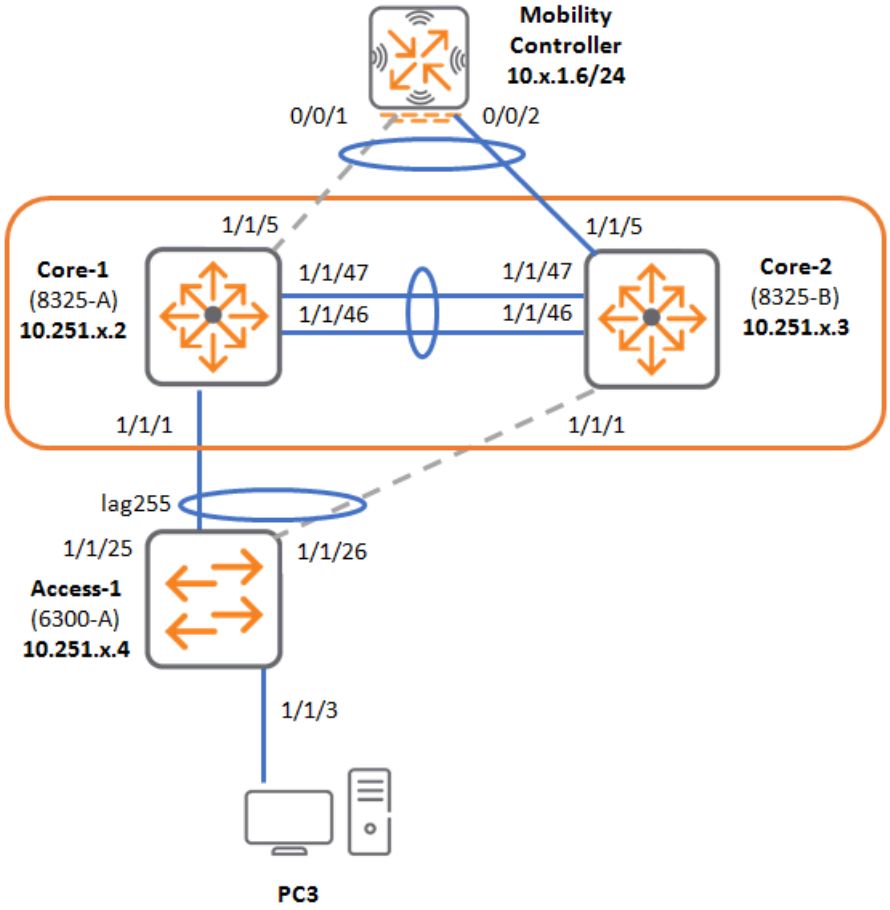

In this section, traffic will be forced over the ISL LAG and an ISL member port failure will be tested.

The traffic will be forced over the ISL link by:

- On the Access1, disabling uplink 2, so Access1 will send all traffic to Core1.
- On the Core1, disabling port 1/1/5, so the MC will be forced to send all traffic to Core2.

This way, the traffic will no longer use the 'local path' on each Core switch, but it must take the ISL path to reach the destination.

Once this has been done, 1 of the ISL member ports will be disabled to verify the ISL LAG failover. Do not shutdown the entire ISL at this point, that will result in a split-brain scenario, this will be tested in the next task.

Prepare the setup

9. On Access1, disable the uplink to Core2.

```
ICX-Tx-Access1(config)# interface 1/1/26
```

```
ICX-Tx-Access1(config-if)# shutdown
```

10. On Core1, disable the local link to the MC.

```
ICX-Tx-Core1# configure terminal
ICX-Tx-Core1(config)# interface 1/1/5
ICX-Tx-Core1(config-if)# shutdown
ICX-Tx-Core1(config-if)# exit
```

11. On PC3, verify the ping to the MC is still working fine.

Test ISL LAG failover

Now the ISL LAG failover can be tested.

12. On the Core1, disable ISL member port 1 (1/1/46).

```
ICX-Tx-Core1(config)# interface 1/1/46
ICX-Tx-Core1(config-if)# shutdown
```

13. On PC3, verify the ping is still operating fine.
14. On the Core1, enable ISL member port 1, wait a few seconds, next verify that the **LAG 256** is UP again.

```
ICX-Tx-Core1(config-if)# no shutdown

ICX-Tx-Core1(config-if)# show lacp interfaces
```

15. Next disable ISL member port 2 (1/1/47).

```
ICX-Tx-Core1(config-if)# interface 1/1/47
ICX-Tx-Core1(config-if)# shutdown
```

16. On PC3, verify the ping is still operating fine.
17. On Core1, enable the ISL member port 2 again and verify the status.

```
ICX-Tx-Core1(config-if)# no shutdown
ICX-Tx-Core1(config-if)# exit
```

```
ICX-Tx-Core1(config)# show lacp interfaces
```

18. On PC3, verify the ping is still operating fine.

This demonstrates the ISL LAG member port failover.

Restore the uplink connections on Access1 and Core1.

19. On Access1, restore the uplink to Core2 (1/1/26), the session is probably still in the interface 1/1/26 context.

```
ICX-Tx-Access1(config-if)# no shutdown
ICX-Tx-Access1(config-if)# exit
```

20. On Core1, restore the link to the MC (port 1/1/5).

```
ICX-Tx-Core1(config)# int 1/1/5
ICX-Tx-Core1(config-if)# no shutdown
ICX-Tx-Core1(config-if)# exit
```

21. On Core1, the LACP states for all LAGs are UP.

```
ICX-Tx-Core1(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     ALFNCD  02:01:00:00:12:00 65534  1    up
1/1/2   lag2(mc)   2     1     ALFNCD  02:01:00:00:12:00 65534  2    up
1/1/5   lag5(mc)   5     1     ALFNCD  02:01:00:00:12:00 65534  5    up
1/1/46  lag256     47    1     ALFNCD  90:20:c2:bc:17:00 65534  256  up
1/1/47  lag256     48    1     ALFNCD  90:20:c2:bc:17:00 65534  256  up
```

This concludes the ISL LAG failover test and the optional Task.

Task 7: VSX Split-Brain Handling

Diagram

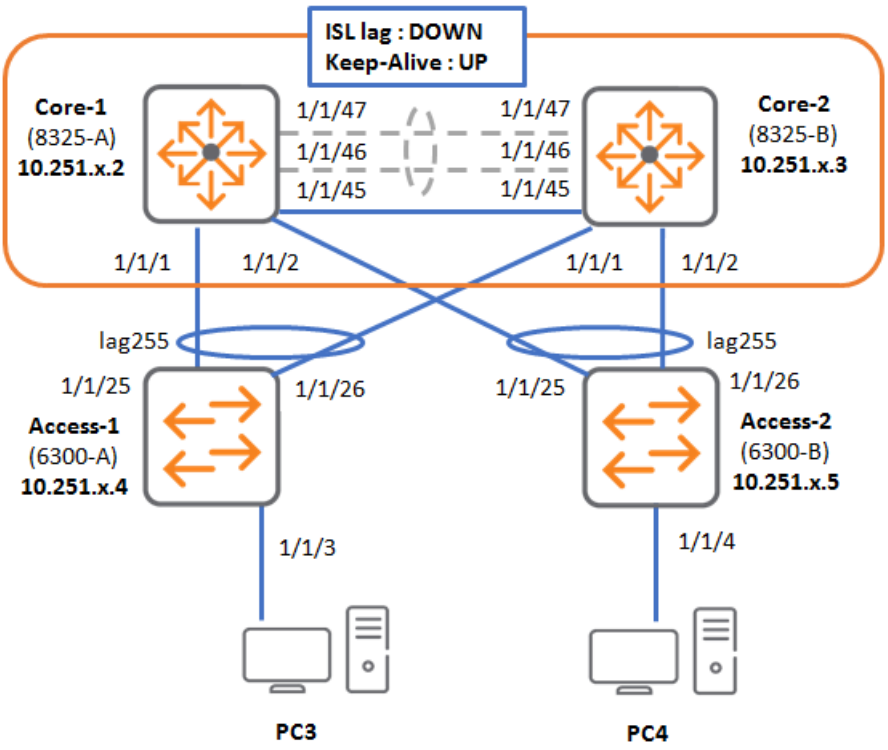

Objectives

- Split-brain scenario, review result / output / logs
- Split-brain recovery
- Link up delay timer

Handling a split-brain scenario

Primary ISL protection

The ISL between the VSX members should be primarily protected against link failures by using a LAG of multiple member ports, so when a single port between the VSX members fails, this will have no impact on the VSX functionality. Only the link capacity between the VSX members will be reduced.

Split Brain problem

However, in case all the members ports of the ISL link would be down, there will be no more VSX state replication between the VSX members, so the primary VSX member may have different MAC/ARP entries compared to the secondary VSX member, and they will not be able to share the STP state between them.

Solution

This scenario would lead to unpredictable results, so the solution will be to have the secondary node disable its interfaces. The result will be that only the primary node will still be on the network, so the network will operate in a predictable way again.

The secondary node must have a means to detect that the primary node is still active. In case the primary node suffered a complete power failure, the ISL link will also be down, but the secondary is supposed to take over complete control. In this case, we want to have the secondary active on the network.

Only when the primary *and* the secondary node are online at the same time, do we want to have the secondary disable (shutdown) its interfaces.

How - Keepalive

This can be achieved by configuring a keepalive between the primary and the secondary node. When the ISL is down, and the secondary still gets a response over the Keepalive, the secondary knows both are still online, and it will disable the VSX LAG interfaces.

In case the ISL goes down, and the keepalive does not respond, the secondary knows the primary is no longer online, so it will keep the VSX LAG interfaces up.

Steps

The keepalive has been configured as part of the basic VSX configuration task, in these steps, the configuration is reviewed.

1. Open a console connection to Core1.
2. The keepalive is an IP connection between the Core1 and Core2 switches. To make sure this connection is never in conflict or impacted by the normal routing table, a dedicated VRF is typically used.

Review the routed port 1/1/45, this is a dedicated link that was used for the VSX keepalive.

Core1

```
ICX-Tx-Core1# show running-config interface 1/1/45
interface 1/1/45
    no shutdown
    vrf attach KA
    ip address 192.168.0.0/31
    exit
```

Review the VSX running configuration.

```
ICX-Tx-Core1# show running-config vsx
vsx
    system-mac 02:01:00:00:xx:00
```

```
    inter-switch-link lag 256
    role primary
    keepalive peer 192.168.0.1 source 192.168.0.0 vrf KA
interface lag 256
   description ISL link
   no shutdown
   no routing
    vlan trunk native 1 tag
    vlan trunk allowed all
   lacp mode active
<output omitted..>
```

Review the VSX keepalive configuration.

```
ICX-Tx-Core1# show vsx configuration keepalive
Keepalive Interface : 1/1/45
Keepalive VRF       : KA
Source IP Address   : 192.168.0.0
Peer IP Address     : 192.168.0.1
UDP Port            : 7678
Hello Interval      : 1 Seconds
Dead Interval       : 3 Seconds
ICX-Tx-Core1#
```

Verify the keepalive operational status (Keepalive-Established) and the peer role (on Core1, the peer should be secondary).

```
ICX-Tx-Core1# show vsx status keepalive
Keepalive State           : Keepalive-Established
Last Established          : Mon Dec 23 19:15:09 2019
Last Failed               : Mon Dec 23 17:56:15 2019
Peer System Id            : 02:01:00:00:xx:00
Peer Device Role          : secondary

Keepalive Counters
Keepalive Packets Tx      : 352151
Keepalive Packets Rx      : 352020
Keepalive Timeouts        : 0
Keepalive Packets Dropped : 0
ICX-Tx-Core1#
```

Core2

3. Open a console connection to Core2, review the VSX keepalive with the same commands.

```
show running-config interface 1/1/45
show running-config vsx
show vsx configuration keepalive
show vsx status keepalive
```

Force a split brain

Core1

4. On Core1, enter configuration mode and disable the LAG 256 (LAG used for the VSX ISL).

```
ICX-Tx-Core1(config)# interface lag 256
ICX-Tx-Core1(config-lag-if)# shutdown
ICX-Tx-Core1(config-lag-if)# exit
```

5. On Core1, review the VSX status.

```
ICX-Tx-Core1(config)# show vsx status
VSX Operational State
---------------------
  ISL channel             : Out-Of-Sync
  ISL mgmt channel        : inter_switch_link_down
  Config Sync Status      : out-of-sync
  NAE                     : peer_unreachable
  HTTPS Server            : peer_unreachable

Attribute           Local               Peer
------------        --------            --------
ISL link            lag256
ISL version         2
System MAC          02:01:00:00:xx:00   02:01:00:00:xx:00
Platform            8325
Software Version    GL.10.04.0030
Device Role         primary
ICX-Tx-Core1(config)#
```

6. Review the keepalive status.

```
ICX-Tx-Core1(config)# show vsx status keepalive
Keepalive State           : Keepalive-Established
Last Established          : Fri Dec 27 18:07:37 2019
Last Failed               : Mon Dec 23 17:56:15 2019
Peer System Id            : 02:01:00:00:xx:00
Peer Device Role          :

Keepalive Counters
Keepalive Packets Tx      : 352501
Keepalive Packets Rx      : 352370
Keepalive Timeouts        : 0
Keepalive Packets Dropped : 0
```

Q: Why are there no 'packets dropped' at this point?

A: Only the ISL is down at this point, the keep-alive operates just fine, so no keep-alive packets are dropped.

7. On Core1, notice that the 'vsx-peer' commands are no longer available, since these commands use the ISL link

```
ICX-Tx-Core1(config)# show vsx status keepalive vsx-peer
VSX Peer not reachable
ICX-Tx-Core1(config)#
```

Core2 - Split - System Detected

Diagram

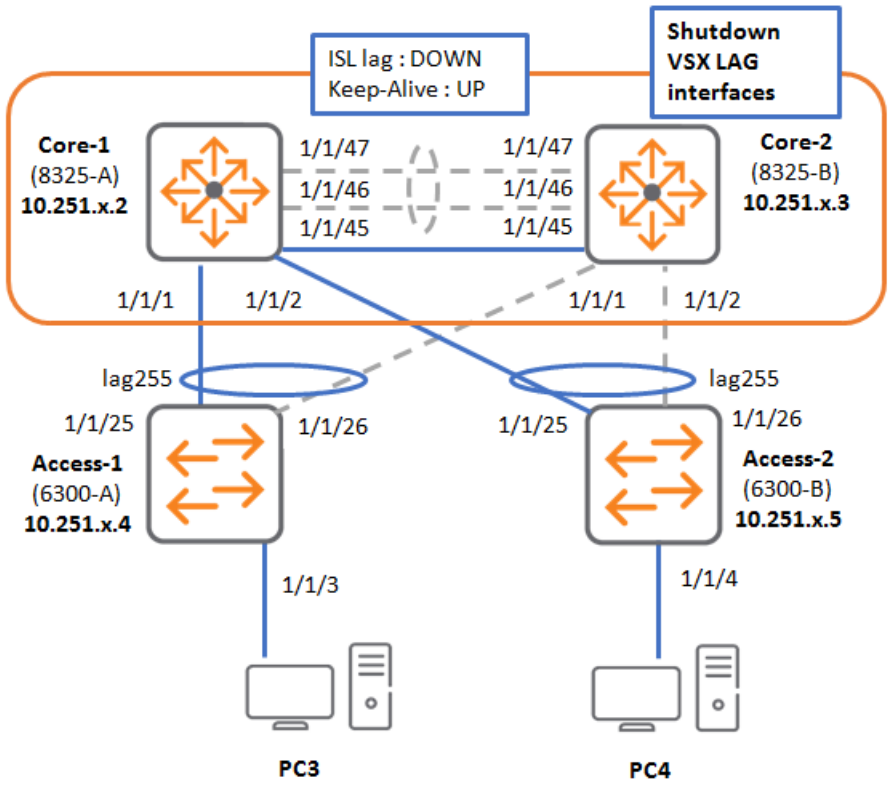

8. On Core2, verify that the keepalive has detected that the primary (Core1) is still responding (established).

```
ICX-Tx-Core2# show vsx status keepalive
Keepalive State           : Keepalive-Established
Last Established          : Fri Dec 27 18:08:18 2019
Last Failed               : Mon Dec 23 17:56:56 2019
Peer System Id            : 02:01:00:00:01:00
Peer Device Role          :
```

```
Keepalive Counters
Keepalive Packets Tx      : 352481
Keepalive Packets Rx      : 352482
Keepalive Timeouts        : 0
Keepalive Packets Dropped : 0
ICX-Tx-Core2#
```

9. Check the log buffer (-r reverse -n number of lines), look for lines containing 'vsx'. This will show an entry that the ISL link is down, but the Keep-Alive has been established. This indicates the 'split-system' scenario.

```
ICX-Tx-Core2# show logging -r -n 30 | include vsx
2020-03-02T08:35:36.312254+00:00 ICX-Tx-Core2 hpe-vsxd[2167]:
Event|7014|LOG_INFO|AMM|1/1|VSX 2 state local down, remote down
2020-03-02T08:35:36.306706+00:00 ICX-Tx-Core2 hpe-vsxd[2167]:
Event|7014|LOG_INFO|AMM|1/1|VSX 5 state local down, remote down
2020-03-02T08:35:36.257432+00:00 ICX-Tx-Core2 hpe-vsxd[2167]:
Event|7020|LOG_INFO|AMM|1/1|ISL out-of-sync and keepalive is in established
2020-03-02T08:35:36.254765+00:00 ICX-Tx-Core2 hpe-vsxkad[8376]:
Event|7006|LOG_INFO|AMM|1/1|VSX Keepalive succeeded
2020-03-02T08:35:35.594193+00:00 ICX-Tx-Core2 vsx-syncd[8800]:
Event|7602|LOG_INFO|AMM|-|Configuration sync update : VSX Inter-Switch-Link is
down.
2020-03-02T08:35:35.575153+00:00 ICX-Tx-Core2 hpe-vsxd[2167]:
Event|7011|LOG_INFO|AMM|1/1|VSX 5 state local up, remote down
2020-03-02T08:35:35.570479+00:00 ICX-Tx-Core2 hpe-vsxd[2167]:
Event|7011|LOG_INFO|AMM|1/1|VSX 2 state local up, remote down
2020-03-02T08:35:35.566540+00:00 ICX-Tx-Core2 hpe-vsxd[2167]:
Event|7014|LOG_INFO|AMM|1/1|VSX 1 state local down, remote down
2020-03-02T08:35:35.559596+00:00 ICX-Tx-Core2 hpe-vsxd[2167]:
Event|7004|LOG_ERR|AMM|1/1|VSX ISL port lag256 is Out-Of-Sync with the peer: link
is down
2020-03-02T08:35:35.559336+00:00 ICX-Tx-Core2 hpe-vsxd[2167]:
Event|7015|LOG_INFO|AMM|1/1|VSX ISL sliding window parameters are reset.
2020-03-02T08:35:35.559097+00:00 ICX-Tx-Core2 hpe-vsxd[2167]:
Event|7001|LOG_INFO|AMM|1/1|VSX ISL port lag256 is down
ICX-Tx-Core2#
```

10. Due to this condition, Core2 has disabled its local VSX LAG interfaces.

```
ICX-Tx-Core2# show interface brief
--------------------------------------------------------------------------------
----
Port      Native  Mode   Type          Enabled Status  Reason
Speed
          VLAN
(Mb/s)
--------------------------------------------------------------------------------
----
1/1/1     1       trunk  SFP+DAC1      yes     down    Disabled by VSX        --
1/1/2     1       trunk  SFP+DAC1      yes     down    Disabled by VSX        --
1/1/3     --      routed --            no      down    No XCVR installed      --
1/1/4     --      routed SFP-BT        no      down    Administratively down  --
```

```
1/1/5     1       trunk  SFP-BT          yes     down    Disabled by VSX         --
1/1/6     --      routed --              no      down    No XCVR installed       --
<output omitted..>

vlan1     --              --             yes     down    Disabled by VSX         --
vlan11    --              --             yes     down    Disabled by VSX         --
vlan12    --              --             yes     down    Disabled by VSX         --
<output omitted..>
```

11. And the VSX state will show as ' Split-System-Secondary '.

```
ICX-Tx-Core2# show vsx brief
ISL State                              : Out-Of-Sync
Device State                           : Split-System-Secondary
Keepalive State                        : Keepalive-Established
Device Role                            : secondary
Number of Multi-chassis LAG interfaces : 3
ICX-Tx-Core2#
```

Split-brain recovery - Link up delay

When the ISL gets restored, the secondary node will automatically re-enable the disabled interfaces after it has completed the VSX hardware table synchronization.

To provide enough time to synchronize all the tables and allow the secondary node time to learn routes from possible OSPF or BGP peers, VSX has a default Link up delay timer of 180 seconds (3 minutes).

Verify this default behavior in the next steps.

Core1

12. On the Core1, enable the LAG256 again.

```
ICX-Tx-Core1(config)# interface lag 256
ICX-Tx-Core1(config-lag-if)# no shut
ICX-Tx-Core1(config-lag-if)# exit
ICX-Tx-Core1(config)#
```

13. On Core2, review the VSX and interface status.

```
ICX-Tx-Core2# show vsx brief
ISL State                              : In-Sync
Device State                           : Sync-Secondary-Linkup-Delay
Keepalive State                        : Keepalive-Established
Device Role                            : secondary
Number of Multi-chassis LAG interfaces : 3
```

```
ICX-Tx-Core2# show interface brief
--------------------------------------------------------------------------------
----
```

```
Port      Native  Mode   Type          Enabled Status  Reason
Speed
          VLAN
(Mb/s)
--------------------------------------------------------------------------------
----
1/1/1     1       trunk  SFP+DAC1      yes     down    Disabled by VSX        --
1/1/2     1       trunk  SFP+DAC1      yes     down    Disabled by VSX        --
1/1/3     --      routed --            no      down    No XCVR installed      --
1/1/4     --      routed SFP-BT        no      down    Administratively down  --
1/1/5     --      routed SFP-BT        no      down    Administratively down  --
1/1/6     --      routed --            no      down    No XCVR installed      --
```

14. It may take up to 3 minutes before the VSX LAG interfaces are enabled again, wait for the interfaces to come up before continuing the lab.

NOTE: When the two switches have completed the synchronization of the forwarding tables before the timer expires, the interfaces will be enabled again. This is why it will not actually take 3 minutes in the lab, but typically less than a minute.

NOTE: It is possible for the administrator to manually set the link-up delay timer. This may be useful for very large deployments when many table entries need to be synchronized between the primary and secondary nodes. Consult the VSX best practice guide for more information.

Task 8: Finalize Configuration for Future Labs

Diagram

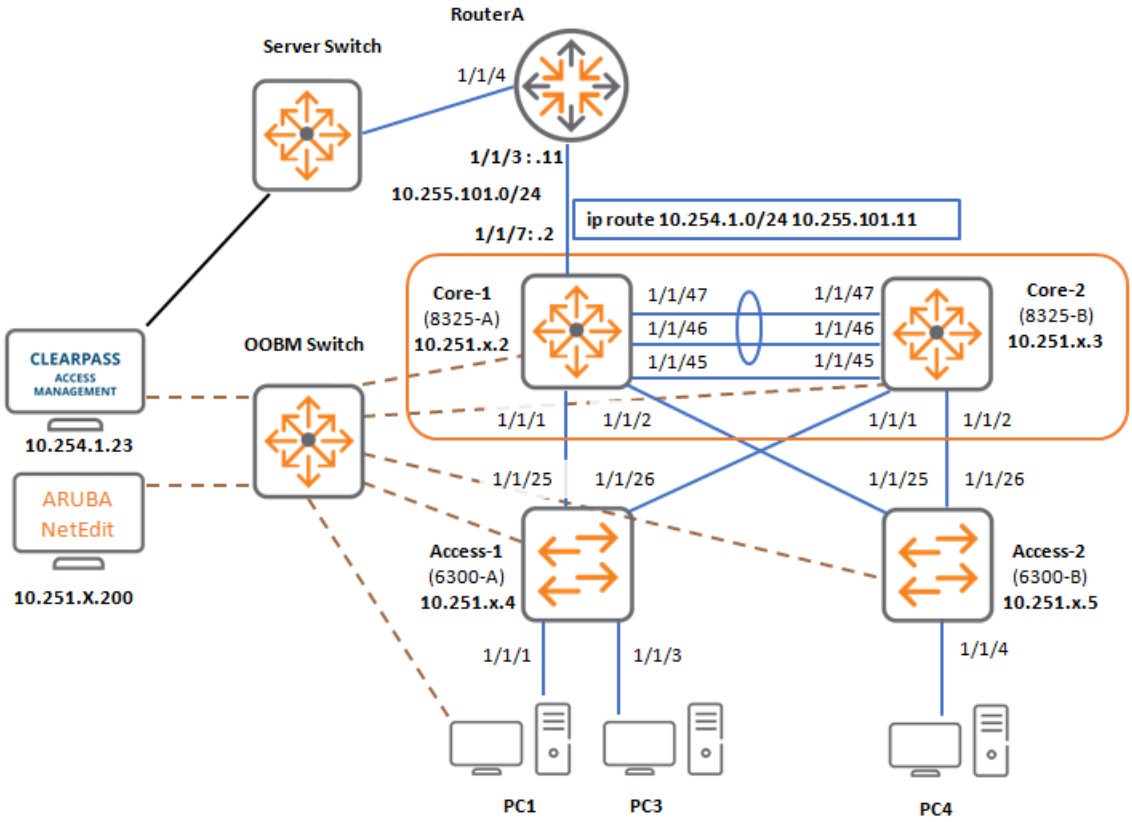

Objectives

In this task, you will complete the IP configuration to reach the ClearPass server. This will be required in future labs.

The ClearPass server is reachable via the RouterA device, that is connected to the Core1 port 1/1/7.

In this task, the port 1/1/7 will be enabled, an IP address will be assigned and a static route will be set for the ClearPass IP subnet.

Core2 will be configured with a static route to the ClearPass server via Core1.

The last step will be to save a new checkpoint, so you can revert to this configuration in future labs.

Steps

Core1

1. Access a terminal on Core1, enter the configuration mode.

2. Define the uplink connection to RouterA and configured the static route.

NOTE: This subnet is local for each table, so there is no 'x' used in these IP addresses.

```
ICX-Tx-Core1(config)# interface 1/1/7
ICX-Tx-Core1(config-if)# ip address 10.255.101.2/24
ICX-Tx-Core1(config-if)# no shutdown
ICX-Tx-Core1(config-if)# exit
ICX-Tx-Core1(config-)# ip route 10.254.1.0/24 10.255.101.11
```

3. Verify access to the ClearPass server.

```
ICX-Tx-Core1(config)# do ping 10.254.1.23 source vlan1
PING 10.254.1.23 (10.254.1.23) from 10.12.1.2 : 100(128) bytes of data.
108 bytes from 10.254.1.23: icmp_seq=1 ttl=61 time=4.72 ms
108 bytes from 10.254.1.23: icmp_seq=2 ttl=61 time=3.10 ms
```

NOTE: If the ping to the ClearPass host is not successful, verify the connection to the RouterA device.

```
ICX-Tx-Core1(config)# do ping 10.255.101.11
```

If this ping is not successful, contact your instructor.

Core2

4. Access a terminal on Core2, enter the configuration mode.
5. Configure a static route for the CPPM subnet to next-hop Core1.

```
ICX-Tx-Core2(config)# ip route 10.254.1.0/24 10.x.1.2
```

6. Verify connectivity to the ClearPass host.

```
ICX-Tx-Core2(config)# do ping 10.254.1.23 source vlan1
PING 10.254.1.23 (10.254.1.23) from 10.12.1.3 : 100(128) bytes of data.
108 bytes from 10.254.1.23: icmp_seq=1 ttl=60 time=3.67 ms
108 bytes from 10.254.1.23: icmp_seq=2 ttl=60 time=3.26 ms
...
ICX-Tx-Core2(config)# end
```

Save configuration checkpoints for future labs

You have now completed the VSX configuration lab.

IMPORTANT: This configuration will be the base configuration of several future labs, so a checkpoint **MUST** be created to complete the rest of the course lab activities.

7. On Core1, save the current configuration as a checkpoint.

```
ICX-Tx-Core1(config)# end
ICX-Tx-Core1# copy run checkpoint icx-lab04-vsx
Configuration changes will take time to process, please be patient.
```

8. Repeat this on Core2.

```
ICX-Tx-Core2(config)# end
ICX-Tx-Core2# copy run checkpoint icx-lab04-vsx
Configuration changes will take time to process, please be patient.
```

9. Repeat this on Access1.

```
ICX-Tx-Access2# copy run checkpoint icx-lab04-vsx
```

10. Repeat this on Access2.

```
ICX-Tx-Access2# copy run checkpoint icx-lab04-vsx
```

Verify the checkpoint

11. Verify on each switch that the checkpoint exists in the checkpoint list.

```
# show checkpoint list | include vsx
icx-lab04-vsx
```

You have completed Lab 4!

Lab 05: ACLs- Access-lists

Lab Diagram

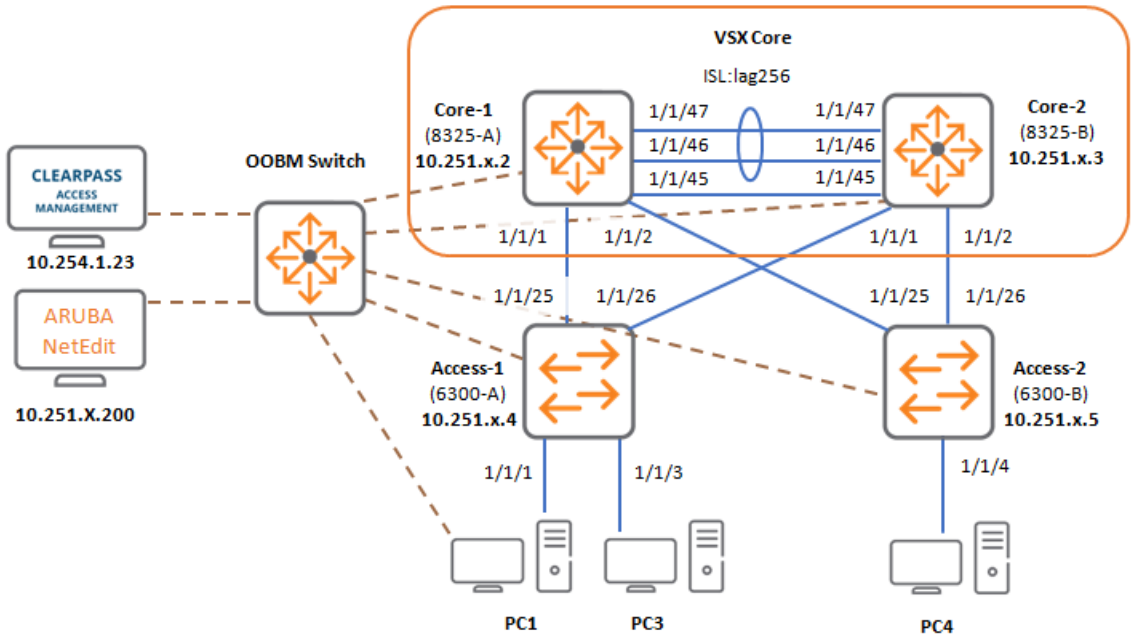

Overview

In this lab activity, Access Lists are configured and tested on the switch.

First, an Access List is defined and then access list rules are reviewed.

Next, the Access List is applied to a switch port and tested.

Next, the object group feature will be demonstrated. This allows the grouping of multiple IP addresses or ports in a logical group, so the Access List rules can be simplified.

The last part of the lab will review the resource utilization of the Access Lists in the hardware tables.

Objectives

- Define an Access List
- Understand the Access List rules, remove a rule and resequence the rules
- Apply an Access List on the switch
- Use object groups for IP addresses and ports
- Review the hardware resource utilization

Task 1: Verify Lab Start Configuration

Objectives

- If you have just completed Lab 04 - VSX, you can skip this task and move to the next task.
- If you have completed any other lab and are performing this lab again, complete these steps to get the base configuration on the devices.

Steps

1. Open a console connection to the 6300A. Login using **admin/aruba123**.

```
ICX-TX-Access1# copy checkpoint icx-lab04-vsx running-config
Configuration changes will take time to process, please be patient.
ICX-TX-Access1#
```

1. Open a console connection to the 6300B. Login using **admin/aruba123**.

```
ICX-TX-Access2# copy checkpoint icx-lab04-vsx running-config
Configuration changes will take time to process, please be patient.
ICX-TX-Access2#
```

1. Open a console connection to the 8325A. Login using **admin/aruba123**.

```
ICX-TX-Core1# copy checkpoint icx-lab04-vsx running-config
Configuration changes will take time to process, please be patient.
ICX-TX-Core1#
```

1. Open a console connection to the 8325A. Login using **admin/aruba123**.

```
ICX-TX-Core2# copy checkpoint icx-lab04-vsx running-config
Configuration changes will take time to process, please be patient.
ICX-TX-Core2#
```

Task 2: Port ACLs

Objectives

- Configure an Access List.
- Understand the Access List syntax and rules.
- Apply an Access List to an interface.

Steps

1. Open a terminal connection to Access2, enter the configuration mode.
2. Enable the port 1/1/4 to the PC4.

```
ICX-TX-Access2(config)# interface 1/1/4
ICX-TX-Access2(config-if)# no shutdown
ICX-TX-Access2(config-if)# exit
```

3. On the PC connected to Access2, verify that a DHCP IP address from the 10.x.12.0/24 subnet was received. (cmd > ipconfig)

Define access-list

4. On Access2, define a new access-list named **'pc'**.

```
ICX-TX-Access2(config)# access-list ip pc
```

5. Allow DHCP client traffic.

```
ICX-TX-Access2(config-acl-ip)# permit udp any eq 68 any eq 67
```

6. Allow traffic to the local subnet. The wildcard mask is not using inverse bits, but normal subnet mask logic, so '1' to match, '0' to ignore.

In this example any destination IP address that has '10' in the first byte, and '11' in the third byte will match the rule (255). The value in the second or fourth byte will be ignored (0). Enable the 'log' and 'count' function as well.

```
ICX-TX-Access2(config-acl-ip)# permit any any 10.0.11.0/255.0.255.0 log count
```

7. No further lines will be added. Any traffic that is not matching the access-list will be handled by the 'implicit deny' rule, so it will be blocked.
8. Review the access-list. Notice that each line now has a line number, so the order of the lines is important. The line numbers can be used to remove an individual line or to insert a new line between two existing lines.

```
ICX-TX-Access2(config-acl-ip)# show running-config current-context
access-list ip pc
    10 permit udp any eq 68 any eq 67
    20 permit any any 10.0.11.0/255.0.255.0 log count
```

9. Add a new line to match destination 10.0.12.0/255.0.255.0.

```
ICX-TX-Access2(config-acl-ip)# permit ip any 10.0.12.0/255.0.255.0 log count
```

10. Next verify the configuration again. This shows that entering the same source/destination information does not overwrite the existing line, but simply adds a line to the access-list.

```
ICX-TX-Access2(config-acl-ip)# show running-config current-context
access-list ip pc
    10 permit udp any eq 68 any eq 67
    20 permit any any 10.0.11.0/255.0.255.0 log count
    30 permit any any 10.0.12.0/255.0.255.0 log count
```

11. Next remove line 20 and review the configuration.

```
ICX-TX-Access2(config-acl-ip)# no 20
ICX-TX-Access2(config-acl-ip)# show running-config current-context
access-list ip pc
    10 permit udp any eq 68 any eq 67
    30 permit any any 10.0.12.0/255.0.255.0 log count
```

12. The switch supports resequencing the line numbers, this must be done from the global configuration. There is a start number and an increment number.

```
ICX-TX-Access2(config-acl-ip)# exit
ICX-TX-Access2(config)# access-list ip pc resequence 1000 100
```

13. Review the access-list configuration with the updated line numbers.

```
ICX-TX-Access2(config)# show access-list configuration commands
access-list ip pc
    1000 permit udp any eq 68 any eq 67
    1100 permit any any 10.0.12.0/255.0.255.0 log count
```

Apply the access-list

14. Enter the interface 1/1/4 context and apply the access-list to the port.

```
ICX-TX-Access2(config)# interface 1/1/4
ICX-TX-Access2(config-if)# apply access-list ip pc in
ICX-TX-Access2(config-if)# exit
```

15. Review the access-list configuration, notice how the access-list commands and the interface commands are combined in the output.

```
ICX-TX-Access2(config)# show access-list configuration commands
access-list ip pc
    1000 permit udp any eq 68 any eq 67
    1100 permit any any 10.0.12.0/255.0.255.0 log count
interface 1/1/4
    apply access-list ip pc in
```

Access-list applications

The next steps are for reference only, no actual configuration must be performed.

16. Reference only: Access-lists can be applied on switched (non-routed) ports in the inbound and outbound direction.

```
ICX-TX-Access2(config)# int 1/1/4
ICX-TX-Access2(config-if)# show run current-context
interface 1/1/4
    no shutdown
    apply access-list ip pc in
    no routing
    vlan access 12
    exit
ICX-TX-Access2(config-if)# apply access-list ip pc ?
  in   Inbound (ingress) traffic
  out  Outbound (egress) traffic
```

17. Reference only: Access-list can also be applied to routed ports (example only, no actual configuration required).

```
ICX-TX-Access2(config)# interface 1/1/2
ICX-TX-Access2(config-if)# show run cur
interface 1/1/2
    shutdown
    no routing
    vlan access 1
    exit
ICX-TX-Access2(config-if)# routing
ICX-TX-Access2(config-if)# show run cur
interface 1/1/2
    shutdown
    routing
    exit
ICX-TX-Access2(config-if)# apply access-list ip pc ?
  in   Inbound (ingress) traffic
  out  Outbound (egress) traffic
```

18. Reference only: Access-lists can also be applied to a switched VLAN (example only).

```
ICX-TX-Access2(config)# vlan 2
ICX-TX-Access2(config-vlan-2)# apply access-list ip pc ?
  in   Inbound (ingress) traffic
  out  Outbound (egress) traffic
```

NOTE: Access-lists *cannot* be applied to an SVI interface (interface VLAN x). If an access-list is required on an SVI interface, a traffic policy can be used on the VLAN interface.

```
ICX-TX-Access2(config)# int vlan 2
ICX-TX-Access2(config-if-vlan)# apply ?
  policy  Classifier policy
```

Test the access-list

19. On PC4, that is connected to Access2 port 1/1/4, open a command prompt with administrator privileges (cmd.exe) and attempt to release and renew the IP address. This should succeed, since this is allowed.

```
C:\Users\student>ipconfig /release
C:\Users\student>ipconfig /renew
```

20. Attempt to ping to 10.x.12.1. This should succeed since this is allowed.

```
C:\Users\student>ping 10.x.12.1

Pinging 10.x.12.1 with 32 bytes of data:
Reply from 10.x.12.1: bytes=32 time<1ms TTL=64
```

21. Attempt to ping to 10.x.11.1 This should fail, due to implicit deny, so when traffic does not match the access-list, it will be rejected.

```
C:\Users\student>ping 10.x.11.1

Pinging 10.12.11.1 with 32 bytes of data:
Request timed out.
```

22. On the Access2, review the access-list hitcount. In the example output, 9 packets have matched the permit rule.

```
ICX-TX-Access2(config)# show access-list hitcounts ip pc
Statistics for ACL pc (ipv4):
Interface 1/1/4* (in):
           Hit Count  Configuration
                   -  1000 permit udp any eq 68 any eq 67
                   9  1100 permit any any 10.0.12.0/255.0.255.0 log count
* access-list statistics are shared among each combination of
  context type (interface, VLAN, VRF) and direction (in, out, control-plane).
  Use 'access-list TYPE NAME copy' to create a uniquely named access-list.
```

23. Next review the log file on the switch using the command. This shows the log in reverse order, with 10 as the number of lines to display.

```
show logging -r -n 10
```

The log should show two type of entries, one for the line that was matched (event 10002), and one for the actual packet info that matched the line (event 10001).

```
ICX-TX-Access2(config)# show logging -r -n 10
-------------------------------------------------
Event logs from current boot
-------------------------------------------------
2020-02-12T12:45:50.356044+00:00 ICX-TX-Access2 ops-switchd[656]:
Event|10002|LOG_INFO|MSTR|1|pc on 1/1/4 (in):            2 1100 permit any any
10.0.12.0/255.0.255.0 log count
2020-02-12T12:44:02.854143+00:00 ICX-TX-Access2 hpe-restd[726]:
Event|4646|LOG_ERR|AMM|-|Aruba Activate server https://devices-
v2.arubanetworks.com is not reachable through any supported VRF.
2020-02-12T12:40:50.109684+00:00 ICX-TX-Access2 ops-switchd[656]:
Event|10001|LOG_INFO|MSTR|1|List pc, seq# 1100 permitted icmp 10.x.12.37 ->
10.x.12.1 type 8 code 0, on vlan 12, port 1/1/4, direction in
2020-02-12T12:39:02.854417+00:00 ICX-TX-Access2 hpe-restd[726]:
Event|4646|LOG_ERR|AMM|-|Aruba Activate server https://devices-
v2.arubanetworks.com is not reachable through any supported VRF.
2020-02-12T12:34:02.854338+00:00 ICX-TX-Access2 hpe-restd[726]:
Event|4646|LOG_ERR|AMM|-|Aruba Activate server https://devices-
v2.arubanetworks.com is not reachable through any supported VRF.
2020-02-12T12:32:54.822361+00:00 ICX-TX-Access2 ops-switchd[656]:
Event|10002|LOG_INFO|MSTR|1|pc on 1/1/4 (in):            1 1100 permit any any
10.0.12.0/255.0.255.0 log count
2020-02-12T12:29:02.854213+00:00 ICX-TX-Access2 hpe-restd[726]:
Event|4646|LOG_ERR|AMM|-|Aruba Activate server https://devices-
v2.arubanetworks.com is not reachable through any supported VRF.
2020-02-12T12:27:54.765734+00:00 ICX-TX-Access2 ops-switchd[656]:
Event|10001|LOG_INFO|MSTR|1|List pc, seq# 1100 permitted udp 10.x.12.37(138) ->
10.x.12.255(138) on vlan 12, port 1/1/4, direction in
2020-02-12T12:24:02.854226+00:00 ICX-TX-Access2 hpe-restd[726]:
Event|4646|LOG_ERR|AMM|-|Aruba Activate server https://devices-
v2.arubanetworks.com is not reachable through any supported VRF.
2020-02-12T12:22:54.698250+00:00 ICX-TX-Access2 ops-switchd[656]:
Event|10002|LOG_INFO|MSTR|1|pc on 1/1/4 (in):            1 1100 permit any any
10.0.12.0/255.0.255.0 log count
ICX-TX-Access2(config)#
```

24. To see the last packets that matched an access-list with the 'log' option, use the event ID to filter.

```
ICX-TX-Access2(config)# show logging -r -n 100 | i 10001
2020-02-12T12:40:50.109684+00:00 ICX-TX-Access2 ops-switchd[656]:
Event|10001|LOG_INFO|MSTR|1|List pc, seq# 1100 permitted icmp 10.x.12.37 ->
10.x.12.1 type 8 code 0, on vlan 12, port 1/1/4, direction in
```

```
2020-02-12T12:27:54.765734+00:00 ICX-TX-Access2 ops-switchd[656]:
Event|10001|LOG_INFO|MSTR|1|List pc, seq# 1100 permitted udp 10.x.12.37(138) ->
10.x.12.255(138) on vlan 12, port 1/1/4, direction in
2020-02-12T12:17:54.631987+00:00 ICX-TX-Access2 ops-switchd[656]:
Event|10001|LOG_INFO|MSTR|1|List pc, seq# 1100 permitted udp 10.x.12.37(138) ->
10.x.12.255(138) on vlan 12, port 1/1/4, direction in
2020-02-12T12:08:04.493868+00:00 ICX-TX-Access2 ops-switchd[656]:
Event|10001|LOG_INFO|MSTR|1|List pc, seq# 1100 permitted udp 10.x.12.37(138) ->
10.x.12.255(138) on vlan 12, port 1/1/4, direction in
2020-02-12T12:02:54.413634+00:00 ICX-TX-Access2 ops-switchd[656]:
Event|10001|LOG_INFO|MSTR|1|List pc, seq# 20 permitted udp 10.x.12.37(138) ->
10.x.12.255(138) on vlan 12, port 1/1/4, direction in
2020-02-12T11:54:59.289694+00:00 ICX-TX-Access2 ops-switchd[656]:
Event|10001|LOG_INFO|MSTR|1|List pc, seq# 20 permitted icmp 10.x.12.37 ->
10.x.12.1 type 8 code 0, on vlan 12, port 1/1/4, direction in
```

Task 3: Using object groups

Objectives

In this task, object groups for access-lists will be introduced.

Object groups group objects that would otherwise be used directly in access-lists.

For example, suppose a customer has two Web servers and they want to allow TCP port 80 and 443 (http/https) access to these two Web server IP addresses.

With a traditional ACL, 4 lines would be required:

- allow tcp 80 to IP1
- allow tcp 80 to IP2
- allow tcp 443 to IP1
- allow tcp 443 to IP2

With object groups, the customer can first assign both IP addresses to an object group (e.g. web-servers), and then use this group in the lines:

- allow tcp 80 to web-servers
- allow tcp 443 to web-servers

The administrator can also combine ports into an object group, so both 80 and 443 can be combined into a group (e.g. web-ports). The result would be:

- allow web-ports to web-servers

Whenever new ports or new IPs must be added, only the groups need to be updated, the ACL entry remains the same.

Steps

Define group for Core switches IP

1. Open the terminal connection to Access2, enter the configuration mode.
2. Define a new object group 'core-switch'.

```
ICX-TX-Access2(config)# object-group ip address core-switch
```

3. Check the possible object types.

```
ICX-TX-Access2(config-addrgroup-ip)# ?
  <1-4294967295>   Object Group sequence number
  A.B.C.D          Specify IP host address
  A.B.C.D/M        Specify IP network address with prefix length
  A.B.C.D/W.X.Y.Z  Specify IP network address with network mask
  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
  show            Show running system information
```

Q: What types of addresses can be added to the object group?

_____

A: A single IP address, a network address with a prefix mask, a network address with a network mask.

4. Add the two IP addresses of the 2 VSX Core switches in VLAN 12.

```
ICX-TX-Access2(config-addrgroup-ip)# 10.x.12.2
ICX-TX-Access2(config-addrgroup-ip)# 10.x.12.3
```

> **NOTE:** Do not add 10.x.12.1, this address will be used in a validation test, so it should not belong to the object-group.

5. Review the object group, notice the line numbers that are automatically generated.

```
ICX-TX-Access2(config-addrgroup-ip)# show run cur
object-group ip address core-switch
    10 10.x.12.2
    20 10.x.12.3
ICX-TX-Access2(config-addrgroup-ip)# exit
```

6. For reference only: Just like the regular access-list, entries in the object-group can be resequenced. This step is for reference only, no actual resequence is needed.

```
ICX-TX-Access2(config)# object-group ip address core-switch ?
  resequence  Re-number entries
  reset       Reset configuration
  <cr>
```

Define group for port 80 and 443

7. Define a new object group, named 'switch-ports'.

```
ICX-TX-Access2(config)# object-group port switch-ports
```

8. Review the object group options.

```
ICX-TX-Access2(config-portgroup)# ?
```

```
  <1-4294967295>  Object Group sequence number
  end             End current mode and change to enable mode.
  eq              Layer 4 port equal to
  exit            Exit current mode and change to previous mode
  gt              Layer 4 port greater than
  list            Print command list
  lt              Layer 4 port less than
  no              Negate a command or set its defaults
  range           Layer 4 port range
  show            Show running system information
```

9. Add the ports 80 and 443 to the object group.

```
ICX-TX-Access2(config-portgroup)# eq 80
ICX-TX-Access2(config-portgroup)# eq 443
```

10. Review the current configuration.

```
ICX-TX-Access2(config-portgroup)# show run cur
object-group port switch-ports
    10 eq 80
    20 eq 443
ICX-TX-Access2(config-portgroup)# exit
```

Combine the groups into an access-list

11. Enter the access-list 'pc' context, review the current settings.

```
ICX-TX-Access2(config)# access-list ip pc
ICX-TX-Access2(config-acl-ip)# show run cur
access-list ip pc
    1000 permit udp any eq 68 any eq 67
    1100 permit any any 10.0.12.0/255.0.255.0 log count
```

12. Overwrite the current line 1100 with the new line based on the object-groups.

```
ICX-TX-Access2(config-acl-ip)# 1100 permit tcp any core-switch group switch-ports
```

13. Verify the new configuration.

```
ICX-TX-Access2(config-acl-ip)# show run cur
access-list ip pc
    1000 permit udp any eq 68 any eq 67
    1100 permit tcp any core-switch group switch-ports
ICX-TX-Access2(config-acl-ip)# exit
```

Verify the operation

14. Using PC4, connected to Access2, open a browser and navigate to **https://10.x.12.2**. This should succeed. There is no need to login.
15. Navigate to **https://10.x.12.3**. This should succeed.
16. Attempt to navigate to **https://10.x.12.1**. This should fail.

Task 4: Resource usage

Objectives

Access-lists and traffic classifiers consume resources in the switch TCAM tables.

When large or many access-lists are configured, the administrator can review the current resource utilization of the TCAMs.

In this task, the resources will be reviewed, and the impact of object groups will be demonstrated.

Steps

- Review the current resource usage.
- Adjust the object groups, to increase the TCAM use.
- Verify the updated resource usage.

Review the current resource usage

Access2

1. On Access2, review the current usage.

```
ICX-TX-Access2(config)# show resources

Resource Usage:

Mod  Description
         Resource                                        Used     Free
-------------------------------------------------------------------------
1/1  Ingress IP Port ACL Lookup
         Ingress TCAM Entries                              6     5107
     Total
         Ingress Lookups                                   1        4
         Egress Lookups                                    0        4

Resource data is updated every 10 seconds.
```

2. For reference only: This is the output of Core1 (model 8325) resources.

```
ICX-TX-Core1(config)# show resources

Resource Usage:

Mod  Description
         Resource                              Width    Used Reserved    Free
-------------------------------------------------------------------------
1/1  Ingress Control Plane Policing
         Ingress TCAM Entries                      3     231     2304
     Egress Control Plane Policing
```

```
        Egress TCAM Entries                    2      84      512
    Total
        Ingress TCAM Entries                         231     2304    6912
        Egress TCAM Entries                           84      512    1536
        Policers                                       0             6144
        Ingress L4 Port Ranges                         0               32

Resource data is updated every 10 seconds.
ICX-TX-Core1(config)#
```

3. Enter the object-group ip address core-switch and add one IP address.

```
ICX-TX-Access2(config)# object-group ip address core-switch
ICX-TX-Access2(config-addrgroup-ip)# 10.x.12.1
```

4. Wait about 10 seconds, then review the resources.

```
ICX-TX-Access2(config-addrgroup-ip)# show resources

Resource Usage:

Mod  Description
        Resource                                        Used     Free
-------------------------------------------------------------------------------
1/1  Ingress IP Port ACL Lookup
        Ingress TCAM Entries                             8     5105
    Total
        Ingress Lookups                                  1        4
        Egress Lookups                                   0        4

Resource data is updated every 10 seconds.
```

5. Next add two more IP addresses to the group.

```
ICX-TX-Access2(config-addrgroup-ip)# 10.x.12.4
ICX-TX-Access2(config-addrgroup-ip)# 10.x.12.5
ICX-TX-Access2(config-addrgroup-ip)# exit
```

6. Wait about 10 seconds, then review the resources.

```
ICX-TX-Access2(config)# show resources

Resource Usage:

Mod  Description
        Resource                                        Used     Free
-------------------------------------------------------------------------------
1/1  Ingress IP Port ACL Lookup
        Ingress TCAM Entries                            12     5101
    Total
```

```
        Ingress Lookups                                 1        4
        Egress Lookups                                  0        4
```

7. Enter the port object-group and add a port and port range to the group.

```
ICX-TX-Access2(config)# object-group port switch-ports
ICX-TX-Access2(config-portgroup)# eq 22
ICX-TX-Access2(config-portgroup)# range 1024 2048
ICX-TX-Access2(config-portgroup)# exit
```

8. Wait about 10 seconds, then review the resources.

```
ICX-TX-Access2(config)# show resources

Resource Usage:

Mod  Description
        Resource                                      Used     Free
-------------------------------------------------------------------------
1/1  Ingress IP Port ACL Lookup
        Ingress TCAM Entries                           27     5086
     Total
        Ingress Lookups                                 1        4
        Egress Lookups                                  0        4

Resource data is updated every 10 seconds.
```

This shows that the object-groups are very convenient for the administrator, but internally they are extracted so they still consume the same hardware resources as if individual lines would have been created.

You have completed Lab 5

Lab 06.1: OSPF Single Area setup- Basic OSPF configuration

Lab Diagram

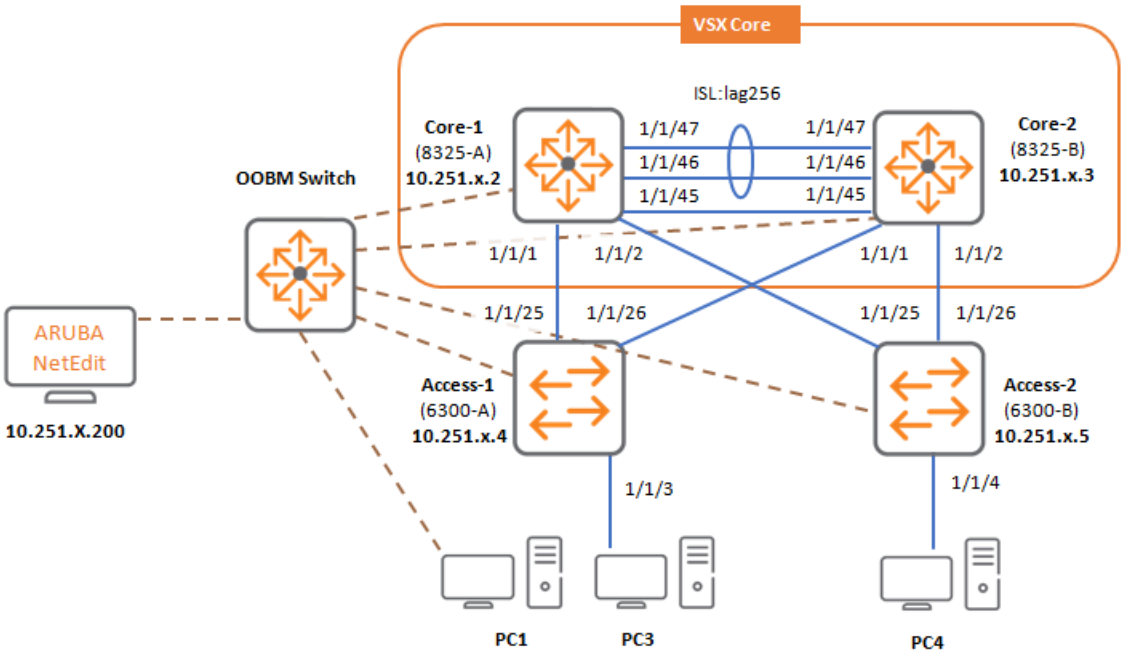

Overview

In this lab activity, OSPF will be configured using a single backbone area (0.0.0.0) on the devices in the lab.

In following lab activities, additional areas will be introduced.

Objectives

- Use loopback interfaces
- Configure single area OSPF on the AOS-CX switches
- Review OSPF passive interfaces

Task 1: Verify Lab Start Configuration

Objectives

- If you have just completed Lab 04 - VSX, you can skip this task and move to the next task.
- If you have completed any other lab and are performing this lab again, complete these steps to get the base configuration on the devices.

Steps

1. Open a console connection to the 6300A. Login using admin, password aruba123.

```
ICX-Tx-Access1# copy checkpoint icx-lab04-vsx running-config
Configuration changes will take time to process, please be patient.
ICX-Tx-Access1#
```

2. Open a console connection to the 6300B. Login using admin, password aruba123.

```
ICX-Tx-Access2# copy checkpoint icx-lab04-vsx running-config
Configuration changes will take time to process, please be patient.
ICX-Tx-Access2#
```

3. Open a console connection to the 8325A. Login using admin, password aruba123.

```
ICX-Tx-Core1# copy checkpoint icx-lab04-vsx running-config
Configuration changes will take time to process, please be patient.
ICX-Tx-Core1#
```

4. Open a console connection to the 8325B. Login using admin, password aruba123.

```
ICX-Tx-Core2# copy checkpoint icx-lab04-vsx running-config
Configuration changes will take time to process, please be patient.
ICX-Tx-Core2#
```

Task 2: Basic OSPF Setup on Core Area 0

Diagram L3

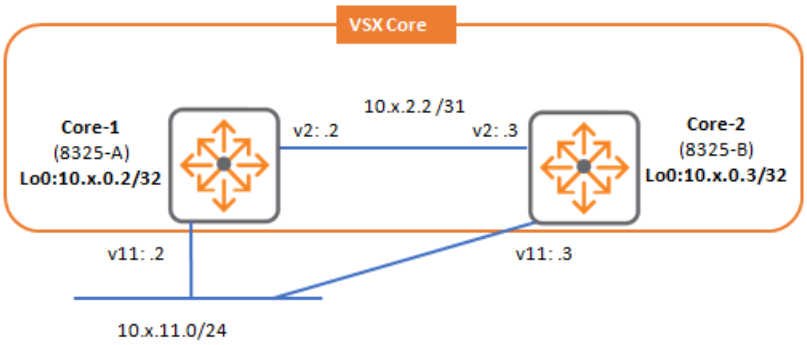

Objectives

- Use Loopback interfaces
- Refresh the OSPF single area configuration
- Setup OSPF in a VSX cluster
- Use OSPF point to point connections
- Use OSPF passive interface feature

Steps

Core1

1. Open a terminal connection to Core1, enter configuration mode.
2. Define a new IP Loopback interface and assign the IP address. Replace X with your table number.

```
ICX-Tx-Core1(config)# interface loopback 0
ICX-Tx-Core1(config-loopback-if)# ip address 10.x.0.2/32
ICX-Tx-Core1(config-loopback-if)# exit
```

3. Start a new OSPF process with process ID 1. Assign same IP as the Loopback as the router ID to ensure stability for the router ID.

   - The process ID is locally significant.

   - By default, the process is linked to the VRF 'default'

```
ICX-Tx-Core1(config)# router ospf 1
ICX-Tx-Core1(config-ospf-1)# router-id 10.x.0.2
```

4. Review current status.

```
ICX-Tx-Core1(config-ospf-1)# show ip ospf
Routing Process 1 with ID : 10.x.0.2 VRF default
-------------------------------------------------

OSPFv2 Protocol is enabled
Graceful-restart is configured
Restart Interval: 120, State: inactive
Last Graceful Restart Exit Status: none
SPF: Start Time: 200ms, Hold Time: 1000ms, Max Wait Time: 5000ms
Maximum Paths to Destination: 4
Number of external LSAs 0, checksum sum 0
Number of areas is 0, 0 normal, 0 stub, 0 NSSA
Number of active areas is 0, 0 normal, 0 stub, 0 NSSA
BFD is disabled
Reference Bandwidth: 100000 Mbps
```

5. Define area 0.

```
ICX-Tx-Core1(config-ospf-1)# area 0
```

6. Review current OSPF LSDB. Since no interfaces have been enabled for OSPF, the database does not contain any LSA information yet.

```
ICX-Tx-Core1(config-ospf-1)# show ip ospf lsdb
No OSPF LSAs found on VRF default.
ICX-Tx-Core1(config-ospf-1)# exit
```

7. Enable loopback interface for OSPF.

```
ICX-Tx-Core1(config)# interface loopback 0
ICX-Tx-Core1(config-loopback-if)# ip ospf 1 area 0
ICX-Tx-Core1(config-loopback-if)# exit
```

NOTE: The 'ip ospf' command allows fine control over which interface is assigned to which OSPF process and in which area.

8. Review the LSDB again.

```
ICX-Tx-Core1(config)# show ip ospf lsdb
OSPF Router with ID (10.x.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.x.0.2       10.x.0.2       27        0x80000001 0x00009773     1
```

Q: What type of LSA has been added into the LSDB?

_____

A: A Type1 LSA, known as a 'router' LSA.

Q: How many links (IP subnets) are announced by this LSA?

_____

A: Currently, only the loopback interface is announced in the LSA, so 1 link.

Dedicated routed link between Core1 and Core2

It is best practice to have a dedicated routed subnet between the VSX members. Dedicated means that there are no other devices on this subnet.

This subnet will be used as the transit network for any routes that would exist on only one of the VSX members.

While technically any existing subnet could be used for this, it provides a cleaner design to have a dedicated subnet for this purpose, a subnet that cannot be impacted by any other end-point devices.

Core1

9. On Core1, define VLAN 2 interface. It will be used as the transit network. Since only Core1 and Core2 will be using this subnet, a /31 mask is used.

> **NOTE:** It is best practice to use a /31 subnet for point to point links when both sides support the /31 mask.

```
ICX-Tx-Core1(config)# vlan 2
ICX-Tx-Core1(config-vlan-2)# vsx-sync
ICX-Tx-Core1(config-vlan-2)# exit
ICX-Tx-Core1(config)# interface vlan 2
ICX-Tx-Core1(config-if-vlan)# ip address 10.x.2.2/31
```

Core2

10. Open a terminal connection to Core2, enter configuration mode.
11. Define the VLAN 2 interface.

```
ICX-Tx-Core2(config)# interface vlan 2
ICX-Tx-Core2(config-if-vlan)# ip address 10.x.2.3/31
```

12. Verify connectivity to Core1 with a ping.

```
ICX-Tx-Core2(config-if-vlan)# do ping 10.x.2.2
PING 10.x.2.2 (10.x.2.2) 100(128) bytes of data.
108 bytes from 10.x.2.2: icmp_seq=1 ttl=64 time=14.1 ms
108 bytes from 10.x.2.2: icmp_seq=2 ttl=64 time=0.231 ms
...
ICX-Tx-Core2(config-if-vlan)# exit
```

Core1

13. Enable OSPF on interface VLAN 2 for process 1 area 0.

```
ICX-Tx-Core1(config-if-vlan)# ip ospf 1 area 0
```

NOTE: It is best practice to use the lowest possible cost for this link. By default, the cost is calculated based on the interface bandwidth, using 100Gbps as the reference. This results in a cost 100 by default for the VLAN interfaces.

For the VSX transit link, the lowest possible cost will be set, that is cost 1.

14. On Core1, apply this cost to the interface VLAN 2.

```
ICX-Tx-Core1(config-if-vlan)# ip ospf cost 1
ICX-Tx-Core1(config)# exit
```

15. Review the LSDB again.

```
ICX-Tx-Core1(config)# show ip ospf lsdb
OSPF Router with ID (10.x.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.x.0.2       10.x.0.2       91        0x80000002 0x0000322f     2

Network Link State Advertisements (Area 0.0.0.0)
------------------------------------------------

LSID            ADV Router      Age       Seq#       Checksum
-------------------------------------------------------------
10.x.2.3       10.x.12.3      45        0x80000003 0x00008f4c
```

Q: How many links are active for the router LSA 10.x.0.2?

A: There are 2 links now: loopback and VLAN2.

Core2

16. Configure Core2 for OSPF. First review the current OSPF configuration.

```
ICX-Tx-Core2(config)# show running-config ospf
router ospf 1
   area 0.0.0.0
interface vlan2
    ip ospf 1 area 0.0.0.0
    ip ospf cost 1
```

Q: Why is there already an OSPF configuration on Core2?

_____

A: VSX-sync is enabled for OSPF, so many OSPF configurations are automatically synchronized to Core2.

17. Configure the loopback and set the router-id.
    - IP Loopback 0 10.x.0.3/32
    - OSPF process 1, router-id 10.x.0.3/32, define area0

```
ICX-Tx-Core2(config)# interface loopback 0
ICX-Tx-Core2(config-loopback-if)# ip address 10.x.0.3/32
ICX-Tx-Core2(config-loopback-if)# exit
```

```
ICX-Tx-Core2(config)# router ospf 1
ICX-Tx-Core2(config-ospf-1)# router-id 10.x.0.3
OSPFv2 protocol will be reset.
Do you want to continue (y/n)? y
ICX-Tx-Core2(config-ospf-1)# exit
```

NOTE: Changing the OSPF router-id requires the OSPF process to restart. Make sure to confirm the restart with 'y'.

18. On Core2, check the current configuration of the Loopback 0 interface.

```
ICX-Tx-Core2(config)# show run interface loopback 0
```

Q: Is OSPF enabled on the interface? If so, why?

_____

A: Yes, OSPF is enabled. VSX OSPF configuration sync took care of this.

Core1

19. On Core1, review the OSPF LSDB.

```
ICX-Tx-Core1(config)# show ip ospf lsdb
OSPF Router with ID (10.x.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.x.0.2       10.x.0.2       601       0x80000005 0x00000c55     2
10.x.0.3       10.x.0.3       602       0x80000002 0x0000263a     2

Network Link State Advertisements (Area 0.0.0.0)
------------------------------------------------

LSID            ADV Router      Age       Seq#       Checksum
-------------------------------------------------------------
10.x.2.2       10.x.0.2       606       0x80000001 0x0000688f
```

Q: What are the 2 LSA types in this LSDB?

_____

A: Router (type1) and Network (type2) .

Q: What does the network LSA represent?

_____

A: A network LSA is representing a multi-access network between two or more OSPF routers. Since the VLAN 2 Ethernet interface is a type 'broadcast' interface, by default OSPF will start the DR/BDR election process and the DR will generate the Network LSA.

20. Review the current OSPF neighbor list.

```
ICX-Tx-Core1(config)# show ip ospf neighbors
OSPF Process ID 1 VRF default
=============================

Total Number of Neighbors: 1
```

```
Neighbor ID      Priority  State             Nbr Address       Interface
-------------------------------------------------------------------------
10.x.0.3         1         FULL/BDR          10.x.2.3          vlan2

ICX-Tx-Core1(config)#
```

Q: What does the DR/BDR stand for?

_____

A: Designated Router and Backup Designated Router are elected on a multi-access network, such as an Ethernet broadcast network. This is the default operation of OSPF.

OSPF Point-to-Point Connection

A network type broadcast means that more than two devices could come online on the subnet, so this is why OSPF prepares itself and elects the DR/BDR for the subnet. This election takes some time during the link setup phase.

When the design ensures that only two OSPF routers will ever be online on a subnet that is transported on a broadcast network, the administrator can change the OSPF network type to 'point to point'.

For this reason, even when the physical network is still broadcast, such as an ethernet VLAN, this configuration will instruct OSPF to assume that only one peer is available on this subnet. This will effectively skip the DR/BDR election process. This will also eliminate the type 2 Network LSA from the LSDB.

It is best practice to set the OSPF network type to point to point for point to point links.

Core1

21. On the Core1, change the VLAN 2 interface OSPF network type to 'point to point'.

```
ICX-Tx-Core1(config)# interface vlan 2
ICX-Tx-Core1(config-if-vlan)# ip ospf network point-to-point
ICX-Tx-Core1(config-if-vlan)# exit
```

22. Review the configuration changes that were pushed to Core2 by VSX sync.

```
ICX-Tx-Core1(config)# show run int vlan2 vsx-peer
interface vlan2
    ip address 10.x.2.3/31
    ip ospf 1 area 0.0.0.0
    ip ospf cost 1
    ip ospf network point-to-point
    exit
```

23. On Core1, review the OSPF LSDB. The network LSA that was previously present for the broadcast transit network has now been removed.

```
ICX-Tx-Core1(config)# show ip ospf lsdb
OSPF Router with ID (10.x.0.2) (Process ID 1 VRF default)
============================================================

Router Link State Advertisements (Area 0.0.0.0)
-----------------------------------------------

LSID            ADV Router      Age       Seq#       Checksum
-----------------------------------------------------------------
10.x.0.2       10.x.0.2       130       0x80000004 0x00007163
10.x.0.3       10.x.0.3       131       0x80000004 0x0000874a

ICX-Tx-Core1(config)#
```

Route verification

24. Verify there is an OSPF route in the routing table for the peer Core switch loopback interface using the 'show ip route ospf' command.

> **NOTE:** The 'show ip route' command has many filtering options. In real world environments with many routes, these filter options can help the administrator to narrow down the output of the command.

```
ICX-Tx-Core1(config)# show ip route ospf

Displaying ipv4 routes selected for forwarding

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

10.x.0.3/32, vrf default
        via  10.x.2.3,  [110/1],  ospf

ICX-Tx-Core1(config)#
```

25. Verify connectivity between the loopback interfaces with a ping from the local Loopback IP address to the peer switch Loopback IP address.

```
ICX-Tx-Core1(config)# do ping 10.x.0.3 source loopback0
PING 10.x.0.3 (10.x.0.3) from 10.x.0.2 : 100(128) bytes of data.
108 bytes from 10.x.0.3: icmp_seq=1 ttl=64 time=0.179 ms
108 bytes from 10.x.0.3: icmp_seq=2 ttl=64 time=0.174 ms
108 bytes from 10.x.0.3: icmp_seq=3 ttl=64 time=0.186 ms
108 bytes from 10.x.0.3: icmp_seq=4 ttl=64 time=0.205 ms
108 bytes from 10.x.0.3: icmp_seq=5 ttl=64 time=0.248 ms
```

```
--- 10.x.0.3 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4086ms
rtt min/avg/max/mdev = 0.174/0.198/0.248/0.029 ms
ICX-Tx-Core1(config)#
```

Q: Why is the loopback IP address configured as the source IP for this test?

_____

A: By default, a router will use the IP address of the outbound interface as the source IP. In this case, that would have been the VLAN 1 interface. If the administrator wants to test if the remote router has an entry in the routing table for the local loopback, the loopback can be set as the source IP.

26. Review the OSPF neighbors.

```
ICX-Tx-Core1(config)# show ip ospf neighbors
OSPF Process ID 1 VRF default
==============================

Total Number of Neighbors: 1

Neighbor ID      Priority  State              Nbr Address       Interface
-------------------------------------------------------------------------
10.x.0.3         n/a       FULL               10.x.2.3          vlan2

ICX-Tx-Core1(config)#
```

Review the detailed OSPF neighbor state.

```
ICX-Tx-Core1(config)# show ip ospf neighbors 10.x.0.3
Neighbor 10.x.0.3, interface address 10.x.2.3
-----------------------------------------------

Process ID 1 VRF default, in area 0.0.0.0 via interface vlan2
Neighbor priority is n/a, State is FULL
Options is 0x42
Dead timer due in 00:00:35
Time since last state change 00h:07m:42s
ICX-Tx-Core1(config)#
```

27. Review the OSPF interfaces, note the 'State' of each interface.

```
ICX-Tx-Core1(config)# show ip ospf interface
Interface loopback0 is up, line protocol is up
-----------------------------------------------

IP address 10.x.0.2/32, Process ID 1 VRF default, area 0.0.0.0
    State Loopback, Status up, Network type Loopback
    Link Speed: NA
```

```
    Cost Configured NA, Calculated NA
    Transit delay 1 sec, Router priority 1
    No designated router on this network
    No backup designated router on this network
    Timer Intervals: Hello 10, Dead 40, Retransmit 5
    No authentication
    Number of Link LSAs: 0, checksum sum 0
    BFD is disabled

Interface vlan2 is up, line protocol is up
------------------------------------------

IP address 10.x.2.2/31, Process ID 1 VRF default, area 0.0.0.0
    State Point-to-point, Status up, Network type Point-to-point
    Link Speed: 1000 Mbps
    Cost Configured 1, Calculated 1
    Transit delay 1 sec, Router priority n/a
    No designated router on this network
    No backup designated router on this network
    Timer Intervals: Hello 10, Dead 40, Retransmit 5
    No authentication
    Number of Link LSAs: 0, checksum sum 0
    BFD is disabled
```

Task 3: OSPF Address Advertisements and Control

Diagram L3

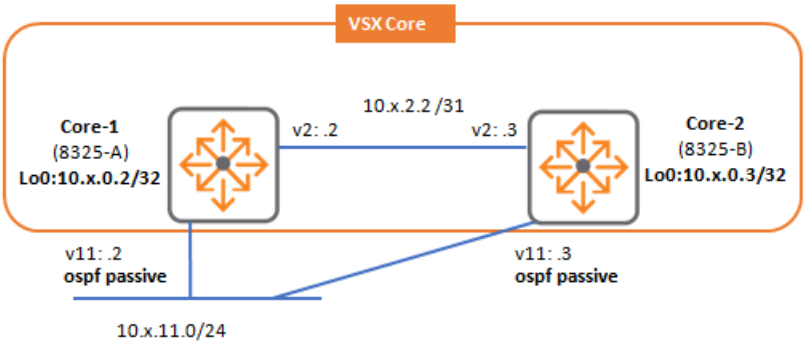

Objectives

- Enable end-host subnet advertisement into OSPF
- Configure a passive interface
- Use the passive interface default
- Review the administrative distance of OSPF

Steps

Core1

1. On Core1, enable the user VLAN 11 for OSPF in area 0.

```
ICX-Tx-Core1(config)# interface vlan11
ICX-Tx-Core1(config-if-vlan)# ip ospf 1 area 0
ICX-Tx-Core1(config-if-vlan)# exit
```

2. Review the running configuration of interface VLAN 11 on both the local device and the VSX peer (Core2).

```
ICX-Tx-Core1(config)# show running interface vlan11
interface vlan11
    vsx-sync active-gateways policies
    ip address 10.x.11.2/24
    active-gateway ip mac 02:02:00:00:12:00
    active-gateway ip 10.x.11.1
    ip helper-address 10.x.1.6
    l3-counters
    ip ospf 1 area 0.0.0.0
    exit
```

```
ICX-Tx-Core1(config)# show running interface vlan11 vsx-peer
interface vlan11
```

```
    vsx-sync active-gateways policies
    ip address 10.x.11.3/24
    active-gateway ip mac 02:02:00:00:12:00
    active-gateway ip 10.x.11.1
    ip helper-address 10.x.1.6
    l3-counters
    ip ospf 1 area 0.0.0.0
    exit
ICX-Tx-Core1(config)#
```

NOTE: For the VSX OSPF synchronization to work properly, the administrator should use the same physical uplink interface. For instance, if Core1 has a routed uplink port 1/1/45, the Core2 should also use interface 1/1/45 as the routed uplink.

3. Review the OSPF neighbors.

```
ICX-Tx-Core1(config)# show ip ospf neighbors
OSPF Process ID 1 VRF default
=============================

Total Number of Neighbors: 3

Neighbor ID      Priority  State             Nbr Address       Interface
-------------------------------------------------------------------------
10.x.0.3         1         FULL/DR           10.x.11.3         vlan11
10.x.0.3         n/a       FULL              10.x.2.3          vlan2
ICX-Tx-Core1(config)#
```

Q: Why did OSPF establish a neighbor relationship over the user vlan?

_____

A: When an interface is enabled for OSPF, it automatically starts sending hello packets.

Q: Is it desired to have OSPF communication over user subnets?

_____

A: Typically, OSPF communication is not enabled on user subnets to protect the OSPF protocol against malicious neighbors that may be connected to the user subnets.

Passive Interfaces

Passive interfaces provide a method to have an interface enabled for OSPF, so the subnet will be included in the router LSA announcement. But at the same time, the OSPF protocol will be disabled on this interface, so no OSPF hello packets will be transmitted or received by this interface.

4. Configure the VLAN 11 interface as passive/silent interfaces.

```
ICX-Tx-Core1(config)# interface vlan 11
ICX-Tx-Core1(config-if-vlan)# ip ospf passive
ICX-Tx-Core1(config-if-vlan)# exit
ICX-Tx-Core1(config)#
```

5. Verify that only the VLAN2 transit link is now used for the OSPF peering.

```
ICX-Tx-Core1(config)# show ip ospf neighbors
OSPF Process ID 1 VRF default
=============================

Total Number of Neighbors: 1

Neighbor ID      Priority  State             Nbr Address       Interface
-------------------------------------------------------------------------
10.x.0.3        n/a       FULL              10.x.2.3          vlan2
```

Passive Interfaces default

If a network has many user subnets, it may be more convenient to set new OSPF interfaces to passive by default.

This means that only the OSPF uplink interfaces must be configured as 'not passive'.

The other advantage is that there is no security risk when adding a new user VLAN, since it will automatically have OSPF disabled.

In the lab setup, the VLANs 1,11,12 can have endpoints connected to them, so they can benefit from the 'passive default' feature.

The VLAN 2 interface is a dedicated routed link, so the passive feature should be turned off for this interface.

6. On Core1, enter the OSPF configuration context.

```
ICX-Tx-Core1(config)# router ospf 1
ICX-Tx-Core1(config-ospf-1)# passive-interface default
ICX-Tx-Core1(config-ospf-1)# exit
```

IMPORTANT: This change will impact the current OSPF adjacencies, so this change should be made during a maintenance window.

7. Enable the VLAN 2 interface again using 'no ip ospf passive'.

```
ICX-Tx-Core1(config)# interface vlan 2
ICX-Tx-Core1(config-if-vlan)# no ip ospf passive
ICX-Tx-Core1(config-if-vlan)# exit
```

8. Add VLAN interfaces 1 and 12 as OSPF enabled interfaces.

```
ICX-Tx-Core1(config)# interface vlan 1,12
ICX-Tx-Core1(config-if-vlan-<1,12>)# ip ospf 1 area 0.0.0.0
ICX-Tx-Core1(config-if-vlan-<1,12>)# exit
```

9. Verify that the new interfaces vlan1 and 12 are now enabled for OSPF.

```
ICX-Tx-Core1(config)# show ip ospf interface brief
OSPF Process ID 1 VRF default
==============================

Total Number of Interfaces: 5

Interface     Area             IP Address/Mask     Cost    State           Status
-------------------------------------------------------------------------------
-
loopback0     0.0.0.0          10.x.0.2/32         0       Loopback        up
vlan1         0.0.0.0          10.x.1.2/24         100     Dr-other        up
vlan11        0.0.0.0          10.x.11.2/24        100     Dr-other        up
vlan12        0.0.0.0          10.x.12.2/24        100     Dr-other        up
vlan2         0.0.0.0          10.x.2.2/31         1       Point-to-point  up
```

10. But the only OSPF adjacency was formed over VLAN 2.

```
ICX-Tx-Core1(config)# show ip ospf neighbors
OSPF Process ID 1 VRF default
==============================

Total Number of Neighbors: 1

Neighbor ID      Priority  State             Nbr Address       Interface
-------------------------------------------------------------------------
10.x.0.3        n/a       FULL              10.x.2.3          vlan2
```

11. Review the interface statistics of VLAN 11 to confirm that no OSPF packets have been transmitted.

```
ICX-Tx-Core1(config)# show ip ospf statistics interface vlan11 | include Tx
Tx Hello Packets      : 0           Rx Hello Packets      : 0
Tx Hello Bytes        : 0           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
```

Review Administrative Distance

12. First, review the routes learned and calculated by OSPF through the LSDB. These routes will be presented by OSPF to the IP routing table. Notice that routes that are local on Core1 have been learned as 'i' (intra-area) routes based on the Router LSA of Core1.

```
ICX-Tx-Core1(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 : 5

10.x.0.3/32       (i) area: 0.0.0.0
     via 10.x.2.3 interface vlan2, cost 1 distance 110
10.x.1.0/24       (i) area: 0.0.0.0
     directly attached to interface vlan1, cost 100 distance 110
10.x.2.2/31       (i) area: 0.0.0.0
     directly attached to interface vlan2, cost 1 distance 110

<...output omitted...>
```

Q: What is the distance applied to the route 10.x.1.0/24?

A: OSPF routes get an administrative distance of 110 by default.

13. Now review the IP routing table.

```
ICX-Tx-Core1(config)# show ip route

Displaying ipv4 routes selected for forwarding

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

10.x.0.2/32, vrf default
        via  loopback0,  [0/0],  local
10.x.0.3/32, vrf default
        via  10.x.2.3,  [110/1],  ospf
10.x.1.0/24, vrf default
        via  vlan1,  [0/0],  connected
10.x.1.2/32, vrf default
        via  vlan1,  [0/0],  local
```

```
<...output omitted...>
```

Q: The output shows that the 10.x.1.0/24 route was not learned by OSPF. What is the source of the 10.x.1.0/24 route?

_____

A: The source is 'connected', this indicates an IP subnet that is locally connected on the router.

Q: Why was this route selected by the routing table over the OSPF route?

_____

A: The administrative distance of directly connected routes is 0, so these routes will always win from any routing protocol. OSPF has a default distance of 110.

Task 4: OSPF Peering Using VSX LAG

Diagram L3

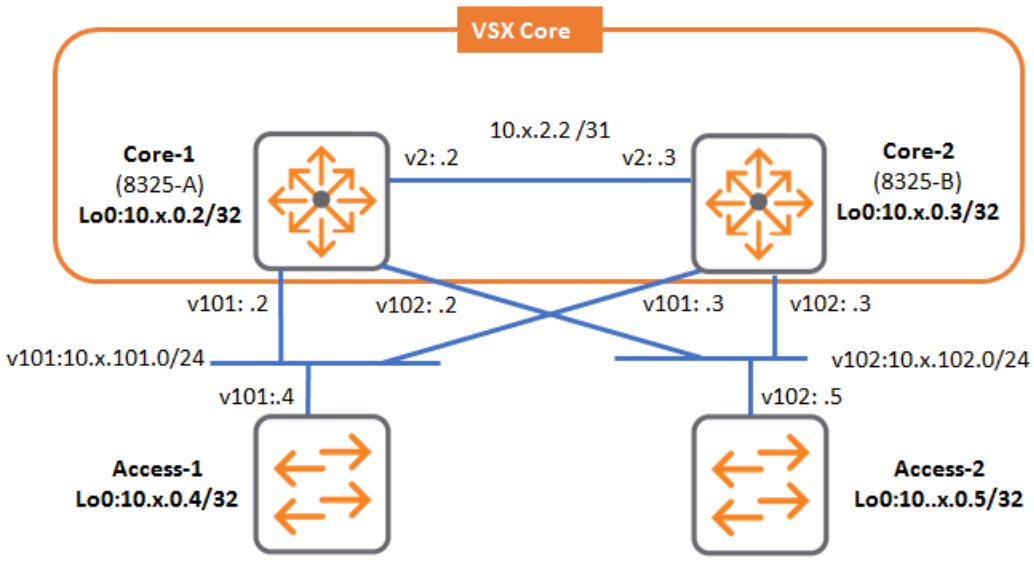

Objectives

- Configure VSX LAG with OSPF to neighbor routers
- Understand and configure the Active Forwarding feature

Steps

Define a new routed VLAN between Access and the VSX Core.

Enable OSPF between Access1 and the Core VSX.

Routed VLAN between Access and Core VSX

Core1

1. On Core1, define VLAN 101 (to Access1) and 102 (to Access2).

```
ICX-Tx-Core1(config)# vlan 101
ICX-Tx-Core1(config-vlan-101)# vsx-sync
ICX-Tx-Core1(config-vlan-101)# exit

ICX-Tx-Core1(config)# vlan 102
ICX-Tx-Core1(config-vlan-102)# vsx-sync
ICX-Tx-Core1(config-vlan-102)# exit
ICX-Tx-Core1(config)#
```

```
ICX-Tx-Core1(config)# interface lag 1
ICX-Tx-Core1(config-lag-if)# vlan trunk allowed 101
ICX-Tx-Core1(config-lag-if)# exit

ICX-Tx-Core1(config)# interface lag 2
ICX-Tx-Core1(config-lag-if)# vlan trunk allowed 102
ICX-Tx-Core1(config-lag-if)# exit
```

2. For VLAN 101, configure the IP address and enable OSPF.

```
ICX-Tx-Core1(config)# interface vlan 101
ICX-Tx-Core1(config-if-vlan)# description L3-access1
ICX-Tx-Core1(config-if-vlan)# ip address 10.x.101.2/24
ICX-Tx-Core1(config-if-vlan)# ip ospf 1 area 0
ICX-Tx-Core1(config-if-vlan)# exit
```

3. VLAN 102 is only prepared with an IP address at this point, it is **not** enabled for OSPF at this point, this will be done in the next Lab.

```
ICX-Tx-Core1(config)# interface vlan 102
ICX-Tx-Core1(config-if-vlan)# description L3-Access2
ICX-Tx-Core1(config-if-vlan)# ip address 10.x.102.2/24
ICX-Tx-Core1(config-if-vlan)# exit
```

Core2

4. On Core2, define VLAN 101 and 102 IP addresses (L2 VLAN and OSFP are automatically enabled by vsx-sync).

```
ICX-Tx-Core2(config)# int vlan 101
ICX-Tx-Core2(config-if-vlan)# ip address 10.x.101.3/24
ICX-Tx-Core2(config-if-vlan)# description L3-access1
ICX-Tx-Core2(config-if-vlan)# exit
```

```
ICX-Tx-Core2(config)# int vlan 102
ICX-Tx-Core2(config-if-vlan)# ip address 10.x.102.3/24
ICX-Tx-Core2(config-if-vlan)# description L3-access2
ICX-Tx-Core2(config-if-vlan)# exit
```

Access1

On Access1, define the VLAN 101 uplink interface.

5. Access a terminal session on Access1, enter the configuration mode.
6. Add VLAN 101, define the IP interface.

> **NOTE:** In the VSX lab activity, the LAG255 was set to allow all vlans, this is why there is no step to allow the VLAN 101 on the VLAN trunk.

```
ICX-Tx-Access1(config)# vlan 101
ICX-Tx-Access1(config-vlan-101)# exit
```

```
ICX-Tx-Access1(config)# interface vlan 101
ICX-Tx-Access1(config-if-vlan)# description L3-core
ICX-Tx-Access1(config-if-vlan)# ip address 10.x.101.4/24
ICX-Tx-Access1(config-if-vlan)# exit
```

7. Verify ping works between the Access1 and Core1/Core2.

```
ICX-Tx-Access1(config)# do ping 10.X.101.2
PING 10.x.101.2 (10.x.101.2) 100(128) bytes of data.
108 bytes from 10.x.101.2: icmp_seq=1 ttl=64 time=16.7 ms
...

ICX-Tx-Access1(config)# do ping 10.X.101.3
PING 10.x.101.3 (10.x.101.3) 100(128) bytes of data.
108 bytes from 10.x.101.3: icmp_seq=1 ttl=64 time=15.4 ms
...
ICX-Tx-Access1(config)#
```

8. Configure OSPF on Access1.

   - OSPF router id, area 0

   - Loopback interface with IP 10.x.0.4/32 and enable for OSPF

   - Enable VLAN 101 for OSPF

```
ICX-Tx-Access1(config)# router ospf 1
ICX-Tx-Access1(config-ospf-1)# router-id 10.x.0.4
ICX-Tx-Access1(config-ospf-1)# enable
ICX-Tx-Access1(config-ospf-1)# area 0
ICX-Tx-Access1(config-ospf-1)# exit
```

```
ICX-Tx-Access1(config)# interface loopback 0
ICX-Tx-Access1(config-loopback-if)# ip address 10.x.0.4/32
ICX-Tx-Access1(config-loopback-if)# ip ospf 1 area 0
ICX-Tx-Access1(config-loopback-if)# exit
```

```
ICX-Tx-Access1(config)# interface vlan 101
ICX-Tx-Access1(config-if-vlan)# ip ospf 1 area 0
ICX-Tx-Access1(config-if-vlan)# exit
ICX-Tx-Access1(config)#
```

9. Verify OSPF neighbors, the expectation would be to see Core1 and Core2 in FULL state.

```
ICX-Tx-Access1(config)# show ip ospf neighbors
No OSPF neighbor found on VRF default.
```

Q: What could be wrong with the configuration?

A: Try to troubleshoot this using these commands on the Access1 and Core1.

```
show ip ospf neighbors
show ip ospf interface
show ip ospf statistics interface vlan101
```

Example output for the statistics:

```
ICX-Tx-Access1(config-ospf-1)# show ip ospf statistics interface vlan101
OSPF Process ID 1 VRF default, interface vlan101 statistics (cleared 0h4m6s ago)
================================================================================

Tx Hello Packets      : 25          Rx Hello Packets      : 0
Tx Hello Bytes        : 1600        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
```

The troubleshooting commands should have shown that there are no Receive statistics for the OSPF interface VLAN 101. This is because the passive-interface default was configured on the VSX Core side, so any new interface on the VSX Core are passive by default.

Core1

10. On the Core1, enable the new routed VLAN interface for OSPF.

```
ICX-Tx-Core1(config)# interface vlan 101
ICX-Tx-Core1(config-if-vlan)# no ip ospf passive
ICX-Tx-Core1(config-if-vlan)# exit
```

Access1

11. On Access1, check the OSPF neighbors again.

```
ICX-Tx-Access1(config)# show ip ospf neighbors
OSPF Process ID 1 VRF default
==============================

Total Number of Neighbors: 2

Neighbor ID      Priority  State             Nbr Address       Interface
--------------------------------------------------------------------------
```

```
10.x.0.2          1           FULL/DR           10.x.101.2        vlan101

10.x.0.3          1           FULL/BDR          10.x.101.3        vlan101

ICX-Tx-Access1(config)#
```

NOTE: The Role (DR / BDR / DROther) may be different in your output, but the state should be FULL.

You have completed Lab 6.1!

Lab 06.2: OSPF Multi-Area- Configuring OSPF with Multiple Areas

Lab Diagram

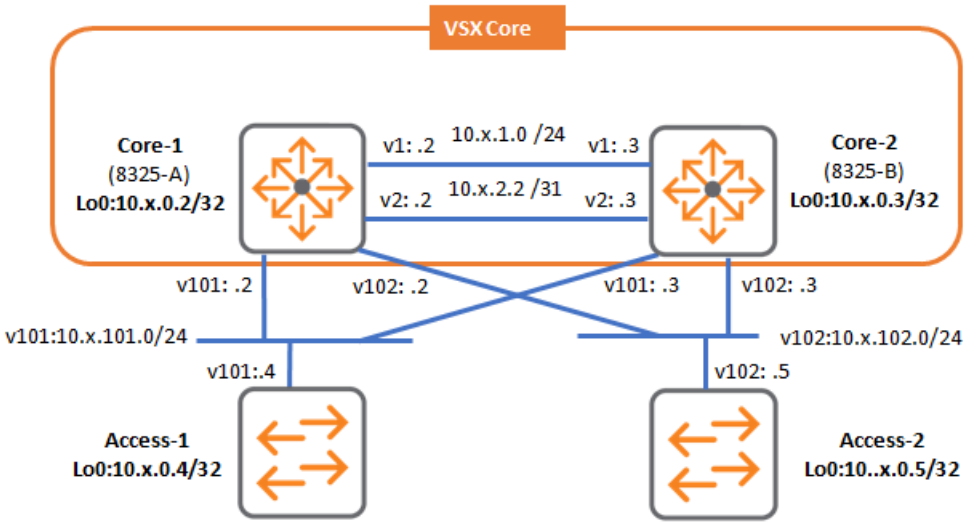

Overview

This lab activity will introduce multiple areas for OSPF. The lab will demonstrate the benefits of using multiple areas, such as route aggregation and filtering.

Objectives

- Learn how to configure an OSPF Area Border Router (ABR)
- Understand how routes from 1 area are announced into other areas
- Understand the need for the backbone area
- Apply route aggregation and route filtering at the ABR

Task 1: Assign Access1 to OSPF Area 1

Diagram

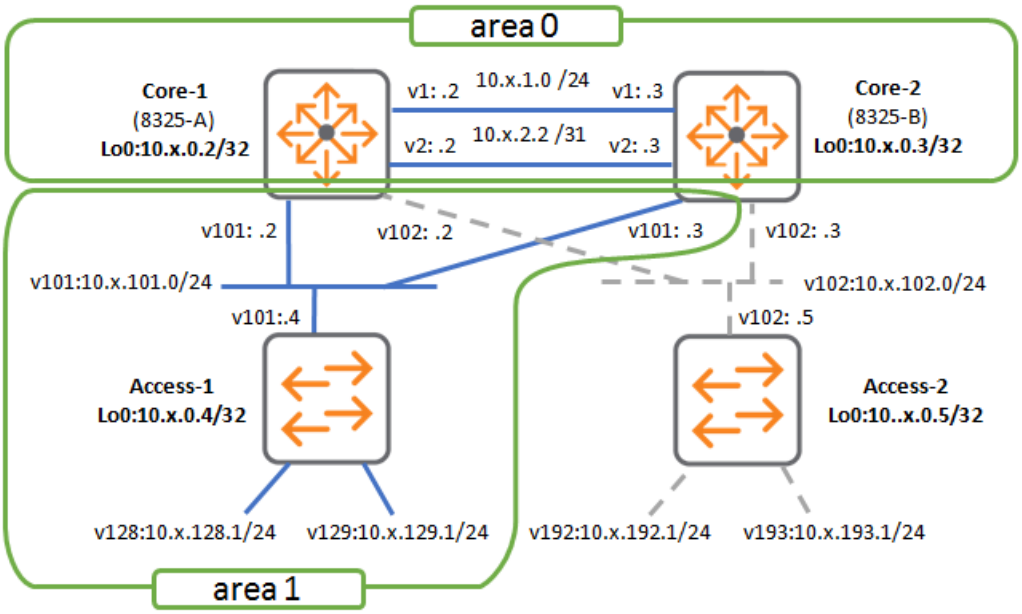

Objectives

- Setup ABR link between Core1 and Access1
- Verify LSA changes
- Setup ABR link between Core2 and Access1

Steps

On Core1, define area 1 and assign interface VLAN 101 to area 1.

On Access1, remove area 0, define area 1. Assign interface VLAN 101 to area 1.

Verify adjacency and operation.

Verify area 1 setup between Access1 and Core2 and verify operation.

Core1

1. Open a terminal connection to Core1, enter the configuration mode.
2. Under the router ospf context, define a new area with id 1.

```
ICX-Tx-Core1(config)# router ospf 1
ICX-Tx-Core1(config-ospf-1)# area 1
ICX-Tx-Core1(config-ospf-1)# exit
```

3. Under the VLAN 101 context, assign area 1 for this interface.

```
ICX-Tx-Core1(config)# interface vlan 101
ICX-Tx-Core1(config-if-vlan)# ip ospf 1 area 1
```

4. Verify the current configuration of this interface.

```
ICX-Tx-Core1(config-if-vlan)# show running current
interface vlan101
    description L3-ccess1
    ip address 10.x.101.2/24
    ip ospf 1 area 0.0.0.1
    no ip ospf passive
ICX-Tx-Core1(config-if-vlan)# exit
```

Verify the LSDB Changes on Core1

In each area database, Core1 has inserted a router LSA that contains the links for that area. You will see in the area 1 database, there will be a router LSA with one link (the VLAN 101 interface). Due to the VSX synchronization, Core2 has also inserted a router LSA in this area 1.

5. Review the LSDB for area 1.

```
ICX-Tx-Core1(config)# show ip ospf lsdb area 1
OSPF Router with ID (10.x.0.2) (Process ID 1 VRF default)
============================================================

Router Link State Advertisements (Area 0.0.0.1)
-----------------------------------------------

LSID            ADV Router      Age       Seq#       Checksum       Link Count
-------------------------------------------------------------------------------
10.x.0.2       10.x.0.2       18        0x80000002 0x0000744d     1
10.x.0.3       10.x.0.3       19        0x80000002 0x0000724c     1

Network Link State Advertisements (Area 0.0.0.1)
------------------------------------------------

LSID            ADV Router      Age       Seq#       Checksum
-------------------------------------------------------------
10.x.101.3     10.x.0.3       24        0x80000001 0x00000c85

Inter-area Summary Link State Advertisements (Area 0.0.0.1)
-----------------------------------------------------------

LSID            ADV Router      Age       Seq#       Checksum
-------------------------------------------------------------
10.x.0.2       10.x.0.2       63        0x80000001 0x00007daf
10.x.0.2       10.x.0.3       63        0x80000001 0x000081a9
10.x.0.3       10.x.0.2       63        0x80000001 0x00007dad
10.x.0.3       10.x.0.3       63        0x80000001 0x00006dbd
10.x.1.0       10.x.0.2       63        0x80000001 0x00007257
10.x.1.0       10.x.0.3       63        0x80000001 0x00006c5c
10.x.2.2       10.x.0.2       63        0x80000001 0x00006bbf
```

```
10.x.2.2        10.x.0.3        63         0x80000001 0x000065c4
10.x.11.0       10.x.0.2        63         0x80000001 0x000004bb
10.x.11.0       10.x.0.3        63         0x80000001 0x0000fdc0
10.x.12.0       10.x.0.2        63         0x80000001 0x0000f8c5
10.x.12.0       10.x.0.3        63         0x80000001 0x0000f2ca
```

Q: How many links are there in the router LSA for Core1 inside area 1 (10.x.0.2)?

_____

A: The router LSA inside the area 1 only contains the VLAN101 interface, so only one link will be present in the router LSA at this point.

The ABR is now responsible to convert the routes that were learned through the Router (type1) and Network (type2) LSAs between the areas. This means that the VLAN 2,11,12 IP prefixes that are known on the Router LSA in area 0, have been converted into a 'Summary' LSA (type3) in area 1.

6. In the same output, verify that the IP prefixes of area 0 are known as inter-area summary (type3) LSAs in area 1.

> **NOTE:** 'Summary' does not indicate route summarization, it only indicates that this LSA is not used for the topology calculation of this area. Only Router and Network LSAs are used to build the topology of an area.

```
ICX-Tx-Core1(config-ospf-1)# show ip ospf lsdb area 1
OSPF Router with ID (10.x.0.2) (Process ID 1 VRF default)
============================================================

Router Link State Advertisements (Area 0.0.0.1)
-----------------------------------------------

LSID            ADV Router      Age       Seq#       Checksum       Link Count
-------------------------------------------------------------------------------
10.x.0.2       10.x.0.2       18        0x80000002 0x0000744d     1
10.x.0.3       10.x.0.3       19        0x80000002 0x0000724c     1

Network Link State Advertisements (Area 0.0.0.1)
------------------------------------------------

LSID            ADV Router      Age       Seq#       Checksum
-------------------------------------------------------------
10.x.101.3     10.x.0.3       24        0x80000001 0x00000c85

Inter-area Summary Link State Advertisements (Area 0.0.0.1)
-----------------------------------------------------------
```

```
LSID            ADV Router      Age       Seq#       Checksum
-------------------------------------------------------------
10.x.0.2        10.x.0.2        63        0x80000001 0x00007daf
10.x.0.2        10.x.0.3        63        0x80000001 0x000081a9
10.x.0.3        10.x.0.2        63        0x80000001 0x00007dad
10.x.0.3        10.x.0.3        63        0x80000001 0x00006dbd
10.x.1.0        10.x.0.2        63        0x80000001 0x00007257
10.x.1.0        10.x.0.3        63        0x80000001 0x00006c5c
10.x.2.2        10.x.0.2        63        0x80000001 0x00006bbf
10.x.2.2        10.x.0.3        63        0x80000001 0x000065c4
10.x.11.0       10.x.0.2        63        0x80000001 0x000004bb
10.x.11.0       10.x.0.3        63        0x80000001 0x0000fdc0
10.x.12.0       10.x.0.2        63        0x80000001 0x0000f8c5
10.x.12.0       10.x.0.3        63        0x80000001 0x0000f2ca
```

7. Review the LSDB for area 0.

```
ICX-Tx-Core1(config)# show ip ospf lsdb area 0
OSPF Router with ID (10.x.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.x.0.2        10.x.0.2        421       0x80000014 0x0000369c     6
10.x.0.3        10.x.0.3        422       0x80000010 0x0000785b     6
10.x.0.4        10.x.0.4        1069      0x80000003 0x0000eca3     2

Network Link State Advertisements (Area 0.0.0.0)
------------------------------------------------

LSID            ADV Router      Age       Seq#       Checksum
-------------------------------------------------------------
10.x.101.4      10.x.0.4        1069      0x80000002 0x0000a1ce

Inter-area Summary Link State Advertisements (Area 0.0.0.0)
-----------------------------------------------------------

LSID            ADV Router      Age       Seq#       Checksum
-------------------------------------------------------------
10.x.101.0      10.x.0.2        416       0x80000002 0x00002044
10.x.101.0      10.x.0.3        417       0x80000002 0x00001a49
```

NOTE: You may still notice the Access1 (10.x.0.4) Router LSA in the area 0 LSDB, even though the link to Access1 was already assigned to area 1.
This is because that Router LSA still needs to age-out from the LSDB, so it will disappear automatically after 3600 seconds.

Q: The interface VLAN 101 IP subnet is now active in area 1. Is it known as a 'summary' LSA in area 0?

A: Yes, the 10.x.101.0 is listed as a summary LSA.

Area Border routers can be easily found in the LSDB since they are advertising the 'Summary' LSAs. The 'ADV Router' indicates the advertising router ID.

Q: Which routers are currently ABR?

A: Both 10.x.0.2 and 10.x.0.3 are currently advertising a summary LSA.

Q: The configuration was only performed on Core1. How is it possible that Core2 (10.x.0.3) is also advertising the 10.x.101.0 Summary LSA?

A: Due to the VSX configuration synchronization, area 1 was also defined on Core2, and the interface VLAN 101 was also moved into this area 1. Therefore, without accessing the Core2 configuration, it also became an ABR.

On Access1, remove area 0, define area 1. Assign interface VLAN 101 to area 1

Access1

8. On Access1, enter the configuration mode.
9. Remove area 0 and define area 1.

```
ICX-Tx-Access1(config)# router ospf 1
ICX-Tx-Access1(config-ospf-1)# no area 0
ICX-Tx-Access1(config-ospf-1)# area 1
ICX-Tx-Access1(config-ospf-1)# exit
```

10. Assign interface VLAN 101 to area 1.

```
ICX-Tx-Access1(config)# interface vlan 101
ICX-Tx-Access1(config-if-vlan)# ip ospf 1 area 1
ICX-Tx-Access1(config-if-vlan)# exit
```

11. On Access1, define some local IP subnets. These subnets simulate client subnets and will be used to test routing.

The Access1 router will be the only host in these test subnets.

The test subnets for area 1 are assigned in the 10.x.128-191 block, so it will be possible to summarize these in a later step.

12. Define the test VLANs. Enable them as VLAN trunk on the port to the client PC. This is only done to ensure the VLAN is 'up'.

NOTE: If a VLAN does not contain any port in the 'up' state, the Layer3 IP interface remains down as well. As a workaround to bring these VLANs 'UP' , these 2 VLANs are enabled on the port to the PC.

```
ICX-Tx-Access1(config)# vlan 128,129
ICX-Tx-Access1(config-vlan-<128,129>)# exit

ICX-Tx-Access1(config)# interface 1/1/3
ICX-Tx-Access1(config-if)# no shutdown
ICX-Tx-Access1(config-if)# vlan trunk allowed 128,129
ICX-Tx-Access1(config-if)# exit
ICX-Tx-Access1(config)#
```

13. Define the test subnets.

```
ICX-Tx-Access1(config)# interface vlan 128
ICX-Tx-Access1(config-if-vlan)# ip address 10.x.128.1/24
ICX-Tx-Access1(config-if-vlan)# ip ospf 1 area 1
ICX-Tx-Access1(config-if-vlan)# exit

ICX-Tx-Access1(config)# interface vlan 129
ICX-Tx-Access1(config-if-vlan)# ip address 10.x.129.1/24
ICX-Tx-Access1(config-if-vlan)# ip ospf 1 area 1
ICX-Tx-Access1(config-if-vlan)# exit
ICX-Tx-Access1(config)#
```

14. On Access1, review the LSDB for area 1, note the link count for the Router LSA of Access1.

```
ICX-Tx-Access1(config)# show ip ospf lsdb area 1
OSPF Router with ID (10.x.0.4) (Process ID 1 VRF default)
============================================================

Router Link State Advertisements (Area 0.0.0.1)
-----------------------------------------------
```

```
LSID            ADV Router      Age       Seq#       Checksum       Link Count
-------------------------------------------------------------------------------
10.x.0.2        10.x.0.2        757       0x80000002 0x0000744d     1
10.x.0.3        10.x.0.3        132       0x80000003 0x0000704d     1
10.x.0.4        10.x.0.4        11        0x80000005 0x00004b57     3

Network Link State Advertisements (Area 0.0.0.1)
------------------------------------------------

LSID            ADV Router      Age       Seq#       Checksum
-------------------------------------------------------------
10.x.101.3      10.x.0.3        132       0x80000002 0x0000adc4

Inter-area Summary Link State Advertisements (Area 0.0.0.1)
-----------------------------------------------------------

LSID            ADV Router      Age       Seq#       Checksum
-------------------------------------------------------------
10.x.0.2        10.x.0.2        802       0x80000001 0x00007daf
10.x.0.2        10.x.0.3        802       0x80000001 0x000081a9
10.x.0.3        10.x.0.2        802       0x80000001 0x00007dad
10.x.0.3        10.x.0.3        802       0x80000001 0x00006dbd
10.x.1.0        10.x.0.2        802       0x80000001 0x00007257
10.x.1.0        10.x.0.3        802       0x80000001 0x00006c5c
10.x.2.2        10.x.0.2        802       0x80000001 0x00006bbf
10.x.2.2        10.x.0.3        802       0x80000001 0x000065c4
10.x.11.0       10.x.0.2        802       0x80000001 0x000004bb
10.x.11.0       10.x.0.3        802       0x80000001 0x0000fdc0
10.x.12.0       10.x.0.2        802       0x80000001 0x0000f8c5
10.x.12.0       10.x.0.3        802       0x80000001 0x0000f2ca
```

Core1

15. On Core1, verify that the same information is known in the area 1 LSDB. So Core1 also knows about the three links of the Access1 router in the area 1 database.

> **NOTE:** Remember that an OSPF area LSDB contains the same information on all routers of that area.

```
ICX-Tx-Core1(config)# show ip ospf lsdb area 1
OSPF Router with ID (10.x.0.2) (Process ID 1 VRF default)
=========================================================

Router Link State Advertisements (Area 0.0.0.1)
-----------------------------------------------

LSID            ADV Router      Age       Seq#       Checksum       Link Count
-------------------------------------------------------------------------------
10.x.0.2        10.x.0.2        808       0x80000002 0x0000744d     1
10.x.0.3        10.x.0.3        184       0x80000003 0x0000704d     1
10.x.0.4        10.x.0.4        24        0x80000007 0x00004759     3
```

```
Network Link State Advertisements (Area 0.0.0.1)
-------------------------------------------------

LSID            ADV Router      Age       Seq#       Checksum
-------------------------------------------------------------
10.x.101.3     10.x.0.3       184       0x80000002 0x0000adc4

Inter-area Summary Link State Advertisements (Area 0.0.0.1)
-----------------------------------------------------------

LSID            ADV Router      Age       Seq#       Checksum
-------------------------------------------------------------
<…output omitted…>
```

16. The Core1 Area Border has converted the routes of this Router LSA from area 1 to Summary Routes into area 0. Verify this by reviewing the area 0 database.

```
ICX-Tx-Core1(config)# show ip ospf lsdb area 0
OSPF Router with ID (10.x.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.x.0.2       10.x.0.2       596       0x80000005 0x00000332     6
10.x.0.3       10.x.0.3       151       0x8000000d 0x00006bbd     6

Network Link State Advertisements (Area 0.0.0.0)
-------------------------------------------------

LSID            ADV Router      Age       Seq#       Checksum
-------------------------------------------------------------
10.x.102.3     10.x.0.3       600       0x80000002 0x0000fe90

Inter-area Summary Link State Advertisements (Area 0.0.0.0)
-----------------------------------------------------------

LSID            ADV Router      Age       Seq#       Checksum
-------------------------------------------------------------
10.x.101.0     10.x.0.2       641       0x80000002 0x00002044
10.x.101.0     10.x.0.3       150       0x80000003 0x0000184a
10.x.128.0     10.x.0.2       339       0x80000001 0x0000e302
10.x.128.0     10.x.0.3       340       0x80000001 0x0000dd07
10.x.129.0     10.x.0.2       329       0x80000001 0x0000d80c
10.x.129.0     10.x.0.3       329       0x80000001 0x0000d211
```

NOTE: The Core2 ABR has performed the same action, this is why the 10.x.128.0

and 10.x.129.0 Summary LSAs are listed two times.

17. On Core1, review learned OSPF routes and notice the OSPF route code.

```
ICX-Tx-Core1(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 : 8

10.x.0.3/32       (i) area: 0.0.0.0
     via 10.x.2.3 interface vlan2, cost 1 distance 110
10.x.1.0/24       (i) area: 0.0.0.0
     directly attached to interface vlan1, cost 100 distance 110
10.x.2.2/31       (i) area: 0.0.0.0
     directly attached to interface vlan2, cost 1 distance 110
10.x.11.0/24      (i) area: 0.0.0.0
     directly attached to interface vlan11, cost 100 distance 110
10.x.12.0/24      (i) area: 0.0.0.0
     directly attached to interface vlan12, cost 100 distance 110
10.x.101.0/24     (i) area: 0.0.0.1
     directly attached to interface vlan101, cost 100 distance 110
10.x.128.0/24     (i) area: 0.0.0.1
     via 10.x.101.4 interface vlan101, cost 200 distance 110
10.x.129.0/24     (i) area: 0.0.0.1
     via 10.x.101.4 interface vlan101, cost 200 distance 110
```

Q: Why are all OSPF routes on Core1 listed as (i) > intra-area route?

_____

A: In the current state of the lab, Core1 has direct access to both area 0 and area 1, so from the Core1 perspective, all learned routes are intra-area routes.

18. On Core1, review the IP routing table. The 10.x.128.0/24 and 10.x.129.0/24 subnets should be visible.

```
ICX-Tx-Core1(config)# show ip route ospf

Displaying ipv4 routes selected for forwarding

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

10.x.0.3/32, vrf default
        via  10.x.2.3,  [110/1],  ospf
10.x.128.0/24, vrf default
        via  10.x.101.4,  [110/200],  ospf
```

```
10.x.129.0/24, vrf default
        via  10.x.101.4,  [110/200],  ospf
```

This concludes the introduction of OSPF area 1 in the lab. Proceed to the next task.

Task 2: Assign Access2 to OSPF Area 2

Diagram

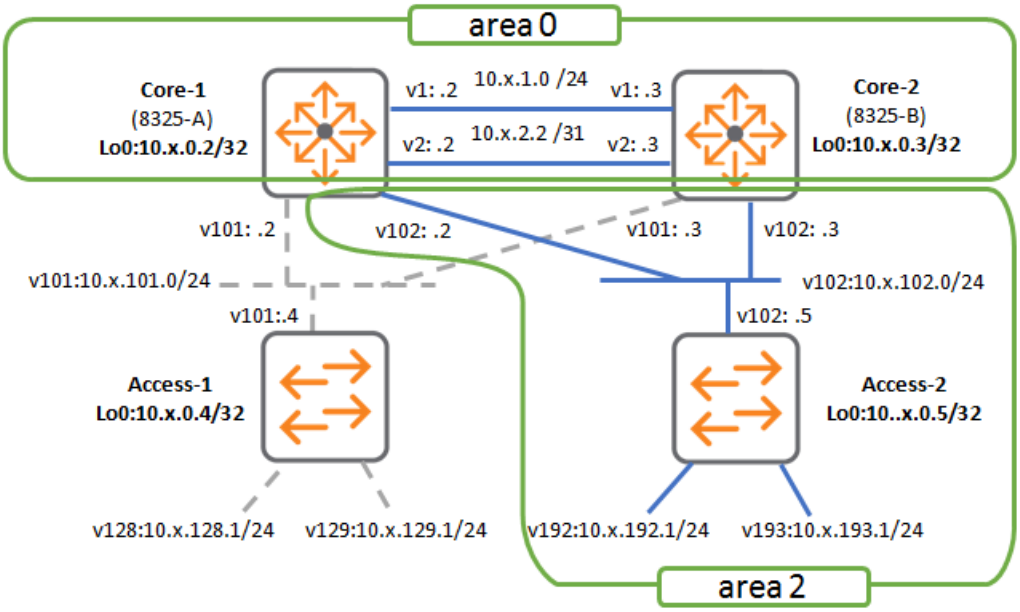

Objectives

- In this task, Access2 will be configured into its own area 2.
- Once area 2 has been defined, the LSA propagation will be reviewed.
- Setup ABR link
- Verify how the LSA changes when traversing through area 2, to area 0, then into area 1.

Steps

Core1

1. On Core1, define area 2, assign interface VLAN 102 to area 2.

```
ICX-Tx-Core1(config)# router ospf 1
ICX-Tx-Core1(config-ospf-1)# area 2
ICX-Tx-Core1(config-ospf-1)# exit

ICX-Tx-Core1(config)# interface vlan 102
ICX-Tx-Core1(config-if-vlan)# ip ospf 1 area 2
ICX-Tx-Core1(config-if-vlan)# no ip ospf passive
ICX-Tx-Core1(config-if-vlan)# exit
ICX-Tx-Core1(config)#
```

Access2

2. Open a terminal connection to Access2, enter the configuration mode.
3. On Access2, enable OSPF with the Access2 router-id and define area 2.

```
ICX-Tx-Access2(config)# router ospf 1
ICX-Tx-Access2(config-ospf-1)# router-id 10.x.0.5
ICX-Tx-Access2(config-ospf-1)# area 2
ICX-Tx-Access2(config-ospf-1)# exit
ICX-Tx-Access2(config)#
```

4. On Access2, define the VLAN 102.

> **NOTE:** In the VSX lab activity, the LAG255 was set to allow all vlans, this is why there is no step to allow the VLAN 101 on the VLAN trunk.

```
ICX-Tx-Access2(config)# vlan 102
ICX-Tx-Access2(config-vlan-102)# exit
```

5. Enable IP and verify connectivity to Core1 and Core2.

```
ICX-Tx-Access2(config)# interface vlan 102
ICX-Tx-Access2(config-vlan)# ip address 10.x.102.5/24
ICX-Tx-Access2(config-vlan)# ip ospf 1 area 2
```

```
ICX-Tx-Access2(config)# do ping 10.x.102.2
PING 10.x.102.2 (10.x.102.2) 100(128) bytes of data.
108 bytes from 10.x.102.2: icmp_seq=1 ttl=64 time=16.7 ms
108 bytes from 10.x.102.2: icmp_seq=2 ttl=64 time=0.173 ms
…
ICX-Tx-Access2(config)# do ping 10.x.102.3
PING 10.x.102.3 (10.x.102.3) 100(128) bytes of data.
108 bytes from 10.x.102.3: icmp_seq=1 ttl=64 time=0.182 ms
108 bytes from 10.x.102.3: icmp_seq=2 ttl=64 time=0.175 ms
```

6. Configure the loopback interface.

```
ICX-Tx-Access2(config)# interface loopback 0
ICX-Tx-Access2(config-loopback-if)# ip address 10.x.0.5/32
ICX-Tx-Access2(config-loopback-if)# ip ospf 1 area 2
ICX-Tx-Access2(config-loopback-if)# exit
```

7. Configure some local test subnets.

```
ICX-Tx-Access2(config)# vlan 192,193
ICX-Tx-Access2(config-vlan-<192,193>)# exit
```

```
ICX-Tx-Access2(config)# interface 1/1/4
ICX-Tx-Access2(config-if)# no shutdown
ICX-Tx-Access2(config-if)# vlan trunk allowed 192,193
ICX-Tx-Access2(config-if)# exit
ICX-Tx-Access2(config)#
```

```
ICX-Tx-Access2(config)# interface vlan 192
ICX-Tx-Access2(config-if-vlan)# ip address 10.x.192.1/24
ICX-Tx-Access2(config-if-vlan)# ip ospf 1 area 2
ICX-Tx-Access2(config-if-vlan)# exit
```

```
ICX-Tx-Access2(config)# interface vlan 193
ICX-Tx-Access2(config-if-vlan)# ip address 10.x.193.1/24
ICX-Tx-Access2(config-if-vlan)# ip ospf 1 area 2
ICX-Tx-Access2(config-if-vlan)# exit
ICX-Tx-Access2(config)#
```

8. Verify the OSPF adjacencies have been established with Core1 and Core2.

```
ICX-Tx-Access2(config)# show ip ospf neighbors
OSPF Process ID 1 VRF default
==============================

Total Number of Neighbors: 2

Neighbor ID      Priority  State              Nbr Address       Interface
-------------------------------------------------------------------------
10.x.0.2         1         FULL/BDR           10.x.102.2        vlan102

10.x.0.3         1         FULL/DR            10.x.102.3        vlan102
```

Verify Route Propagation and LSA Changes From Area 2 to Area 0 to Area 1

In the next steps, the route propagation between the areas will be reviewed. For this example, the route 10.x.192.0/24 will be used. This is the test subnet that was defined on Access2 in area 2.

In previous steps, the 'show ip ospf lsdb' command was used to review the LSAs in the LSDB.

In the next steps, the command 'show ip ospf routes' will be used. This command shows the actual routes that have been calculated by OSPF for both intra-area (Router and Network LSA) and inter-area LSAs (Summary LSA).

Access2

9. On Access2, verify the 10.x.192.0/24 route is known as an intra-area route.

NOTE: the lowercase 'i' is used for intra-area, while the uppercase 'I' is used for inter-area.

```
ICX-Tx-Access2(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 : 19

10.x.0.2/32       (I)
     via 10.x.102.2 interface vlan102, cost 100 distance 110
10.x.0.3/32       (I)
     via 10.x.102.3 interface vlan102, cost 100 distance 110
10.x.1.0/24       (I)
     via 10.x.102.2 interface vlan102, cost 200 distance 110
10.x.1.0/24       (I)
     via 10.x.102.3 interface vlan102, cost 200 distance 110
10.x.2.2/31       (I)
     via 10.x.102.2 interface vlan102, cost 101 distance 110
10.x.2.2/31       (I)
     via 10.x.102.3 interface vlan102, cost 101 distance 110
10.x.11.0/24      (I)
     via 10.x.102.2 interface vlan102, cost 200 distance 110
10.x.11.0/24      (I)
     via 10.x.102.3 interface vlan102, cost 200 distance 110
10.x.12.0/24      (I)
     via 10.x.102.2 interface vlan102, cost 200 distance 110
10.x.12.0/24      (I)
     via 10.x.102.3 interface vlan102, cost 200 distance 110
10.x.101.0/24     (I)
     via 10.x.102.2 interface vlan102, cost 200 distance 110
10.x.101.0/24     (I)
     via 10.x.102.3 interface vlan102, cost 200 distance 110
10.x.102.0/24     (i) area: 0.0.0.2
     directly attached to interface vlan102, cost 100 distance 110
10.x.128.0/24     (I)
     via 10.x.102.2 interface vlan102, cost 300 distance 110
10.x.128.0/24     (I)
     via 10.x.102.3 interface vlan102, cost 300 distance 110
10.x.129.0/24     (I)
     via 10.x.102.2 interface vlan102, cost 300 distance 110
10.x.129.0/24     (I)
     via 10.x.102.3 interface vlan102, cost 300 distance 110
10.x.192.0/24     (i) area: 0.0.0.2
     directly attached to interface vlan192, cost 100 distance 110
```

```
10.x.193.0/24     (i) area: 0.0.0.2
     directly attached to interface vlan193, cost 100 distance 110
```

Core1

10. On Core1, review the OSPF routing table, look for the 10.x.192.0/24 subnet.

```
ICX-Tx-Core1(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 : 12

10.x.0.3/32       (i) area: 0.0.0.0
     via 10.x.2.3 interface vlan2, cost 1 distance 110
10.x.0.5/32       (i) area: 0.0.0.2
     via 10.x.102.5 interface vlan102, cost 100 distance 110
10.x.1.0/24       (i) area: 0.0.0.0
     directly attached to interface vlan1, cost 100 distance 110
10.x.2.2/31       (i) area: 0.0.0.0
     directly attached to interface vlan2, cost 1 distance 110
10.x.11.0/24      (i) area: 0.0.0.0
     directly attached to interface vlan11, cost 100 distance 110
10.x.12.0/24      (i) area: 0.0.0.0
     directly attached to interface vlan12, cost 100 distance 110
10.x.101.0/24     (i) area: 0.0.0.1
     directly attached to interface vlan101, cost 100 distance 110
10.x.102.0/24     (i) area: 0.0.0.2
     directly attached to interface vlan102, cost 100 distance 110
10.x.128.0/24     (i) area: 0.0.0.1
     via 10.x.101.4 interface vlan101, cost 200 distance 110
10.x.129.0/24     (i) area: 0.0.0.1
     via 10.x.101.4 interface vlan101, cost 200 distance 110
10.x.192.0/24     (i) area: 0.0.0.2
     via 10.x.102.5 interface vlan102, cost 200 distance 110
10.x.193.0/24     (i) area: 0.0.0.2
     via 10.x.102.5 interface vlan102, cost 200 distance 110
```

Q: Why is the route listed as an intra-area route?

A: Core1 has both area 2 and area 0 databases locally. Both area 2 and area 0 are offering this route to the OSPF routing table. But an intra-area route takes precedence over an inter-area route, so the intra-area route is listed here.

11. Also review the Summary LSAs in the area 0 LSDB.

```
ICX-Tx-Core1(config)# show ip ospf lsdb summary area 0
OSPF Router with ID (10.x.0.2) (Process ID 1 VRF default)
============================================================

Inter-area Summary Link State Advertisements (Area 0.0.0.0)
------------------------------------------------------------

LSID            ADV Router      Age       Seq#       Checksum
------------------------------------------------------------
10.x.0.5       10.x.0.2       454       0x80000001 0x00004b7a
10.x.0.5       10.x.0.3       455       0x80000001 0x0000457f
10.x.101.0     10.x.0.2       404       0x80000003 0x00001e45
10.x.101.0     10.x.0.3       405       0x80000003 0x0000184a
10.x.102.0     10.x.0.2       743       0x80000002 0x0000154e
10.x.102.0     10.x.0.3       740       0x80000002 0x00000f53
10.x.128.0     10.x.0.2       1437      0x80000001 0x0000e302
10.x.128.0     10.x.0.3       1439      0x80000001 0x0000dd07
10.x.129.0     10.x.0.2       1418      0x80000001 0x0000d80c
10.x.129.0     10.x.0.3       1419      0x80000001 0x0000d211
10.x.192.0     10.x.0.2       414       0x80000001 0x00002184
10.x.192.0     10.x.0.3       415       0x80000001 0x00001b89
10.x.193.0     10.x.0.2       398       0x80000001 0x0000168e
10.x.193.0     10.x.0.3       399       0x80000001 0x00001093
```

Q: Which routers are advertising the Summary LSA for the 10.x.192.0/24 subnet?

A: As ABRs, both Core1 and Core2 are advertising this Summary LSA.

12. On Core1, review the area 1 LSDB, look for the 10.x.192.0 Summary LSA.

An ABR will convert Router and Network LSAs from non-backbone areas (not area 0.0.0.0) into Summary LSAs in the backbone area.

These Summary LSAs will be copied as Summary LSAs into any other areas that exist on the ABR.

```
ICX-Tx-Core1(config)# show ip ospf lsdb summary area 1
OSPF Router with ID (10.x.0.2) (Process ID 1 VRF default)
============================================================

Inter-area Summary Link State Advertisements (Area 0.0.0.1)
------------------------------------------------------------

LSID            ADV Router      Age       Seq#       Checksum
------------------------------------------------------------
10.x.0.2       10.x.0.2       459       0x80000002 0x00007bb0
```

```
10.x.0.2        10.x.0.3        460         0x80000002 0x00007faa
10.x.0.3        10.x.0.2        459         0x80000002 0x00007bae
10.x.0.3        10.x.0.3        460         0x80000002 0x00006bbe
10.x.0.5        10.x.0.2        504         0x80000001 0x00004b7a
10.x.0.5        10.x.0.3        505         0x80000001 0x0000457f
10.x.1.0        10.x.0.2        459         0x80000002 0x00007058
10.x.1.0        10.x.0.3        460         0x80000002 0x00006a5d
10.x.2.2        10.x.0.2        459         0x80000002 0x000069c0
10.x.2.2        10.x.0.3        460         0x80000002 0x000063c5
10.x.11.0       10.x.0.2        459         0x80000002 0x000002bc
10.x.11.0       10.x.0.3        460         0x80000002 0x0000fbc1
10.x.12.0       10.x.0.2        459         0x80000002 0x0000f6c6
10.x.12.0       10.x.0.3        460         0x80000002 0x0000f0cb
10.x.102.0      10.x.0.2        793         0x80000002 0x0000154e
10.x.102.0      10.x.0.3        789         0x80000003 0x00000d54
10.x.192.0      10.x.0.2        464         0x80000001 0x00002184
10.x.192.0      10.x.0.3        465         0x80000001 0x00001b89
10.x.193.0      10.x.0.2        448         0x80000001 0x0000168e
10.x.193.0      10.x.0.3        449         0x80000001 0x00001093
```

Q: Which routers are advertising the Summary LSA?

_____

A: Again, as ABRs, both Core1 and Core2 are advertising this Summary LSA.

Access1

13. On Access1, review the OSPF LSDB and the routing table.

```
ICX-Tx-Access1(config)# show ip ospf lsdb summary area 1
OSPF Router with ID (10.x.0.4) (Process ID 1 VRF default)
============================================================

Inter-area Summary Link State Advertisements (Area 0.0.0.1)
------------------------------------------------------------

LSID            ADV Router      Age       Seq#       Checksum
------------------------------------------------------------
10.x.0.2        10.x.0.2        536       0x80000002 0x00007bb0
10.x.0.2        10.x.0.3        535       0x80000002 0x00007faa
10.x.0.3        10.x.0.2        536       0x80000002 0x00007bae
10.x.0.3        10.x.0.3        535       0x80000002 0x00006bbe
10.x.0.5        10.x.0.2        580       0x80000001 0x00004b7a
10.x.0.5        10.x.0.3        580       0x80000001 0x0000457f
10.x.1.0        10.x.0.2        536       0x80000002 0x00007058
10.x.1.0        10.x.0.3        535       0x80000002 0x00006a5d
10.x.2.2        10.x.0.2        536       0x80000002 0x000069c0
10.x.2.2        10.x.0.3        535       0x80000002 0x000063c5
10.x.11.0       10.x.0.2        536       0x80000002 0x000002bc
10.x.11.0       10.x.0.3        535       0x80000002 0x0000fbc1
10.x.12.0       10.x.0.2        536       0x80000002 0x0000f6c6
10.x.12.0       10.x.0.3        535       0x80000002 0x0000f0cb
```

```
10.x.102.0      10.x.0.2        870         0x80000002 0x0000154e
10.x.102.0      10.x.0.3        864         0x80000003 0x00000d54
10.x.192.0      10.x.0.2        540         0x80000001 0x00002184
10.x.192.0      10.x.0.3        540         0x80000001 0x00001b89
10.x.193.0      10.x.0.2        524         0x80000001 0x0000168e
10.x.193.0      10.x.0.3        524         0x80000001 0x00001093
```

14. Review the OSPF routes, the route to 10.x.192.0/24 should be listed as an inter-area route.

```
ICX-Tx-Access1(config)# show ip ospf routes 10.x.192.0/24
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 for prefixes 10.x.192.0/24
------------------------------------------------------------------------

Total Number of Routes : 2

10.x.192.0/24     (I)
     via 10.x.101.2 interface vlan101, cost 300 distance 110
10.x.192.0/24     (I)
     via 10.x.101.3 interface vlan101, cost 300 distance 110

ICX-Tx-Access1(config)#
```

This demonstrates that a route on a Router LSA in area 2 is converted into a Summary LSA by the ABR into the backbone area (0.0.0.0).

An ABR will copy any Summary LSA from the backbone area into any other area LSDB.

NOTE: An ABR will only exchange information between an area and the backbone (0.0.0.0) area LSDB databases. It will never exchange information directly between two non-backbone areas, for example between area 1 and area 2 LSDBs.

Task 3: Route Summarization

Diagram L3

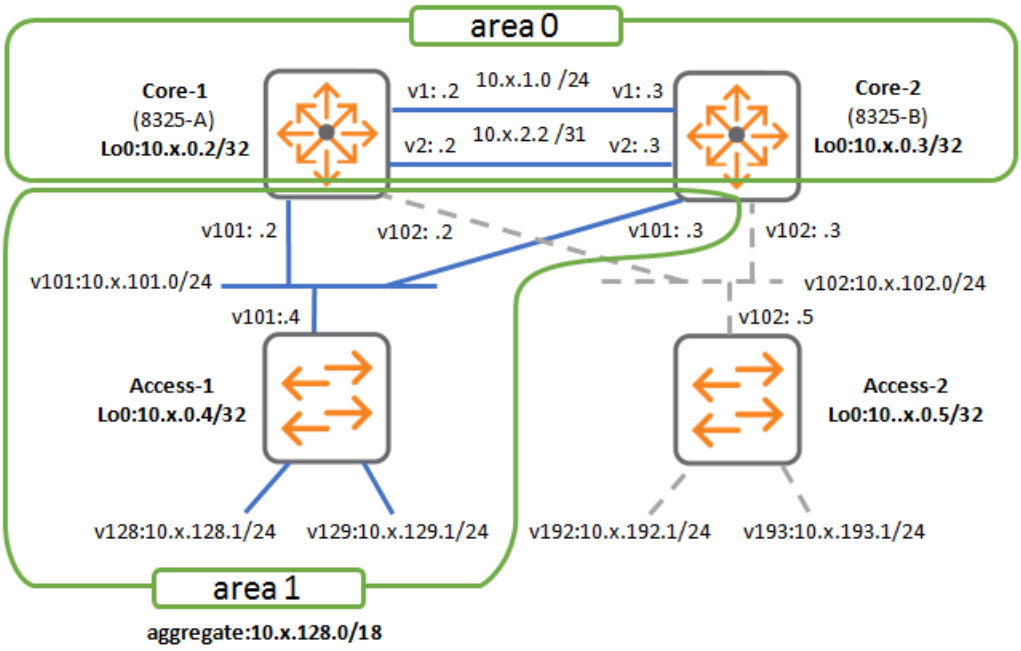

Objectives

- Implement route summarization between the areas

Steps

First Summarize routes injected by area 1 into area 0 on Core1.

Verify in area 2, verify LSA type and cost.

Summarize area 2 test subnets.

Summarize the Area 1 Test Subnets.

In area 1, the test subnets 10.x.128.0/24 and 10.x.129.0/24 can be summarized into a 10.x.128.0/18 prefix.

Core1

1. On Core1, configure the summarization rule in the OSPF area 1 context.

```
ICX-Tx-Core1(config)# router ospf 1
ICX-Tx-Core1(config-ospf-1)# area 1 range 10.x.128.0/18 type inter-area
ICX-Tx-Core1(config-ospf-1)# exit
```

IMPORTANT: The route aggregation must be performed on every ABR that is connected to this area to be effective. An ABR that does not have the aggregation route will still be announcing the more specific routes into the area.
As a result, all traffic to these subnets will go via the 'most specific match', meaning the ABR that does not have the aggregate route configured.
In this lab environment, the VSX configuration synchronization is also synchronizing the OSPF commands, the 'range' command is automatically set on the ABR Core2.

2. Review how this affects the area 0 LSDB on Core1. Previously, these subnets were known as unique Summary LSAs, now only the summarized route is listed.

```
ICX-Tx-Core1(config)# show ip ospf lsdb summary area 0
OSPF Router with ID (10.x.0.2) (Process ID 1 VRF default)
============================================================

Inter-area Summary Link State Advertisements (Area 0.0.0.0)
------------------------------------------------------------

LSID            ADV Router      Age       Seq#       Checksum
------------------------------------------------------------
10.x.0.5       10.x.0.2       962       0x80000001 0x00004b7a
10.x.0.5       10.x.0.3       963       0x80000001 0x0000457f
10.x.101.0     10.x.0.2       912       0x80000003 0x00001e45
10.x.101.0     10.x.0.3       913       0x80000003 0x0000184a
10.x.102.0     10.x.0.2       1251      0x80000002 0x0000154e
10.x.102.0     10.x.0.3       1248      0x80000002 0x00000f53
10.x.128.0     10.x.0.2       162       0x80000002 0x0000a57e
10.x.128.0     10.x.0.3       161       0x80000002 0x00009f83
10.x.192.0     10.x.0.2       922       0x80000001 0x00002184
10.x.192.0     10.x.0.3       923       0x80000001 0x00001b89
10.x.193.0     10.x.0.2       906       0x80000001 0x0000168e
10.x.193.0     10.x.0.3       907       0x80000001 0x00001093
```

3. And since the area 0 only contains the summarized entry, the area 2 LSDB will also reflect this. Review Summary LSAs in the LSDB for area 2.

```
ICX-Tx-Core1(config)# show ip ospf lsdb summary area 2
OSPF Router with ID (10.x.0.2) (Process ID 1 VRF default)
============================================================

Inter-area Summary Link State Advertisements (Area 0.0.0.2)
------------------------------------------------------------

LSID            ADV Router      Age       Seq#       Checksum
------------------------------------------------------------
10.x.0.2       10.x.0.2       1322      0x80000001 0x00007daf
10.x.0.2       10.x.0.3       1319      0x80000001 0x000081a9
10.x.0.3       10.x.0.2       1322      0x80000001 0x00007dad
10.x.0.3       10.x.0.3       1319      0x80000001 0x00006dbd
10.x.1.0       10.x.0.2       1322      0x80000001 0x00007257
10.x.1.0       10.x.0.3       1319      0x80000001 0x00006c5c
```

```
10.x.2.2        10.x.0.2        1322       0x80000001 0x00006bbf
10.x.2.2        10.x.0.3        1319       0x80000001 0x000065c4
10.x.11.0       10.x.0.2        1322       0x80000001 0x000004bb
10.x.11.0       10.x.0.3        1319       0x80000001 0x0000fdc0
10.x.12.0       10.x.0.2        1322       0x80000001 0x0000f8c5
10.x.12.0       10.x.0.3        1319       0x80000001 0x0000f2ca
10.x.101.0      10.x.0.2        1322       0x80000001 0x00002243
10.x.101.0      10.x.0.3        1319       0x80000001 0x00001c48
10.x.128.0      10.x.0.2        233        0x80000002 0x0000a57e
10.x.128.0      10.x.0.3        229        0x80000003 0x00009d84
```

Access2

4. On Access2, this is a router that only knows area 2, review the LSDB. Since all the database for an area must be the same on all the routers of an area, this area 2 LSDB will also contain the summarized Summary LSA.

```
ICX-Tx-Access2(config)# show ip ospf lsdb summary area 2
OSPF Router with ID (10.x.0.5) (Process ID 1 VRF default)
============================================================

Inter-area Summary Link State Advertisements (Area 0.0.0.2)
------------------------------------------------------------

LSID             ADV Router      Age       Seq#       Checksum
------------------------------------------------------------------
10.x.0.2        10.x.0.2        1368       0x80000001 0x00007daf
10.x.0.2        10.x.0.3        1365       0x80000001 0x000081a9
10.x.0.3        10.x.0.2        1368       0x80000001 0x00007dad
10.x.0.3        10.x.0.3        1365       0x80000001 0x00006dbd
10.x.1.0        10.x.0.2        1368       0x80000001 0x00007257
10.x.1.0        10.x.0.3        1365       0x80000001 0x00006c5c
10.x.2.2        10.x.0.2        1368       0x80000001 0x00006bbf
10.x.2.2        10.x.0.3        1365       0x80000001 0x000065c4
10.x.11.0       10.x.0.2        1368       0x80000001 0x000004bb
10.x.11.0       10.x.0.3        1365       0x80000001 0x0000fdc0
10.x.12.0       10.x.0.2        1368       0x80000001 0x0000f8c5
10.x.12.0       10.x.0.3        1365       0x80000001 0x0000f2ca
10.x.101.0      10.x.0.2        1368       0x80000001 0x00002243
10.x.101.0      10.x.0.3        1365       0x80000001 0x00001c48
10.x.128.0      10.x.0.2        281        0x80000002 0x0000a57e
10.x.128.0      10.x.0.3        275        0x80000003 0x00009d84
```

5. Verify the OSPF routing table, it should only contain the summarized entry **10.x.128.0/18**, so the **10.x.128.0/24** and **10.x.129.0/24** subnets should not appear anymore.

```
ICX-Tx-Access2(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 : 17

10.x.0.2/32       (I)
     via 10.x.102.2 interface vlan102, cost 100 distance 110
10.x.0.3/32       (I)
     via 10.x.102.3 interface vlan102, cost 100 distance 110
10.x.1.0/24       (I)
     via 10.x.102.2 interface vlan102, cost 200 distance 110
10.x.1.0/24       (I)
     via 10.x.102.3 interface vlan102, cost 200 distance 110
10.x.2.2/31       (I)
     via 10.x.102.2 interface vlan102, cost 101 distance 110
10.x.2.2/31       (I)
     via 10.x.102.3 interface vlan102, cost 101 distance 110
10.x.11.0/24      (I)
     via 10.x.102.2 interface vlan102, cost 200 distance 110
10.x.11.0/24      (I)
     via 10.x.102.3 interface vlan102, cost 200 distance 110
10.x.12.0/24      (I)
     via 10.x.102.2 interface vlan102, cost 200 distance 110
10.x.12.0/24      (I)
     via 10.x.102.3 interface vlan102, cost 200 distance 110
10.x.101.0/24     (I)
     via 10.x.102.2 interface vlan102, cost 200 distance 110
10.x.101.0/24     (I)
     via 10.x.102.3 interface vlan102, cost 200 distance 110
10.x.102.0/24     (i) area: 0.0.0.2
     directly attached to interface vlan102, cost 100 distance 110
10.x.128.0/18     (I)
     via 10.x.102.2 interface vlan102, cost 300 distance 110
10.x.128.0/18     (I)
     via 10.x.102.3 interface vlan102, cost 300 distance 110
10.x.192.0/24     (i) area: 0.0.0.2
     directly attached to interface vlan192, cost 100 distance 110
10.x.193.0/24     (i) area: 0.0.0.2
     directly attached to interface vlan193, cost 100 distance 110
```

Reject Route for the Summarized Route

When an ABR is announcing a summary route to other routers, it introduces the risk that remote systems may send traffic to a subnet that is simply not online or available.

In this example, Core1 and Core2 ABRs are announcing the 10.x.128.0/18 summarized route to Area 2.

This means that devices connected to Access2 could try to send traffic to, for example, the 10.x.130.0/24 subnet, since this traffic also fits under the 10.x.128.0/18 route.

To prevent this, the ABR router that has the summarized route configured and will automatically get a 'reject' route in its routing table for the summarized route. In this example the 10.x.128.0/18 route is rejected.

While this may seem like it would block all the traffic, the ABR still has the more specific routes for the 10.x.128.0/24 and 10.x.129.0/24 in its routing table, so these routes take precedence over the summarized route.

Core1

6. On Core1, review the IP routing table. Verify that the summarized route is listed as a 'reject' route, while the more specific routes still exist as normal routes.

```
ICX-Tx-Core1(config)# show ip route

Displaying ipv4 routes selected for forwarding

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

...
10.x.128.0/18, vrf default
        reject route,  [110/0],  static
10.x.128.0/24, vrf default
        via  10.x.101.4,  [110/200],  ospf
10.x.129.0/24, vrf default
        via  10.x.101.4,  [110/200],  ospf
...
```

7. Therefore, if any traffic would arrive for the non-existing 10.x.130.0/24 subnet, it would be rejected since it only matches the reject route.

```
ICX-Tx-Core1(config)# show ip route 10.x.130.0

Displaying ipv4 routes selected for forwarding

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

10.x.128.0/18, vrf default
        reject route,  [110/0],  static

ICX-Tx-Core1(config)#
```

NOTE: While this reject route is listed as a static route, it will not be visible on the running configuration. OSFP summarized route entries are automatically added to the operational routing table as a static route with the reject option.

Summarize Area 2 Test Subnets

Diagram L3

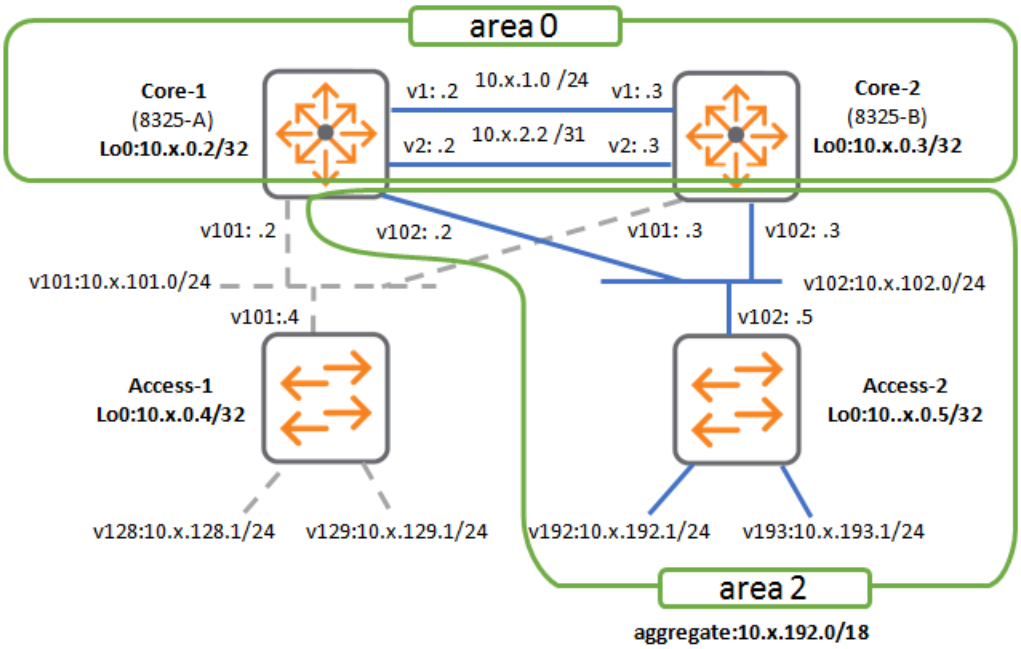

8. Now attempt to summarize the area 2 test subnets 10.x.192.0/24 and 10.x.193.0/24 into a 10.x.192.0/18 summary route.

Try to apply the commands yourself. If you are not sure, the commands can be found in the next steps.

Access1

9. Verify on Access1 (area1) that only the summary route is listed.

```
ICX-Tx-Access1(config)# show ip route 10.x.192.0

Displaying ipv4 routes selected for forwarding

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

10.x.192.0/18, vrf default
        via  10.x.101.3,  [110/300],  ospf
        via  10.x.101.2,  [110/300],  ospf

ICX-Tx-Access1(config)#
```

Solution

On Core1

```
ICX-Tx-Core1(config)# router ospf 1
ICX-Tx-Core1(config-ospf-1)# area 2 range 10.x.192.0/18 type inter-area
ICX-Tx-Core1(config-ospf-1)# exit
ICX-Tx-Core1(config)#
```

Task 4: Verify Route Propagation Impact with Summarization

Using route summarization reduces the routing table size, but it also minimizes the impact of routing updates.

In case where a summarized route is added or removed, the change will only be seen up to the point where the summarization takes place.

In a large OSPF network, this has significant advantages and ensures better network stability due to fewer routing updates.

Objectives

- Verify OSPF LSA update propagates over the entire autonomous system by default.
- Verify LSA propagation can be controlled using summarization.

Steps

Default LSA propagation

In the lab setup, the loopback 0 interface of Access2 will be disabled. Since this interface has IP address 10.x.0.5/32, and this does **not** match the summarization rule, the LSA change will propagate over the entire Autonomous System (AS).

These are the high-level steps:

- First check the OSPF SPF count on Access1.
- On Access2, disable the loopback interface to simulate a route change.
- On Access1, check the SPF count again, it should have increased since the LSA is propagated over the entire AS.

Access1

1. On Access1, review the OSPF SPF count and routing table.

NOTE: In each lab, the number of calculations may be different, the output below is just an example.

```
ICX-Tx-Access1(config)# show ip ospf
Routing Process 1 with ID : 10.x.0.4 VRF default
-------------------------------------------------

OSPFv2 Protocol is enabled
Graceful-restart is configured
Restart Interval: 120, State: inactive
Last Graceful Restart Exit Status: none
SPF: Start Time: 200ms, Hold Time: 1000ms, Max Wait Time: 5000ms
Maximum Paths to Destination: 4
Number of external LSAs 0, checksum sum 0
```

```
Number of areas is 1, 1 normal, 0 stub, 0 NSSA
Number of active areas is 1, 1 normal, 0 stub, 0 NSSA
BFD is disabled
Reference Bandwidth: 100000 Mbps
Area (0.0.0.1) (Active)
  Interfaces in this Area: 3 Active Interfaces: 3
  Passive Interfaces: 0 Loopback Interfaces: 0
  SPF calculation has run 42 times
  Area ranges:
  Number of LSAs: 20, checksum sum 652089
```

2. On Access1, check the OSPF database summary.

```
ICX-Tx-Access1(config)# show ip ospf lsdb database-summary
OSPF Router with ID (10.x.0.4) (Process ID 1 VRF default)
============================================================

Area 0.0.0.1 database summary
------------------------------

LSA Type             Count
---------------------------
Router               3
Network              1
Inter-area Summary   18
ASBR Summary         0
NSSA External        0
Subtotal             22

Process 1 database summary
---------------------------

LSA Type             Count
---------------------------
Router               3
Network              1
Inter-area Summary   18
ASBR Summary         0
NSSA External        0
AS External          0
Total                22
```

Access2

3. On Access2, disable the loopback 0 interface for OSPF.

```
ICX-Tx-Access2(config)# interface loopback 0
ICX-Tx-Access2(config-loopback-if)# ip ospf shutdown
```

Access1

4. On Access1, review the SPF count and the LSDB database summary.

```
ICX-Tx-Access1(config)# show ip ospf | include SPF
```

```
OSPFv2 Protocol is enabled
SPF: Start Time: 200ms, Hold Time: 1000ms, Max Wait Time: 5000ms
  SPF calculation has run 44 times
ICX-Tx-Access1(config)#
```

```
ICX-Tx-Access1(config)# show ip ospf lsdb database-summary
OSPF Router with ID (10.x.0.4) (Process ID 1 VRF default)
============================================================

Area 0.0.0.1 database summary
------------------------------

LSA Type             Count
---------------------------
Router               3
Network              1
Inter-area Summary   16
ASBR Summary         0
NSSA External        0
Subtotal             20

<…output omitted…>
```

Access2

5. On Access2, enable the Loopback 0 for OSPF again.

```
ICX-Tx-Access2(config-loopback-if)# no ip ospf shutdown
ICX-Tx-Access2(config-loopback-if)# exit
ICX-Tx-Access2(config)#
```

This demonstrates that LSA updates are propagated over the entire AS by default.

Verify LSA Propagation Can Be Controlled Using Summarization

Now an interface that was summarized will be disabled. This should have no impact on the LSDB and the SPF count in the other areas.

For this example, on Access2, the interface VLAN 192 will be disabled.

Again, verify the SPF count on Access1 first, then make the change on Access2, and verify the count again on Access1.

Access1

6. On Access1, review the current SPF count.

```
ICX-Tx-Access1(config)# show ip ospf | include SPF
OSPFv2 Protocol is enabled
SPF: Start Time: 200ms, Hold Time: 1000ms, Max Wait Time: 5000ms
  SPF calculation has run 45 times
```

Access2

7. On Access2, disable the VLAN 192 interface. This interface was summarized using the 10.x.192.0/18 range command.

```
ICX-Tx-Access2(config)# interface vlan 192
ICX-Tx-Access2(config-if-vlan)# shutdown
```

Access1

8. On Access1, verify the SPF count has not increased.

```
ICX-Tx-Access1(config)# show ip ospf | include SPF
OSPFv2 Protocol is enabled
SPF: Start Time: 200ms, Hold Time: 1000ms, Max Wait Time: 5000ms
  SPF calculation has run 45 times
```

This demonstrates the advantage of route summarization between OSPF areas.

Access2

9. On Access2, enable the VLAN 192 interface again.

```
ICX-Tx-Access2(config)# interface vlan 192
ICX-Tx-Access2(config-if-vlan)# no shutdown
ICX-Tx-Access2(config-if-vlan)# exit
```

Task 5: ABR Route Filtering

Diagram

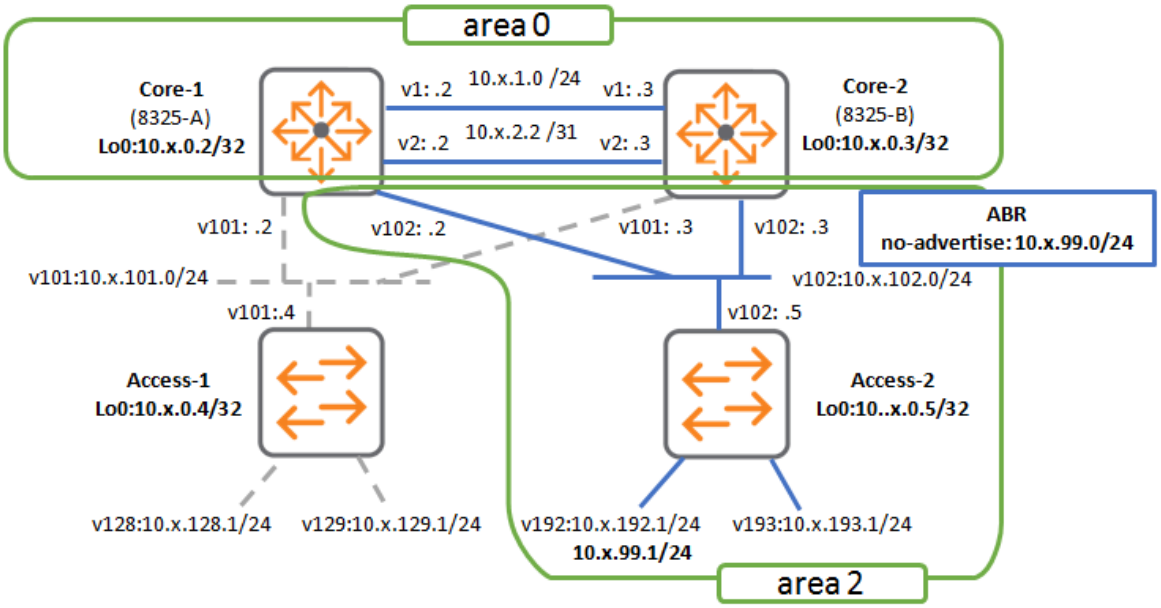

On the ABR, the network administrator can also filter routes.

This feature is useful in case some routes only need to be known within the local area, but they should not be advertised to the rest of the autonomous system.

Objectives

- Understand ABR route filtering
- Configure ABR route filtering

Steps

To demonstrate this, a new IP interface will be enabled on Access2, in area 2.

This route will be visible to the Core1 and Core2 routers, since they also belong to area 2. Therefore, the route can be used by all the routers that belong to this area.

However, the administrator does not want this route to be visible in the other areas, so a route filter will be used.

- Introduce a new subnet on Access2 in area 2.
- On the ABR, filter the route in the OSPF area 2 LSDB.
- Verify the result on Access1, a route in area 1.

Access2

1. On Access2, introduce a new route by adding a secondary IP address on the VLAN 192 interface.

```
ICX-Tx-Access2(config)# interface vlan 192
ICX-Tx-Access2(config-if-vlan)# ip address 10.x.99.1/24 secondary
ICX-Tx-Access2(config-if-vlan)#
```

Access1

2. On Access1, verify the route is visible in the routing table, since no filtering has been applied yet.

```
ICX-Tx-Access1(config)# show ip route 10.x.99.0

Displaying ipv4 routes selected for forwarding

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

10.x.99.0/24, vrf default
        via  10.x.101.3,  [110/200],  ospf
        via  10.x.101.2,  [110/200],  ospf
```

Core1

3. On Core1, review that the 10.x.99.0 Summary LSA exists in the area 0 database.

```
ICX-Tx-Core1(config)# show ip ospf lsdb summary area 0 | include 99.0
10.x.99.0      10.x.0.2       10         0x80000001 0x0000382f
10.x.99.0      10.x.0.3       9          0x80000001 0x00003234
```

4. Apply the filter for area 2 so the 10.x.99.0/24 subnet is not announced anymore.

```
ICX-Tx-Core1(config)# router ospf 1
ICX-Tx-Core1(config-ospf-1)# area 2 range 10.x.99.0/24 type inter-area no-
advertise
ICX-Tx-Core1(config-ospf-1)# exit
```

5. Verify that the 10.x.99.0/24 summary route is no longer listed in the area 0 database.

```
ICX-Tx-Core1(config)# show ip ospf lsdb summary area 0 | include 99.0
ICX-Tx-Core1(config)#
```

6. As a result, it will also be removed from the area 1 database.

```
ICX-Tx-Core1(config)# show ip ospf lsdb summary area 1 | include 99.0
```

```
ICX-Tx-Core1(config)#
```

Access1

7. Verify the result on Access1.

```
ICX-Tx-Access1(config)# show ip route 10.x.99.0

No ipv4 routes configured
```

All devices

8. Save the configurations on all devices using **'write memory**'.

You've completed Lab 6.2!

Lab 06.3: OSPF External Routes- Managing External Routes with OSPF

Lab Diagram

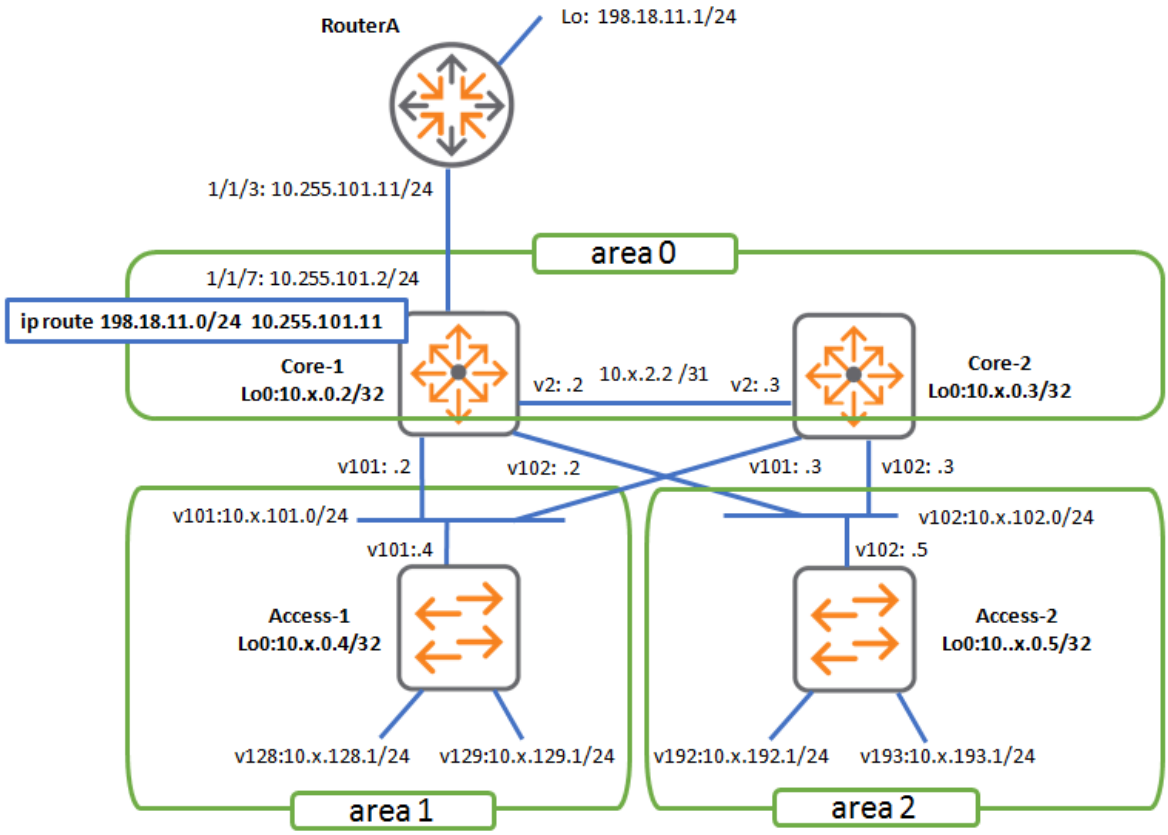

Overview

This lab activity will demonstrate how OSPF can be used to distribute routes to external networks to other OSPF routers in the OSPF Autonomous System (AS).

The router that is performing this redistribution is referred to as the Autonomous System Boundary Router (ASBR). This will be demonstrated using Core1 in the backbone and Access2 in area2.

Once the external routes have been exchanged, the lab will show how a route map can be used to control the cost and the list of routes that can be redistributed.

The last section of the lab will show how area types can be used to optimize which external and OSPF routes will be announced into an area.

Objectives

- Learn how to configure an ASBR to redistribute routes.
- Configure a route map to control external cost types, such as Metric type1 and type2.
- Understand the difference between Stub, Stub no-summary and Not So Stubby Area types.

Task 1: Setup Link to RouterA

Objectives

In this task, the routed connection between the Core1 and RouterA will be verified.

On Core1, a static route will be configured to this RouterA, so Core1 (and only Core1) can reach the external network at this point.

Steps

Verify the routed connection to RouterA using interface 1/1/7.

Core1

1. On Core1, enter the configuration mode.
2. Review the configuration of interface 1/1/7. This was configured during the VSX lab activity.

NOTE: The IP scheme for the upstream routers is local for each table, so there is no 'x' marking in these IP addresses.

```
ICX-Tx-Core1(config)# interface 1/1/7
ICX-Tx-Core1(config-if)# show run current-context
interface 1/1/7
    no shutdown
    ip address 10.255.101.2/24
    exit
ICX-Tx-Core1(config-if)# exit
```

3. Verify connectivity to the RouterA IP.

```
ICX-Tx-Core1(config)# do ping 10.255.101.11
PING 10.255.101.11 (10.255.101.11) 100(128) bytes of data.
108 bytes from 10.255.101.11: icmp_seq=1 ttl=64 time=1.75 ms
...
```

4. Define a static route to the remote subnet. This represents a business partner network outside of this organization OSPF AS.

```
ICX-Tx-Core1(config)# ip route 198.18.11.0/24 10.255.101.11
```

5. Test with a ping to 198.18.11.1, this IP is hosted by the RouterA using a loopback interface.

```
ICX-Tx-Core1(config)# do ping 198.18.11.1
PING 198.18.11.1 (198.18.11.1) 100(128) bytes of data.
108 bytes from 198.18.11.1: icmp_seq=1 ttl=64 time=2.20 ms
…
```

6. Review the IP Routing table on the Core1.

```
ICX-Tx-Core1(config)# show ip route static

Displaying ipv4 routes selected for forwarding

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

10.x.128.0/18, vrf default
        reject route,  [110/0],  static
10.x.192.0/18, vrf default
        reject route,  [110/0],  static
10.254.1.0/24, vrf default
        via  10.255.101.11,  [1/0],  static
198.18.11.0/24, vrf default
        via  10.255.101.11,  [1/0],  static
```

At this point, the Core1 can reach the external network, but the rest of the OSPF AS cannot reach this 198.18.11.0/24 network yet. This will be solved with route redistribution in the next task.

Task 2: Redistribute Static Routes into OSPF

Diagram

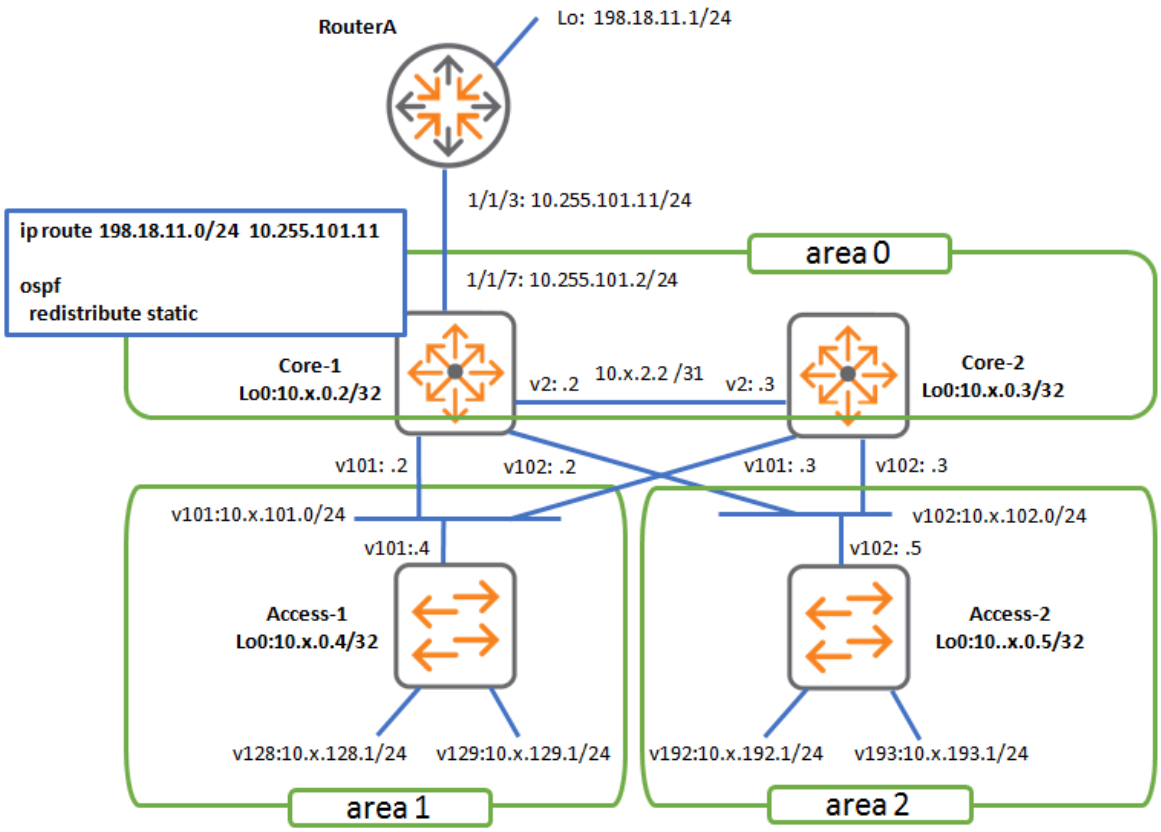

Objectives

In this task, the static route that was installed on the Core1 router will now be distributed to the other routers using OSPF route redistribution.

The router that is inserting 'external' routes into OSPF is known as the ASBR, the autonomous system boundary router.

The ASBR can be a router in the backbone area or a router in another area. Other area types are covered in a later task.

In order to allow other OSPF routers to make a distinction between the routes of the OSPF autonomous system and routes that are injected (so not originated) into the autonomous system (external routes), OSPF uses a different LSA type from the intra-area (Router/Type1 and Network/Type2) and the inter-area (Summary/Type3).

All the routes that are injected by the ASBR are inserted as Type5 LSAs in the LSDB.

Therefore, when a complete AS has two routers that are injecting 100 routes each, the LSDB will contain 200 Type5 LSAs (2x100 routes).

Steps

Redistribute static routes into ospf.

Verify the OSPF LSA Type5 in the backbone and other areas.

Core1

1. Open a terminal session to Core1, enter the configuration mode.
2. Review the current OSPF database summary. There are currently no external routes in the database.

```
ICX-Tx-Core1(config)# show ip ospf lsdb database-summary
...
Process 1 database summary
---------------------------

LSA Type             Count
---------------------------
Router               8
Network              2
Inter-area Summary   44
ASBR Summary         0
NSSA External        0
AS External          0
Total                54
```

3. Enter the OSPF router context, enable route redistribution of the static routes.

This will instruct the OSPF process to look into the local routing table for routes that match the redistribution condition, in this case, the condition is 'static' routes.

```
ICX-Tx-Core1(config)# router ospf 1
ICX-Tx-Core1(config-ospf-1)# redistribute static
ICX-Tx-Core1(config-ospf-1)# exit
```

4. Verify in the LSDB that a new record was created for this external route.

```
ICX-Tx-Core1(config)# show ip ospf lsdb database-summary
...
Process 1 database summary
---------------------------

LSA Type             Count
---------------------------
```

```
Router               8
Network              2
Inter-area Summary   44
ASBR Summary         0
NSSA External        0
AS External          3
Total                57
```

```
ICX-Tx-Core1(config)# show ip ospf lsdb external
OSPF Router with ID (10.x.0.2) (Process ID 1 VRF default)
============================================================

AS External Link State Advertisements
--------------------------------------

LSID            ADV Router      Age       Seq#       Checksum
------------------------------------------------------------
10.254.1.0      10.x.0.2        129       0x80000001 0x00009502
10.254.1.0      10.x.0.3        129       0x80000001 0x00008f07
198.18.11.0     10.x.0.2        129       0x80000001 0x0000ac11
```

NOTE: The 10.254.1.0/24 static route was added during the VSX lab on both Core1 and Core2. It provides access to the ClearPass server subnet. Due to the VSX OSPF sync, Core2 is now also redistributing its static route to the 10.254.1.0/24 subnet. These routes can be ignored in this lab.

Core2

5. Now check the LSDB on Core2 and Access2. Access2 is connected via Core2 to Core1.

Review that the external route has arrived in the LSDB.

```
ICX-Tx-Core2(config)# show ip ospf lsdb external
OSPF Router with ID (10.x.0.3) (Process ID 1 VRF default)
============================================================

AS External Link State Advertisements
--------------------------------------

LSID            ADV Router      Age       Seq#       Checksum
------------------------------------------------------------
10.254.1.0      10.x.0.2        176       0x80000001 0x00009502
10.254.1.0      10.x.0.3        174       0x80000001 0x00008f07
198.18.11.0     10.x.0.2        176       0x80000001 0x0000ac11
```

Access2

```
ICX-Tx-Access2(config)# show ip ospf lsdb external
OSPF Router with ID (10.x.0.5) (Process ID 1 VRF default)
============================================================

AS External Link State Advertisements
--------------------------------------

LSID            ADV Router      Age       Seq#       Checksum
------------------------------------------------------------
10.254.1.0      10.x.0.2        200       0x80000001 0x00009502
10.254.1.0      10.x.0.3        199       0x80000001 0x00008f07
198.18.11.0     10.x.0.2        200       0x80000001 0x0000ac11
```

6. On Core1, this route is a static route in the routing table. Now review the IP routing table on Access2.

```
ICX-Tx-Access2(config)# show ip route 198.18.11.0
...

Displaying ipv4 routes selected for forwarding

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

198.18.11.0/24, vrf default
        via  10.x.102.2,  [110/25],  ospf
```

Q: Which router is the next hop IP of this route?

_____

A: This 10.x.102.2 is the Core1 router. It was calculated by OSPF to be the shortest path to the ASBR that is announcing this route (Core1).

Access1

7. On Access1, attempt to ping the remote host 198.18.11.1 to validate that all the routers on the path to the ASBR know how to route this traffic.

```
ICX-Tx-Access1(config)# do ping 198.18.11.1
PING 198.18.11.1 (198.18.11.1) 100(128) bytes of data.
108 bytes from 198.18.11.1: icmp_seq=1 ttl=63 time=2.52 ms
...
```

NOTE: IP always needs bi-directional communication paths. On the remote router, a static route for 10.0.0.0/8 exists already in the configuration to handle the return traffic.

This demonstrates how OSPF can be used to distribute routes in the OSPF Autonomous System that were not originated by the OSPF routers themselves.

Access2 - External Routes

Diagram L3

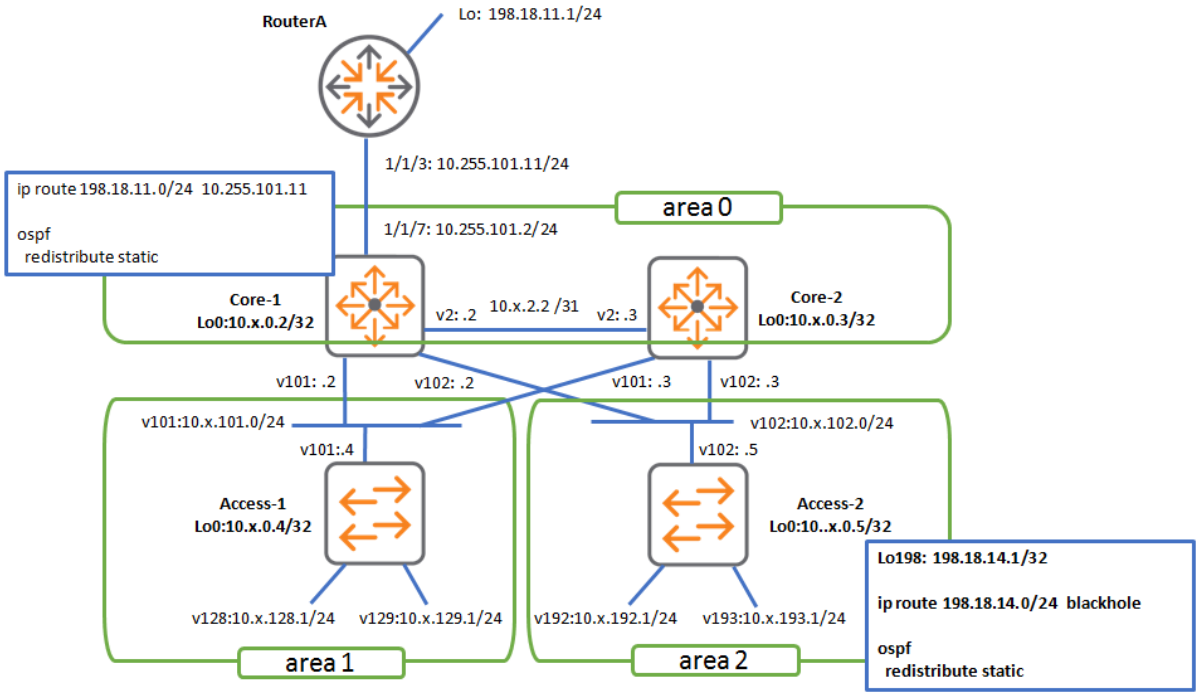

Autonomous System Boundary Routers (ASBR) do not have to be in the backbone area. A router in a normal area can also redistribute routes into the OSPF Autonomous System.

In the next steps, Access2 will be configured to inject an external route.

For testing purposes, Access2 will be configured with a local static route that will be redistributed into OSPF.

To provide a test IP in this subnet, a local loopback interface will be used.

Since this test loopback interface is not enabled for OSPF, it will not be known as a native OSPF route. Only the static route will be known to OSPF due to the redistribution.

Access2

8. On Access2 (OSPF area 2), define a new static route. A normal static route would be configured to a real next-hop IP address. Since this is a test subnet, and there is no real nexthop router, the next-hop will be set to 'blackhole'. This means that all traffic arriving for this subnet will be dropped. In order to perform a test ping, a loopback address will be made in this test subnet.

```
ICX-Tx-Access2(config)# ip route 198.18.14.0/24 blackhole
```

9. Define a new loopback interface, this can be used to test access to the 198.18.14.0/24 subnet.

```
ICX-Tx-Access2(config)# interface loopback 198
ICX-Tx-Access2(config-loopback-if)# ip address 198.18.14.1/32
ICX-Tx-Access2(config-loopback-if)# exit
```

10. Redistribute the static route into OSPF.

```
ICX-Tx-Access2(config)# router ospf 1
ICX-Tx-Access2(config-ospf-1)# redistribute static
ICX-Tx-Access2(config-ospf-1)# exit
ICX-Tx-Access2(config)#
```

Core1

11. Verify in the backbone area that the External route is available.

```
ICX-Tx-Core1(config)# show ip ospf routes 198.18.14.0/24
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 for prefixes 198.18.14.0/24
------------------------------------------------------------------------

Total Number of Routes : 1

198.18.14.0/24     (E2)
     via 10.x.102.5 interface vlan102, cost 25 distance 110
```

Access1

Verify in OSPF area 1 that the new external route is available.

12. On Access1, check the IP routing table for the 198.18.14.0 subnet.

```
ICX-Tx-Access1# show ip route 198.18.14.0

Displaying ipv4 routes selected for forwarding

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

198.18.14.0/24, vrf default
        via  10.x.101.3,  [110/25],  ospf
        via  10.x.101.2,  [110/25],  ospf
```

Q: Why are there 2 paths available for this route?

A: OSPF calculates the best path to reach the ASBR. In this case, there are 2 equal cost paths to reach the Access2 router, so both paths are entered in the routing table.

13. Perform a traceroute to the test IP in the 198.18.14.0/24 subnet. The path should use either Core1 or Core2 to reach the Access2 router.

```
ICX-Tx-Access1# traceroute 198.18.14.1
traceroute to 198.18.14.1 (198.18.14.1), 1 hops min, 30 hops max, 3 sec. timeout,
3 probes
  1   10.x.101.2  0.202ms  0.099ms  0.504ms
  2   198.18.14.1  0.188ms  0.101ms  0.099ms
ICX-Tx-Access1#
```

Task 3: Control Route Redistribution and Metric Types

Introduction to Metric Types

OSPF has 2 metric types to describe the cost for external routes: External type 1 and External type 2 (default).

External Type2

In the case of External type 2, the external route is entered into the OSPF LSDB with a static cost. This static cost is **not** adjusted as the route is distributed by the OSPF routers. Therefore, a router that is several hops away would still see the route with the same cost that was initially set by the redistributing ASBR.

This is convenient to force the entire network to use a specific ASBR to reach some external network, instead of the 'closest' ASBR. This applies to scenarios where the traffic to the external network must pass a firewall and the customer wants to avoid asynchronous routing.

External Type1

When there is no need for a 'fixed' ASBR to the external network, OSPF can be configured to use the path to the closest ASBR. This can be achieved by configuring the external route with a metric type 1.

In this case, the OSPF path cost to reach the ASBR is added to the initial cost of the external route.

Route maps

The configuration of the route metric types is done through a route policy.

Route policies can also be used to filter routes that should not be redistributed.

Diagram L3

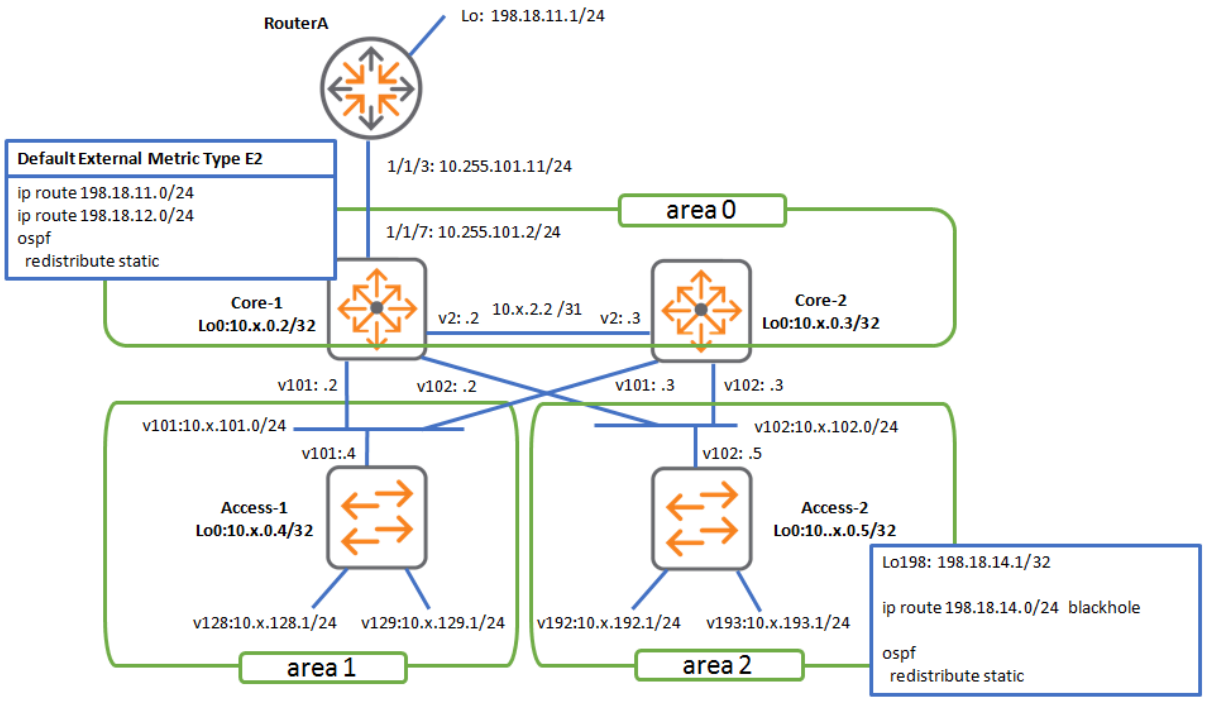

Steps

Review the default cost and metric type of OSPF external routes.

Add a static route on Core1 that should not be redistributed using OSPF.

Define a route policy on Core1 to set the OSPF metric-type 2 and set the initial cost to 100.

Core1

1. On Core1, add a dummy static route. In the current configuration, every static route will be redistributed by OSPF. This route will be filtered by the route policy in the upcoming step.

```
ICX-Tx-Core1(config)# ip route 198.18.12.0/24 blackhole
```

Verify this static route is redistributed in the current configuration.

```
ICX-Tx-Core1(config-ospf-1)# show ip ospf lsdb external
OSPF Router with ID (10.x.0.2) (Process ID 1 VRF default)
============================================================

AS External Link State Advertisements
-------------------------------------

LSID            ADV Router      Age       Seq#       Checksum
------------------------------------------------------------
10.254.1.0      10.x.0.2       213       0x80000003 0x00009104
```

```
10.254.1.0      10.x.0.3        213         0x80000003 0x00008b09
198.18.11.0     10.x.0.2        213         0x80000003 0x0000a813
198.18.12.0     10.x.0.2        7           0x80000001 0x0000a11b
198.18.14.0     10.x.0.5        1661        0x80000002 0x0000773f
```

Access2

2. Review the default cost and metric type. On Access2, review the OSPF routing table, look for the routes **198.18.11.0/24** and **198.18.12.0/24**.

Note the (E2) marking:

```
ICX-Tx-Access2# show ip ospf routes 198.18.11.0/24
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 for prefixes 198.18.11.0/24
------------------------------------------------------------------------

Total Number of Routes : 1

198.18.11.0/24     (E2)
     via 10.x.102.2 interface vlan102, cost 25 distance 110
```

```
ICX-Tx-Access2# show ip ospf routes 198.18.12.0/24
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 for prefixes 198.18.12.0/24
------------------------------------------------------------------------

Total Number of Routes : 1

198.18.12.0/24     (E2)
     via 10.x.102.2 interface vlan102, cost 25 distance 110

ICX-T12-Access2#
```

Q: What is the default cost for the external route?

_____

A: The default cost is 25.

Define Route Map to Control Redistribution

Diagram L3

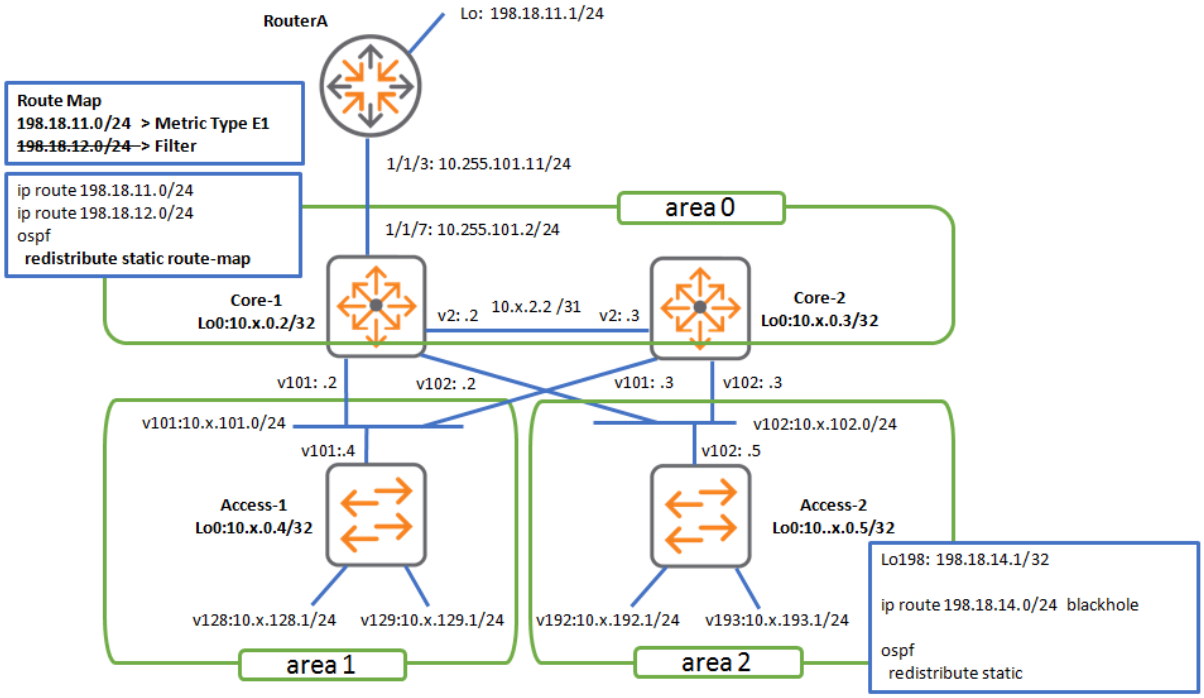

Now a route map will be created on Core1 to change the external route to metric type 1 for the 198.18.11.0/24 prefix. This will also prevent the route 198.18.12.0/24 from being redistributed.

Define Prefix-list

First a prefix-list must be defined, this defines which routes need to be controlled by the policy.

3. On Core1, define a prefix-list that matches the external route to RouterA. Any name can be used, try to use a descriptive name.

A prefix-list can contain many prefixes, and the administrator can control whether the entry should be matched (permit) or ignored (deny). This is why the order sequence (seq) is important if permit/deny combinations would be used.

In this example, only 1 entry is used.

```
ICX-Tx-Core1(config)# ip prefix-list routerAext seq 10 permit 198.18.11.0/24
```

Define Route Map

4. Next, define a route map that references the prefix list and sets the metric-type to External Type 1. Only routes that are permitted in the route policy are passing the route-map. There is an implicit deny at the end of the route-map.

Since the 198.18.12.0.24 does not match the prefix list and does not match any other rule in the route-map, it will not pass the route-map.

```
ICX-Tx-Core1(config)# route-map ospf_from_static permit seq 10
ICX-Tx-Core1(config-route-map-ospf_from_static-10)# match ip address prefix-list routerAext
ICX-Tx-Core1(config-route-map-ospf_from_static-10)#   set   metric-type   external type-1
```

5. Review the configuration.

```
ICX-Tx-Core1(config-route-map-ospf_from_static-10)# show run cur
ip prefix-list routerAext seq 10 permit 198.18.11.0/24 ge 24 le 24
!
!
!
!
route-map ospf_from_static permit seq 10
     match ip address prefix-list routerAext
     set metric-type external type-1
!
```

6. Exit the route-map context.

```
ICX-Tx-Core1(config-route-map-ospf_from_static-10)# exit
```

Activate Route map for redistributed routes

7. Activate the route policy on the OSPF redistribution for static routes.

```
ICX-Tx-Core1(config)# router ospf 1
ICX-Tx-Core1(config-ospf-1)# redistribute static route-map ospf_from_static
ICX-Tx-Core1(config-ospf-1)# exit
```

Access2

8. Verify the result on the Access2 router. The 198.18.11.0/24 route should now be listed as E1, and the cost will reflect the path to reach the Core1 (100) plus the initial cost (25 by default).

```
ICX-Tx-Access2# show ip ospf routes 198.18.11.0/24
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 for prefixes 198.18.11.0/24
------------------------------------------------------------------------

Total Number of Routes : 1
```

```
198.18.11.0/24     (E1)
     via 10.x.102.2 interface vlan102, cost 125 distance 110
```

9. The route to 198.18.12.0/24 should have been removed as well since it does not match the route policy.

```
ICX-Tx-Access2# show ip ospf routes 198.18.12.0/24
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 for prefixes 198.18.12.0/24
------------------------------------------------------------------------

Total Number of Routes : 0
```

Core1

10. Return to Core1 and modify the route policy so the initial cost is now 1000.

```
ICX-Tx-Core1(config)# route-map ospf_from_static permit seq 10
ICX-Tx-Core1(config-route-map-ospf_from_static-10)# set metric 1000
ICX-Tx-Core1(config-route-map-ospf_from_static-10)# exit
```

Access2

11. Verify the result on Access2. The cost should now be 1100, that is 1000 initial cost plus cost 100 to reach the Core1.

```
ICX-Tx-Access2(config)# show ip ospf routes 198.18.11.0/24
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 for prefixes 198.18.11.0/24
------------------------------------------------------------------------

Total Number of Routes : 1

198.18.11.0/24     (E1)
     via 10.x.102.2 interface vlan102, cost 1100 distance 110
```

This demonstrates how a route policy can be used to control external routes.

Task 4: Filter Routes with Stub and Totally Stub Areas

Objectives

In this task, the stub and totally stubby (sometimes referred to as ‘no-summary’) area types will be configured.

Stub

When an area does not have an ASBR inside the area, this means that all the traffic destined to external networks will need to go through the ABR systems.

The logic here is: why distribute all those external routes into the area, if the traffic has to go through the ABR systems anyway?

It would be convenient to 'aggregate' all those external routes using the default route (0.0.0.0/0), that is operating as the ultimate aggregation route.

This is exactly what the Stub area feature provides:

- The ABR will not announce the external Type 5 routes from the backbone into the stub area.

- The ABR will announce a default route (0.0.0.0/0) into the area.

- Routers inside the stub area will see the 0.0.0.0/0 route, but they will not see the external routes anymore.

- Routers inside the stub area will still see all the OSPF intra-area and inter-area (OSPF LSA Type3 Summary).

The stub area feature must be configured on every router in the area.

Stub No-summary

On top of the stub area feature, an administrator can configure the area to be a 'Stub No-summary' area. This extra 'no-summary' feature is only configured on the ABR systems.

Since the routers inside the stub area already have the default route 0.0.0.0/0, this default route could also be used to 'aggregate' all the OSPF inter-area routes (OSPF LSA Type 3 Summary routes). This results in the name 'stub **no-summary**'. This feature is also known as a 'totally stubby area'.

Steps

Make Area 1 an OSPF Stub Area

Diagram

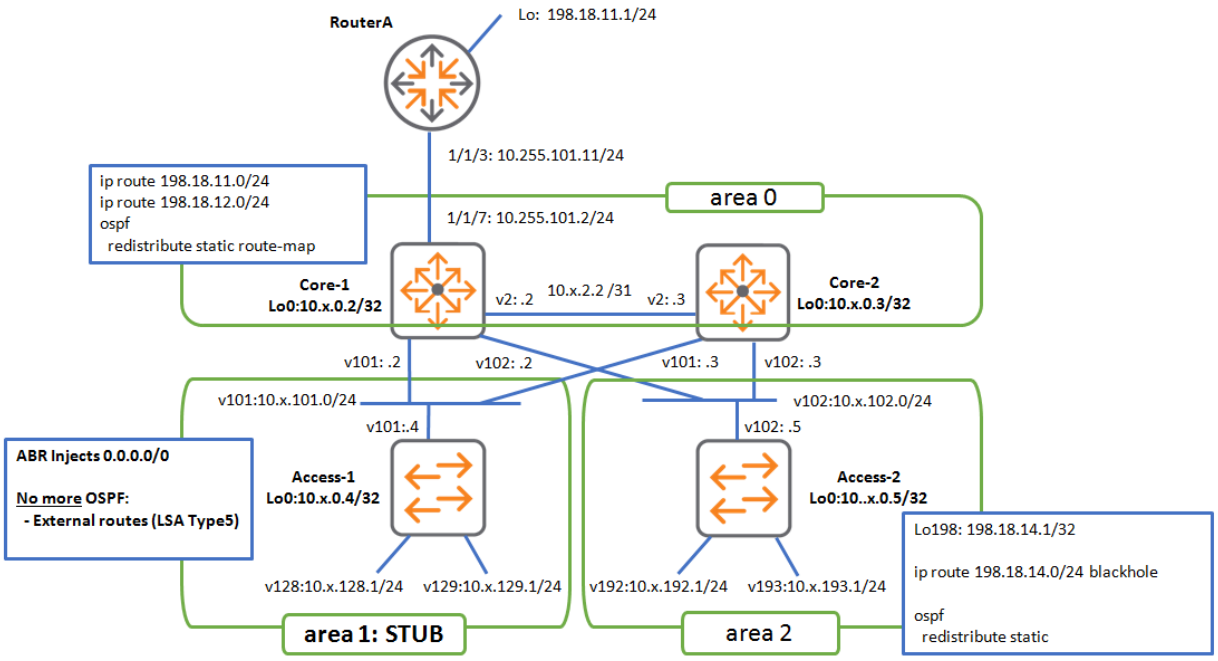

Access1

Review that Access1 in area 1 sees all the external routes.

1. On Access1, check the IP routing table. Notice that there is no default route needed at this point, since all external routes are actually known to the router. So there is no 0.0.0.0/0 entry at this moment.

```
ICX-Tx-Access1# show ip route 0.0.0.0

No ipv4 routes configured
```

Make area 1 a stub area. This will summarize all the external routes into the default route 0.0.0.0/0. This action is done at the ABR and should be done on **each** ABR. In this lab setup, Core1 and Core2 are part of a VSX setup, so making area 1 a stub area on Core1 will automatically be synchronized to Core2 due to the VSX OSPF synchronization.

Core1

2. Open Core1 and make area 1 a stub under the OSPF context.

```
ICX-Tx-Core1(config)# router ospf 1
ICX-Tx-Core1(config-ospf-1)# area 1 stub
```

3. Verify using the vsx-peer command that Core2 also has area 1 set as stub.

```
ICX-Tx-Core1(config-ospf-1)# show run ospf vsx-peer
router ospf 1
   router-id 10.x.0.3
   passive-interface default
   redistribute static route-map ospf_from_static
   area 0.0.0.0
   area 0.0.0.1 stub
   area 0.0.0.1 range 10.x.128.0/18 type inter-area
   area 0.0.0.2
   area 0.0.0.2 range 10.x.99.0/24 type inter-area no-advertise
   area 0.0.0.2 range 10.x.192.0/18 type inter-area
interface loopback 0
    ip ospf 1 area 0.0.0.0
interface vlan1
    ip ospf 1 area 0.0.0.0

<...output omitted...>

ICX-Tx-Core1(config-ospf-1)# exit
```

Access1

4. On Access1 router (router inside area 1), check the ip routing table.

```
ICX-Tx-Access1# show ip route ospf
```

Q: Are there any OSPF routes in the routing table?

_____

A: No, there are no more OSPF routes.

Q: Why did the OSPF routes disappear?

_____

A: The OSPF stub area option must be enabled on all routers in the area. The stub option is checked using the OSPF hello packets, a mismatch will prevent an OSPF neighbor adjacency.

5. Check the OSPF interface statistics. there should be dropped hello packets due to 'Options mismatch'.

```
ICX-Tx-Access1# show ip ospf statistics interface vlan101
...

Reason                         Packets Dropped
-----------------------------------------------
Invalid type                   0
Invalid length                 0
Invalid checksum               0
Invalid version                0
Bad or unknown source          8
Area mismatch                  0
Self-originated                0
Duplicate router ID            0
Interface standby              0
Total Hello packets dropped    62
  Network Mask mismatch        0
  Hello interval mismatch      0
  Dead interval mismatch       0
  Options mismatch             62
  MTU mismatch                 0
```

6. On Access1, change the area 1 type to stub.

```
ICX-Tx-Access1(config)# router ospf 1
ICX-Tx-Access1(config-ospf-1)# area 1 stub
ICX-Tx-Access1(config-ospf-1)# exit
```

7. Review the routing table, the external routes should have been replaced with a default route by the ABR.

NOTE: It may take a few moments for the OSPF adjacencies to re-establish. Within about 30 seconds the OSPF routes should appear.

```
ICX-Tx-Access1(config)# show ip route 0.0.0.0

Displaying ipv4 routes selected for forwarding

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

0.0.0.0/0, vrf default
        via  10.x.101.2,  [110/101],  ospf
        via  10.x.101.3,  [110/101],  ospf
<...output omitted...>
```

Q: What is the cost of the default route?

_____

A: The cost is 101. This is based on 100 to reach the Core1/Core2, the metric for the default route is 1 by default.

8. On Access1, check which route it would take for the traffic to 198.18.11.0. Since the specific route to 198.18.11.0 is no longer available, the default route is used.

```
ICX-Tx-Access1(config)# show ip route 198.18.11.0

Displaying ipv4 routes selected for forwarding

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

0.0.0.0/0, vrf default
        via  10.x.101.2,  [110/101],  ospf
        via  10.x.101.3,  [110/101],  ospf
```

Core1

9. On Core1, review the routing table, look for the default route 0.0.0.0/0.

```
ICX-Tx-Core1(config)# show ip route 0.0.0.0

No ipv4 routes configured
```

Q: Is there a default route on the Core1?

A: No, the ABR will announce the default route into the stub area, but it does not need the default route by itself, since it knows all the external routes.

This demonstrates how the 'stub' area feature of OSPF can automatically summarize all external routes with the default route.

Make Area 1 Stub No-Summary Area (Totally Stubby Area)

Diagram

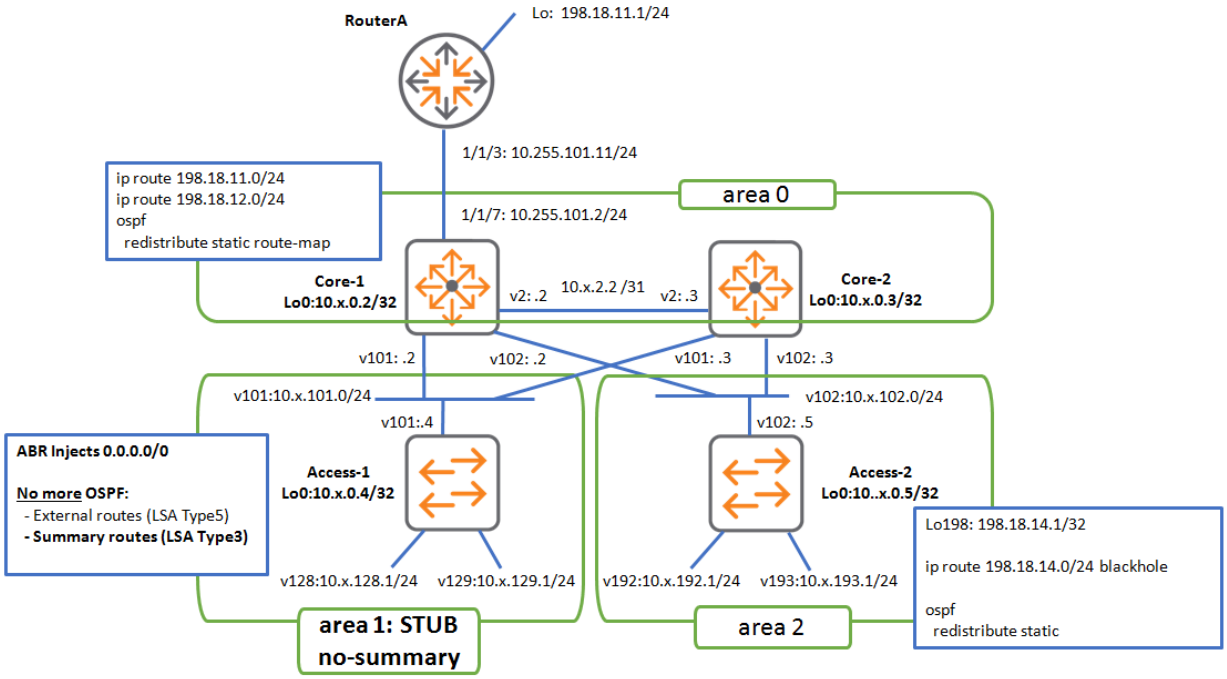

The stub feature in area 1 is summarizing all the external routes for this area into the default route, while it still sees all the detailed OSPF inter-area routes.

It is also possible to summarize all the OSPF inter-area routes into this default route, so all the OSPF summary LSAs (type3) will be removed as well.

Since the OSPF summary LSA is suppressed by this action, this feature is known as 'stub no-summary'.

First verify on Access1 that the OSPF inter-area routes are still visible.

Access1

10. On Access1, review the IP routing table, look for OSPF routes.

```
ICX-Tx-Access1(config)# show ip route ospf

Displaying ipv4 routes selected for forwarding

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

0.0.0.0/0, vrf default
        via  10.x.101.2,  [110/101],  ospf
        via  10.x.101.3,  [110/101],  ospf
10.x.0.3/32, vrf default
        via  10.x.101.3,  [110/100],  ospf
10.x.0.5/32, vrf default
        via  10.x.101.2,  [110/200],  ospf
```

```
        via  10.x.101.3,  [110/200],  ospf
10.x.0.2/32, vrf default
        via  10.x.101.2,  [110/100],  ospf
10.x.1.0/24, vrf default
        via  10.x.101.2,  [110/200],  ospf
        via  10.x.101.3,  [110/200],  ospf
10.x.2.2/31, vrf default
        via  10.x.101.2,  [110/101],  ospf
        via  10.x.101.3,  [110/101],  ospf
10.x.11.0/24, vrf default
        via  10.x.101.2,  [110/200],  ospf
        via  10.x.101.3,  [110/200],  ospf
10.x.12.0/24, vrf default
        via  10.x.101.2,  [110/200],  ospf
        via  10.x.101.3,  [110/200],  ospf
10.x.102.0/24, vrf default
        via  10.x.101.2,  [110/200],  ospf
        via  10.x.101.3,  [110/200],  ospf
10.x.192.0/18, vrf default
        via  10.x.101.2,  [110/300],  ospf
        via  10.x.101.3,  [110/300],  ospf
```

Now make the area a 'stub no-summary' area. This must only be done on the ABR, since that is the router that will filter the routes.

Core1

11. On Core1, change the area 1 to no-summary. This change should be made on every ABR, but due to the VSX synchronization, this is automatically pushed to Core2.

```
ICX-Tx-Core1(config)# router ospf 1
ICX-Tx-Core1(config-ospf-1)# area 1 stub no-summary
ICX-Tx-Core1(config-ospf-1)# exit
```

Access1

12. On Access1, verify the impact on the routing table. All OSPF external and inter-area routes have been replaced with the default route.

```
ICX-Tx-Access1(config)# show ip route ospf

Displaying ipv4 routes selected for forwarding

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

0.0.0.0/0, vrf default
        via  10.x.101.2,  [110/101],  ospf
        via  10.x.101.3,  [110/101],  ospf

ICX-Tx-Access1(config)#
```

13. Verify that Access1 can still reach the external network to RouterA and Access2.

```
ICX-Tx-Access1(config)# do ping 198.18.11.1
PING 198.18.11.1 (198.18.11.1) 100(128) bytes of data.
108 bytes from 198.18.11.1: icmp_seq=1 ttl=63 time=2.39 ms
<...output omitted...>
```

```
ICX-Tx-Access1(config)# do ping 198.18.14.1
PING 198.18.14.1 (198.18.14.1) 100(128) bytes of data.
108 bytes from 198.18.14.1: icmp_seq=1 ttl=63 time=0.183 ms
<...output omitted...>
```

14. Verify Access1 can still reach the loopback IP on Access2 (an inter-area route) .

```
ICX-Tx-Access1(config)# do ping 10.x.0.5
PING 10.x.0.5 (10.x.0.5) 100(128) bytes of data.
108 bytes from 10.x.0.5: icmp_seq=1 ttl=63 time=0.180 ms
<...output omitted...>
```

This demonstrates how the routing table of Access1 can benefit from the summarization while it still maintains full connectivity to the rest of the network.

Task 5: Filter Routes with a Not So Stubby Area (NSSA)

Diagram

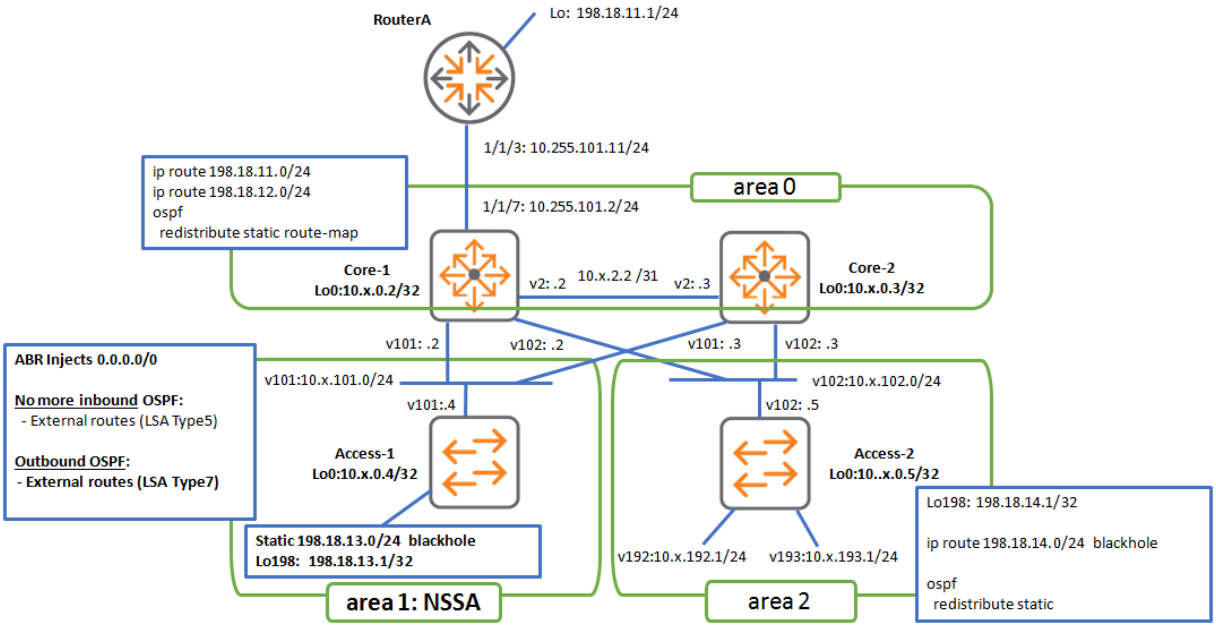

Objectives

Not-So-Stubby-Area (NSSA)

While a stub area benefits from the route summarization that is performed by the ABR, it is also limited now, since it cannot announce any external networks by itself anymore.

When the administrator attempts to redistribute a route on a router in a stub area, OSPF will simply not process the request.

Scenario

In these steps, Access1 will make a connection to an external network and the rest of the OSPF AS needs access to this external network as well.

At the same time, Access1 should still be optimized, so it does not get all the external routes from the backbone area.

Steps

Verify the problem.

Access1

1. On Access1, introduce a new external network. Since there is no real external network in this lab, it will be simulated using a blackhole static route and a local loopback interface. This is similar to the setup on Access2.

```
ICX-Tx-Access1(config)# ip route 198.18.13.0/24 blackhole
ICX-Tx-Access1(config)# interface loopback 198
ICX-Tx-Access1(config-loopback-if)# ip address 198.18.13.1/32
ICX-Tx-Access1(config-loopback-if)# exit
ICX-Tx-Access1(config)#
```

2. Now attempt to redistribute the static route into OSPF.

```
ICX-Tx-Access1(config)# router ospf 1
ICX-Tx-Access1(config-ospf-1)# redistribute static
ICX-Tx-Access1(config-ospf-1)# exit
```

3. Check the local OSPF LSDB, look for external routes.

```
ICX-Tx-Access1(config)# show ip ospf lsdb
OSPF Router with ID (10.x.0.4) (Process ID 1 VRF default)
============================================================

Router Link State Advertisements (Area 0.0.0.1)
-----------------------------------------------

LSID            ADV Router      Age       Seq#       Checksum       Link Count
-------------------------------------------------------------------------------
10.x.0.2       10.x.0.2       166       0x80000006 0x00008a35     1
10.x.0.3       10.x.0.3       2         0x80000009 0x00008237     1
10.x.0.4       10.x.0.4       1         0x8000000d 0x00005943     3

Network Link State Advertisements (Area 0.0.0.1)
------------------------------------------------

LSID            ADV Router      Age       Seq#       Checksum
-------------------------------------------------------------
10.x.101.3     10.x.0.3       2         0x80000002 0x0000cba8

Inter-area Summary Link State Advertisements (Area 0.0.0.1)
-----------------------------------------------------------

LSID            ADV Router      Age       Seq#       Checksum
-------------------------------------------------------------
0.0.0.0         10.x.0.2       1506      0x80000003 0x0000c87b
0.0.0.0         10.x.0.3       1506      0x80000003 0x0000c280
```

Q: Are there any external LSAs in the OSPF LSDB?

A: No, since this is a stub area, OSPF does not actually perform the redistribution. This means that the Core1 and Core2 would never be able to learn about this 198.18.13.0/24 route.

4. A filtered view 'external' on the LSDB shows the same result.

```
ICX-Tx-Access1(config)# show ip ospf lsdb external
No OSPF LSAs found on VRF default.

ICX-Tx-Access1(config)#
```

Solution: NSSA

The solution for this is to convert the area into an NSSA area type. This will allow external routes to be exchanged via the OSPF LSDB as LSA Type7 entries (since LSA type 5 is not allowed by a stub area).

The ABR will convert the Type7 LSAs into the regular Type5 LSAs in the backbone.

Any other areas in the OSPF autonomous system are not aware that this NSSA configuration was used, since they only see the external routes as LSA Type5.

5. On Access1, convert the area type to NSSA

```
ICX-Tx-Access1(config)# router ospf 1
ICX-Tx-Access1(config-ospf-1)# area 1 nssa
ICX-Tx-Access1(config-ospf-1)# exit
```

6. Review the OSPF LSDB. There should be an NSSA external LSA.

```
ICX-Tx-Access1(config)# show ip ospf lsdb
OSPF Router with ID (10.x.0.4) (Process ID 1 VRF default)
============================================================

Router Link State Advertisements (Area 0.0.0.1)
-----------------------------------------------

LSID            ADV Router      Age       Seq#       Checksum       Link Count
-------------------------------------------------------------------------------
10.x.0.4       10.x.0.4       30        0x80000001 0x0000928d     3

NSSA External Link State Advertisements (Area 0.0.0.1)
------------------------------------------------------

LSID            ADV Router      Age       Seq#       Checksum
-------------------------------------------------------------
198.18.13.0     10.x.0.4       30        0x80000001 0x000071c0
```

Q: Why are all the other LSAs gone from the LSDB?

A: The NSSA area type is also included in the OSPF hello. Since Core1 and Core2 still have the stub area type, Access1 no longer has an OSPF adjacency with them. This can be verified using 'show ip ospf neighbors'.

7. Notice that the NSSA external (LSA Type7) and regular external (LSA Type5) are shown separately in the LSDB.

```
ICX-Tx-Access1(config)# show ip ospf lsdb external
No OSPF LSAs found on VRF default.
```

```
ICX-Tx-Access1(config)# show ip ospf lsdb nssa-external
OSPF Router with ID (10.x.0.4) (Process ID 1 VRF default)
============================================================

NSSA External Link State Advertisements (Area 0.0.0.1)
------------------------------------------------------

LSID            ADV Router      Age       Seq#       Checksum
-------------------------------------------------------------
198.18.13.0     10.x.0.4        10        0x80000001 0x000071c0
```

Core1

8. On Core1, change the area 1 area-type to nssa.

```
ICX-Tx-Core1(config)# router ospf 1
ICX-Tx-Core1(config-ospf-1)# area 1 nssa
ICX-Tx-Core1(config-ospf-1)# exit
```

9. Wait a few moments for the Adjacencies to re-established. Verify that NSSA external LSAs are now appearing in the LSDB of area 1.

```
ICX-T12-Core1(config)# show ip ospf lsdb nssa-external area 1
OSPF Router with ID (10.x.0.2) (Process ID 1 VRF default)
============================================================

NSSA External Link State Advertisements (Area 0.0.0.1)
------------------------------------------------------

LSID            ADV Router      Age       Seq#       Checksum
-------------------------------------------------------------
0.0.0.0         10.x.0.2        93        0x80000001 0x00002891
0.0.0.0         10.x.0.3        89        0x80000001 0x00002296
198.18.11.0     10.x.0.2        93        0x80000001 0x00006407
198.18.13.0     10.x.0.4        221       0x80000001 0x000071c0
```

Verify the routes in another area, like area 2.

Access2

10. Now switch to Access2, verify the routing table. The route to the 198.18.13.0 should now be available.

```
ICX-Tx-Access2(config)# show ip route 198.18.13.0

Displaying ipv4 routes selected for forwarding

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

198.18.13.0/24, vrf default
        via  10.x.102.3,  [110/25],  ospf
        via  10.x.102.2,  [110/25],  ospf
```

This demonstrates how Area 1 still benefits from the route summarization of the stub area feature, while at the same time it can announce a route to an external network.

11. Review the OSPF LSDB of area 2, look for the 'external' and 'nssa-external' entries in the LSDB.

```
ICX-Tx-Access2# show ip ospf lsdb nssa-external
No OSPF LSAs found on VRF default.
```

```
ICX-T12-Access2# show ip ospf lsdb external
OSPF Router with ID (10.x.0.5) (Process ID 1 VRF default)
============================================================

AS External Link State Advertisements
--------------------------------------

LSID            ADV Router      Age       Seq#       Checksum
------------------------------------------------------------
198.18.11.0     10.x.0.2       1373      0x80000007 0x0000560f
198.18.13.0     10.x.0.3       193       0x80000001 0x0000ed4d
198.18.14.0     10.x.0.5       29        0x80000006 0x00006f43
```

Q: Why is the entry for the 198.18.13.0/24 network not shown as an NSSA-external route?

_____

A: The NSSA-external only has meaning ***inside*** the NSSA area. When the ABR receives these NSSA (LSA Type7) entries, it converts them into regular 'external' routes (LSA Type5) in the backbone area, so the backbone area and any other areas are not aware of this.

Revert Area 1 to a Normal Area

Area 1 will now be reverted to a normal area.

Core1

12. On Core1, remove the NSSA type for area 1.

```
ICX-Tx-Core1(config)# router ospf 1
ICX-Tx-Core1(config-ospf-1)# no area 1 nssa
ICX-Tx-Core1(config-ospf-1)# exit
```

Access1

13. On Access1, repeat this step.

```
ICX-Tx-Access1(config)# router ospf 1
ICX-Tx-Access1(config-ospf-1)# no area 1 nssa
ICX-Tx-Access1(config-ospf-1)# exit
```

14. On Access1, verify OSPF adjacency is back to both core switches.

```
ICX-Tx-Access1(config)# show ip ospf neighbors
OSPF Process ID 1 VRF default
==============================

Total Number of Neighbors: 2

Neighbor ID      Priority  State              Nbr Address       Interface
-------------------------------------------------------------------------
10.x.0.2         1         FULL/BDR           10.x.101.2        vlan101

10.x.0.3         1         FULL/DROther       10.x.101.3        vlan101
```

NOTE: It may take a few moments for the OSPF adjacencies to re-establish. Repeat the command until the **'FULL'** state is shown.

15. Verify that external routes are available again, check, for example, the OSPF route to 198.18.11.0/24 network.

```
ICX-Tx-Access1(config)# show ip route 198.18.11.0

Displaying ipv4 routes selected for forwarding

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

198.18.11.0/24, vrf default
        via  10.x.101.2,  [110/1100],  ospf
```

This concludes the NSSA task.

16. On each switch, save a checkpoint for the current lab.

```
copy run checkpoint icx-lab63
```

17. Make sure to execute the previous step on Core1, Core2, Access1 and Access2.

You've completed Lab 6.3!

Lab 07: BGP- Basic BGP Peering

Lab Diagram

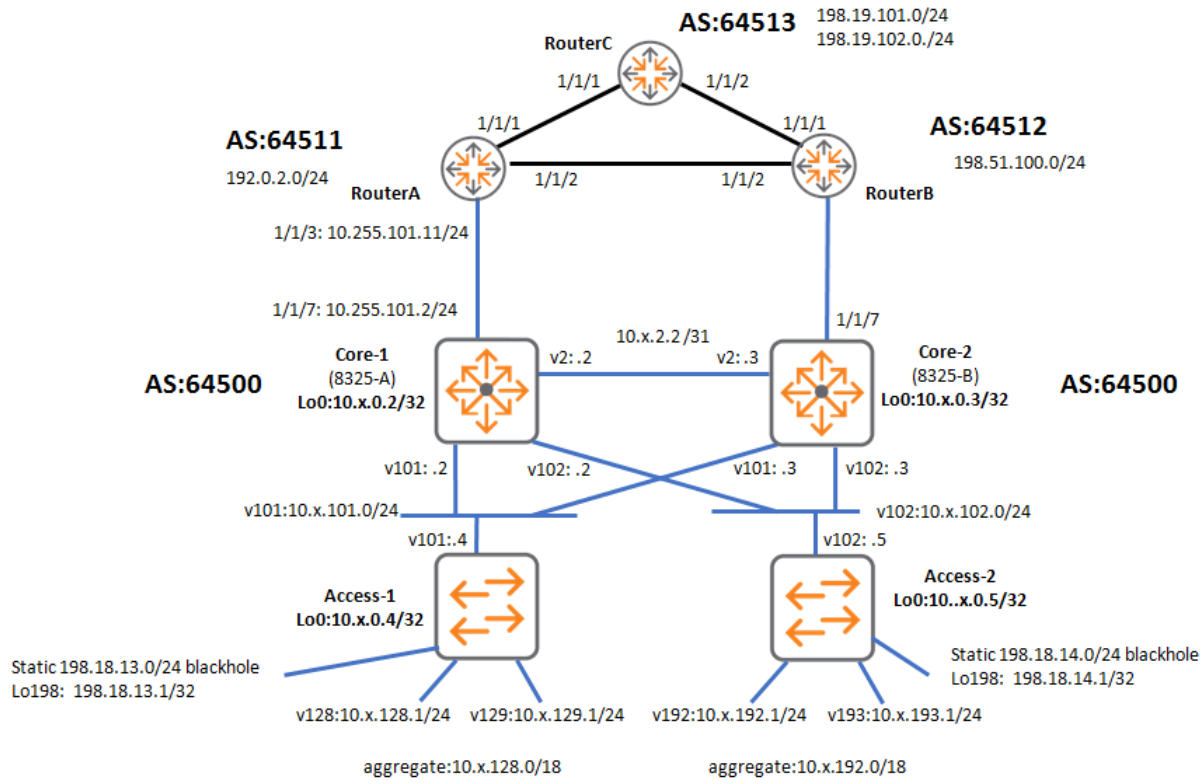

Overview

In this lab activity, basic BGP peering will be demonstrated. The two core switches will be configured to use iBGP between each other. They will use eBGP to communicate with two different ISPs, one connected to each Core switch.

Once the BGP peering has been established, a route advertisement will be configured.

The Core switches will also be configured with a route policy so they do not become a transit AS for traffic between the service providers.

Objectives

- Setup eBGP and iBGP peering
- Advertise routes using BGP
- Configure a route policy to control BGP routing updates

Task 1: Prepare the Lab Setup

This lab requires the completion of the OSPF lab activities (Single area, Multi-area and External routes).

RouterA

Load the checkpoint for bgp lab.

1. Connect to RouterA and login with username **admin**, password **aruba123** .
2. Load the BGP checkpoint to the running configuration. Verify that the 'bgp' checkpoint exists.

```
ICX-RouterA-ospf# show checkpoint list
ospf
bgp
startup-config
<...output omitted...>
```

3. Load the checkpoint to the running configuration.

```
ICX-RouterA-ospf# copy checkpoint bgp running-config
Configuration changes will take time to process, please be patient.
ICX-RouterA-ospf#
```

4. Press 'enter' a few times, until the updated hostname is displayed.

```
ICX-RouterA-ospf#
ICX-RouterA-bgp#
```

Task 2: Core1 eBGP Peering to ISP1

Diagram

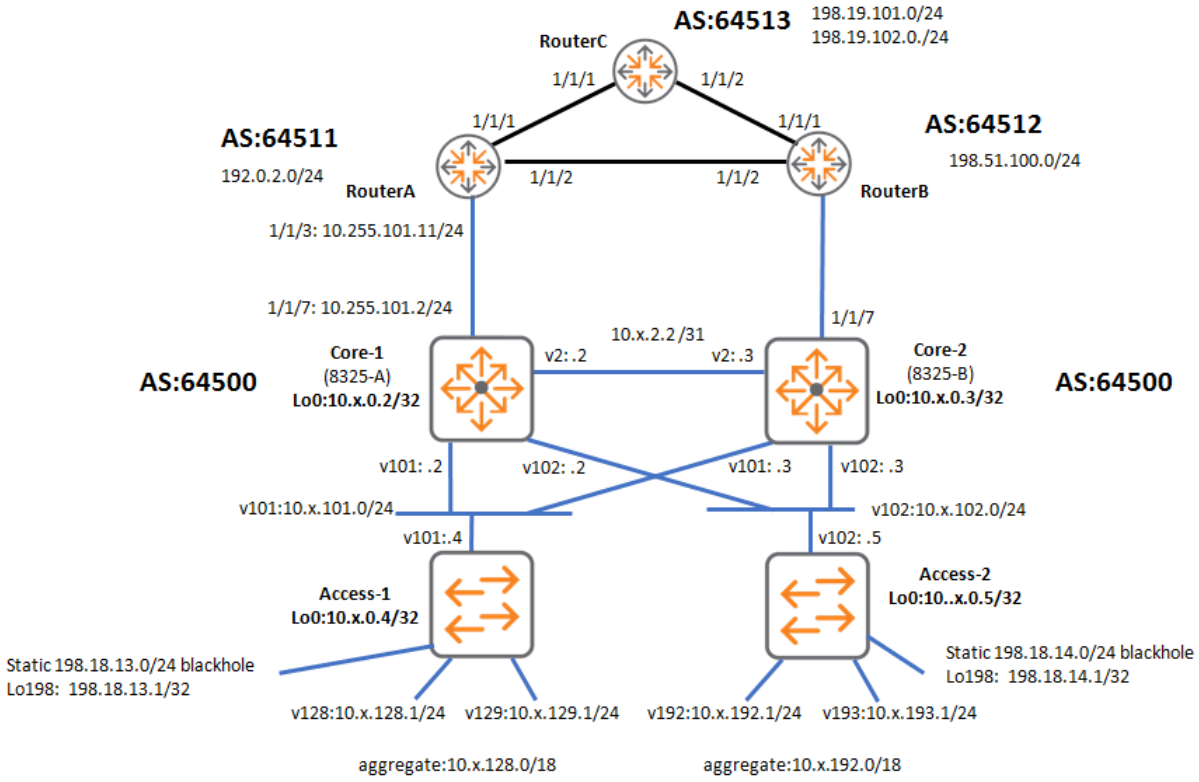

Objectives

In this task, Core1 will be configured to establish an eBGP peering with ISP RouterA. Since the RouterA is running as a VM, there is no direct physical link between the Core1 and RouterA, so it can take more time to detect that the peer is unreachable. This will be handled using BFD, Bidirectional Forward Detection, a simple keep-alive protocol.

Steps

Core1

1. Open a terminal connection to Core1 and enter configuration mode.
2. In the OSPF lab, a routed link was established to RouterA, verify Core1 can reach the Router.

```
ICX-Tx-Core1(config)# do ping 10.255.101.11
PING 10.255.101.11 (10.255.101.11) 100(128) bytes of data.
108 bytes from 10.255.101.11: icmp_seq=1 ttl=64 time=1.49 ms
108 bytes from 10.255.101.11: icmp_seq=2 ttl=64 time=1.79 ms
```

3. Enable BGP for AS 64500.

```
ICX-Tx-Core1(config)# router bgp 64500
```

4. Define RouterA as a BGP neighbor.

```
ICX-Tx-Core1(config-bgp)# neighbor 10.255.101.11 remote-as 64511
```

5. Enter the IPv4 address-family, enable the peer RouterA for IPv4.

```
ICX-Tx-Core1(config-bgp)# address-family ipv4 unicast
ICX-Tx-Core1(config-bgp-ipv4-uc)# neighbor 10.255.101.11 activate
ICX-Tx-Core1(config-bgp-ipv4-uc)# exit
ICX-Tx-Core1(config-bgp)# exit
```

6. Verify the status of the BGP session, it should be established. The state may be 'Idle' initially, wait for the state to change to 'Established'.

```
ICX-Tx-Core1(config)# show bgp all summary
VRF : default
BGP Summary
-----------
 Local AS               : 64500        BGP Router Identifier  : 10.x.0.2
 Peers                  : 1            Log Neighbor Changes   : No
 Cfg. Hold Time         : 180          Cfg. Keep Alive        : 60

Address-family : IPv4 Unicast
------------------------------
 Neighbor        Remote-AS MsgRcvd MsgSent   Up/Down Time State
AdminStatus
 10.255.101.11   64511       10      12      00h:03m:30s  Established   Up

Address-family : IPv6 Unicast
------------------------------

Address-family : L2VPN EVPN
------------------------------

ICX-Tx-Core1(config)#
```

Verify the Impact of the Default BGP Keepalive Interval (60 seconds)

7. On Core1, disable the interface that connects to the RouterA.

```
ICX-Tx-Core1(config)# interface 1/1/7
ICX-Tx-Core1(config-if)# shutdown
ICX-Tx-Core1(config-if)# exit
```

8. Review the status of the BGP neighbor using the command **'show bgp all summary**'.
9. Wait about 30 seconds and review the status again, it should still be 'Established', even when the link is down.

```
ICX-Tx-Core1(config)# show bgp all summary
```

10. **There is no need to wait** for the status to transition to 'Idle'. This output simply shows what the result would be after about three minutes (3 x 60 seconds).

```
ICX-Tx-Core1(config)# show bgp all summary
VRF : default
BGP Summary
-----------
 Local AS               : 64500        BGP Router Identifier  : 10.x.0.2
 Peers                  : 1            Log Neighbor Changes   : No
 Cfg. Hold Time         : 180          Cfg. Keep Alive        : 60

Address-family : IPv4 Unicast
------------------------------
 Neighbor        Remote-AS MsgRcvd MsgSent   Up/Down Time State
AdminStatus
 10.255.101.11   64511       12      18      00h:01m:06s  Idle          Up

Address-family : IPv6 Unicast
------------------------------

Address-family : L2VPN EVPN
------------------------------

ICX-Tx-Core1(config)#
```

This demonstrates that, by default, it can take several minutes before BGP detects that a peer is not reachable. In the next section a solution will be configured.

11. Enable the interface to the RouterA again.

```
ICX-Tx-Core1(config)# interface 1/1/7
ICX-Tx-Core1(config-if)# no shutdown
ICX-Tx-Core1(config-if)# exit
```

Use BFD for Faster Peer Keepalive Detection

12. On Core1, enter the BGP context and enable BFD for the RouterA neighbor. On RouterA, this command should also be enabled. In the lab setup, this has been done already.

```
ICX-Tx-Core1(config)# router bgp 64500
ICX-Tx-Core1(config-bgp)# neighbor 10.255.101.11 fall-over bfd
ICX-Tx-Core1(config-bgp)# exit
```

13. Verify the current BFD sessions.

```
ICX-Tx-Core1(config)# show bfd

Admin status: disabled
Statistics:
Total number of control packets transmitted: 0
Total number of control packets received: 0
Total number of control packets dropped: 0
Session Interface VRF       Source        Destination IP Echo    State Aplication
------- --------- --------- ------------------------------------------------------
1       1/1/7     default   10.255.101.2  10.255.101.11 enabled  admin_down   bgp
ICX-Tx-Core1(config)#
```

```
ICX-Tx-Core1(config)# show bfd session 1

BFD session information - Session 1
VRF: default
Min Tx interval (msec): 3000
Min Rx interval (msec): 3000
Min echo Rx interval (msec): 500
Detect multiplier: 5
Application: bgp
Local discriminator: 13211
Remote discriminator: 366
Echo: enabled
Local diagnostic: administratively_down
Remote diagnostic: no_diagnostic
State flaps: 0
Interface Source IP     Destination IP  State        Pkt Rx   Pkt Tx   Pkt drop
--------- ------------------------------------ ------------------------------
1/1/7     10.255.101.2  10.255.101.11   admin_down   27       26       0
ICX-Tx-Core1(config)#
```

Q: Has a BFD session been established to RouterA, and what would be the reason?

_____

A: No BFD session is active at the moment, since the BFD process is not enabled at the global level (admin down).

14. Enable BFD globally on Core1 and RouterA

```
ICX-Tx-Core1(config)# bfd
ICX-RouterA-bgp(config)# bfd
```

15. Review the state. The state will transition from admin_down to down> init > up.

```
ICX-Tx-Core1(config)# show bfd session 1

BFD session information - Session 1
VRF: default
Min Tx interval (msec): 3000
Min Rx interval (msec): 3000
Min echo Rx interval (msec): 500
Detect multiplier: 5
Application: bgp
Local discriminator: 13211
Remote discriminator: 366
Echo: enabled
Local diagnostic: no_diagnostic
Remote diagnostic: no_diagnostic
State flaps: 0
Interface Source IP     Destination IP State        Pkt Rx   Pkt Tx   Pkt drop
--------- ------------------------------------ ------------------------------
1/1/7     10.255.101.2  10.255.101.11  up           0        0        0
ICX-Tx-Core1(config)#
```

Verify the BFD Operation

16. Now repeat the test by bringing down the interface to RouterA.

```
ICX-Tx-Core1(config)# interface 1/1/7
ICX-Tx-Core1(config-if)# shutdown
ICX-Tx-Core1(config-if)# exit
```

17. Check the BGP neighbor state. Within a few seconds, the state should now transition to 'Idle'.

```
ICX-Tx-Core1(config)# show bgp all summary
VRF : default
BGP Summary
-----------
 Local AS               : 64500        BGP Router Identifier  : 10.x.0.2
 Peers                  : 1            Log Neighbor Changes   : No
 Cfg. Hold Time         : 180          Cfg. Keep Alive        : 60

Address-family : IPv4 Unicast
------------------------------
 Neighbor        Remote-AS MsgRcvd MsgSent   Up/Down Time State
AdminStatus
 10.255.101.11   64511       30      32      00h:00m:00s  Idle          Up

Address-family : IPv6 Unicast
------------------------------

Address-family : L2VPN EVPN
------------------------------

ICX-Tx-Core1(config)#
```

18. Enable the interface to RouterA again.

```
ICX-Tx-Core1(config)# interface 1/1/7
ICX-Tx-Core1(config-if)# no shutdown
ICX-Tx-Core1(config-if)# exit
```

This demonstrates how BFD can help with a faster detection of the peer reachability.

NOTE: For eBGP peers, the switch can also monitor if the link used to reach a directly adjacent peer goes down and then reset the BGP session immediately.
This can be enabled under the 'router bgp ' context using the command 'bgp fast-external-fallover'.

```
router bgp <as>
   bgp fast-external-fallover
```

Verify Received Routes via BGP

19. Now review the routes that have been received via BGP on Core1.

```
ICX-Tx-Core1(config)# show bgp all paths
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed
VRF : default
Local Router-ID 10.x.0.2

Address-family : IPv4 Unicast
------------------------------
    Network            Next Hop         Path
*>e 192.0.2.0/24       10.255.101.11    64511

*>e 198.19.101.0/24    10.255.101.11    64511 64513

*>e 198.19.102.0/24    10.255.101.11    64511 64512 64513

*>e 198.51.100.0/24    10.255.101.11    64511 64512

Total number of entries 4

Address-family : IPv6 Unicast
------------------------------
    Network            Next Hop         Path
Total number of entries 0

ICX-Tx-Core1(config)#
```

Q: Route 192.0.2.0/24 has AS-path 64511. What does this mean?

A: Since the neighbor 10.255.101.11 belongs to AS 64511, this means that route 192.0.2.0/24 is originated by this neighbor AS.

Q: Which AS originates the route 198.19.102.0/24?

_____

A: The last AS in the AS-path indicates the originating AS, so this route is originated by AS 64513, it traverses AS 64512 and then AS 64511.

Q: Both 198.19.101.0/24 and 198.19.102.0/24 originate from AS 64513. Why are they using a different path (different AS-path)?

_____

A: BGP allows the administrator to control the routing paths. In this case, apparently some route policy was defined on Routers A, B or C so that the 198.19.101.0/24 and 198.19.102.0/24 are taking a different path to reach their destination. You only see the result of that policy here.

Task 3: Core1 and Core2 iBGP Peering

Diagram

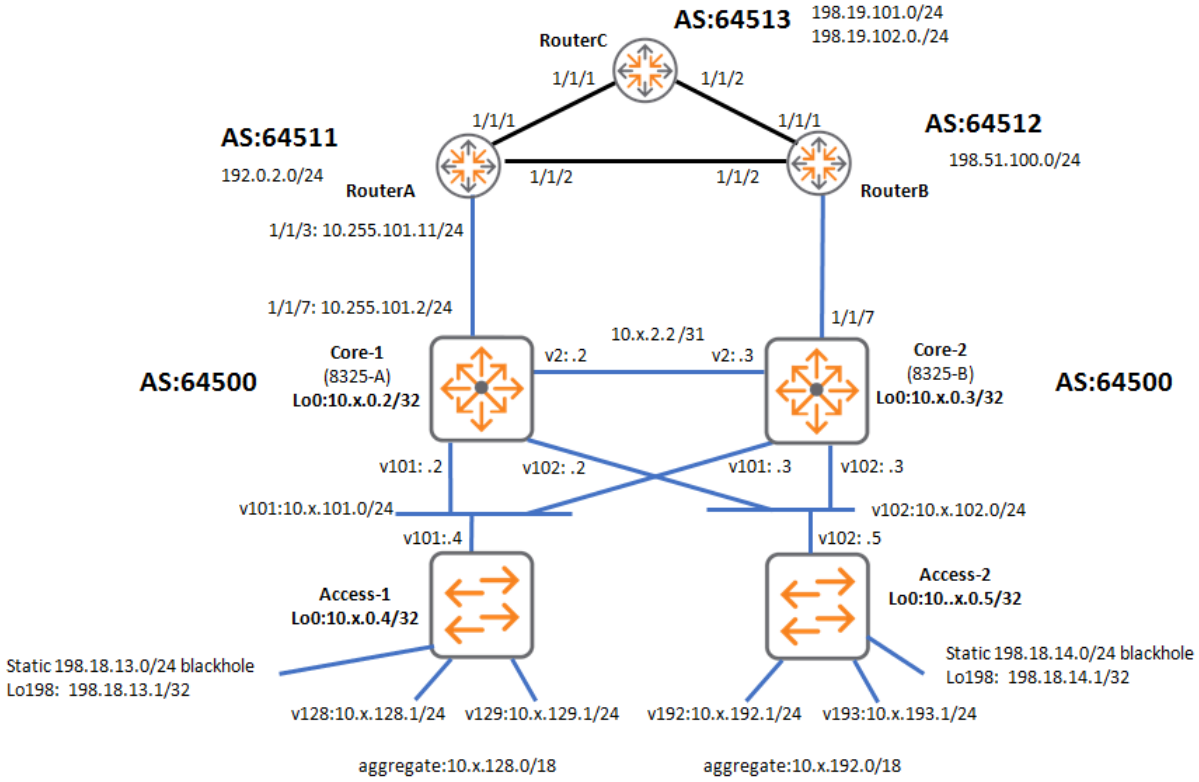

Objectives

In this task, iBGP peering will be configured between the Core1 and the Core2 devices.

The VSX synchronization for BGP will be reviewed and disabled, so each Core switch can have its own BGP configuration.

With the iBGP peering in place between the two core switches, the BGP routes and the next-hop will be analyzed. This will result in the configuration of the next hop self for the iBGP peers.

Steps

Setup iBGP Peering Between Core1 and Core2

Core1

1. On Core1, add a new neighbor for Core2, use the same AS, so BGP knows this is iBGP.

```
ICX-Tx-Core1(config)# router bgp 64500
ICX-Tx-Core1(config-bgp)# neighbor 10.x.0.3 remote-as 64500
```

2. For iBGP, the peering is typically configured between the Loopback IPs of the routers, instead of the outgoing interface IP. Configure the loopback IP as the source IP for this neighbor.

```
ICX-Tx-Core1(config-bgp)# neighbor 10.x.0.3 update-source loopback 0
```

3. Enable the neighbor for IPv4 unicast.

```
ICX-Tx-Core1(config-bgp)# address-family ipv4 unicast
ICX-Tx-Core1(config-bgp-ipv4-uc)# neighbor 10.x.0.3 activate
ICX-Tx-Core1(config-bgp-ipv4-uc)# exit
ICX-Tx-Core1(config-bgp)# exit
```

Core2

4. Switch to Core2, enter the configuration mode if needed, review the current configuration, look for the BGP section.

```
ICX-Tx-Core2(config)# show run bgp
router bgp 64500
    neighbor 10.x.0.3 remote-as 64500
    neighbor 10.x.0.3 update-source loopback 0
    neighbor 10.255.101.11 remote-as 64511
    neighbor 10.255.101.11 fall-over bfd
    address-family ipv4 unicast
        neighbor 10.x.0.3 activate
        neighbor 10.255.101.11 activate
    exit-address-family
!
ICX-Tx-Core2(config)#
```

Q: Why is there already a BGP configuration on Core2?

A: Since Core1 and Core2 are in a VSX cluster, and all the VSX synchronization features had been turned on, the BGP commands are automatically replicated.

This is very useful in case the VSX cluster is in a datacenter and needs to have the same neighbor peering commands on both Core switches.

In this lab environment, the BGP configuration of both Core1 and Core2 will be unique, so the VSX Synchronization will be disabled.

NOTE: In case only minor differences are needed, it is possible to exclude the neighbor configuration synchronization at the neighbor level in the bgp configuration. An example would be:

```
example(config)# router bgp 64500
example(config-bgp)# neighbor 10.1.1.1 vsx-sync-exclude
```

It is not needed to apply this command in the lab, since the entire BGP synchronization will be disabled.

Core1

5. On Core1, disable the VSX synchronization of the BGP feature.

```
ICX-Tx-Core1(config)# vsx
ICX-Tx-Core1(config-vsx)# no vsx-sync bgp
ICX-Tx-Core1(config-vsx)# exit
```

Core2

6. On Core2, verify the BGP configuration.

```
ICX-Tx-Core2(config)# show run bgp
```

Q: Is the BGP configuration still in the running configuration?

A: Yes, it is still there. Only the VSX synchronization was stopped. Now the administrator can make 'local' changes to the Core2.

7. To cleanup, remove the BGP configuration.

```
ICX-Tx-Core2(config)# no router bgp 64500
This will delete all BGP configurations on this device.
Continue (y/n)? y
```

8. Now try to setup the BGP router by yourself, add the Core1 as an iBGP neighbor, use the Loopback IP as the source IP, verify the peering.

9. The commands are only listed here in case you want to verify your commands.

```
ICX-Tx-Core2(config)# router bgp 64500
ICX-Tx-Core2(config-bgp)# neighbor 10.x.0.2 remote-as 64500
ICX-Tx-Core2(config-bgp)# neighbor 10.x.0.2 update-source loopback 0
ICX-Tx-Core2(config-bgp)# address-family ipv4 unicast
ICX-Tx-Core2(config-bgp-ipv4-uc)# neighbor 10.x.0.2 activate
ICX-Tx-Core2(config-bgp-ipv4-uc)# exit-address-family
ICX-Tx-Core2(config-bgp)#
```

10. The result should be an 'Established' iBGP session.

```
ICX-Tx-Core2(config-bgp)# show bgp all summary
VRF : default
BGP Summary
-----------
 Local AS               : 64500        BGP Router Identifier  : 10.x.0.3
 Peers                  : 1            Log Neighbor Changes   : No
 Cfg. Hold Time         : 180          Cfg. Keep Alive        : 60

Address-family : IPv4 Unicast
------------------------------
 Neighbor        Remote-AS MsgRcvd MsgSent   Up/Down Time State
AdminStatus
 10.x.0.2       64500       7       3       00h:00m:37s  Established   Up

Address-family : IPv6 Unicast
------------------------------

Address-family : L2VPN EVPN
------------------------------

ICX-Tx-Core2(config-bgp)#
```

Verify Received Routes

11. Now review the received BGP routes on Core2.

```
ICX-Tx-Core2(config-bgp)# show bgp all path
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed
VRF : default
Local Router-ID 10.x.0.3

Address-family : IPv4 Unicast
------------------------------
    Network            Next Hop         Path
* i 192.0.2.0/24       10.255.101.11    64511

* i 198.19.101.0/24    10.255.101.11    64511 64513

* i 198.19.102.0/24    10.255.101.11    64511 64512 64513
```

```
* i 198.51.100.0/24    10.255.101.11    64511 64512

Total number of entries 4

Address-family : IPv6 Unicast
------------------------------
    Network            Next Hop         Path
Total number of entries 0
```

Q: Do you notice a difference in the flags compared to the output on Core1?

_____

A: There is no 'best' route on Core2, so the '>' flag is missing. In the next steps this will be investigated.

12. Review the IP routing table on Core2. Is there a route to the 198.19.101.0/24 network?

```
ICX-Tx-Core2(config-bgp)# show ip route 198.19.101.0

No ipv4 routes configured
```

Q: What would be the reason that this route that is visible in the BGP routing table is not inserted into the IP routing table?

_____

A: BGP always checks if the next-hop IP address of a route is reachable. In this case, the Core1 announces the external next-hop IP 10.255.101.11. This IP address is not reachable for Core2, so it cannot inject these routes into the IP routing table.

```
ICX-Tx-Core2(config-bgp)# show ip route 10.255.101.11

No ipv4 routes configured
```

Solution: iBGP next-hop-self

The behavior seen in the previous section occurs because, by default, iBGP keeps the next hop of the eBGP peer.

In many cases, this is not desired, so the internal iBGP IP should be used as the next hop for the announced IP routes.

This will now be configured on both Core1 and on Core2 on the neighbor definitions to each other.

Core1

13. Open Core1, enter the IPv4 address-family under the BGP router context, and enable the 'next-hop-self' for the Core2 neighbor.

```
ICX-Tx-Core1(config)# router bgp 64500
ICX-Tx-Core1(config-bgp)# address-family ipv4 unicast
ICX-Tx-Core1(config-bgp-ipv4-uc)# neighbor 10.x.0.3 next-hop-self
ICX-Tx-Core1(config-bgp-ipv4-uc)# exit
```

14. Review the configuration.

```
ICX-Tx-Core1(config-bgp)# show running-config current-context
router bgp 64500
    neighbor 10.x.0.3 remote-as 64500
    neighbor 10.x.0.3 update-source loopback 0
    neighbor 10.255.101.11 remote-as 64511
    neighbor 10.255.101.11 fall-over bfd
    address-family ipv4 unicast
        neighbor 10.x.0.3 activate
        neighbor 10.x.0.3 next-hop-self
        neighbor 10.255.101.11 activate
    exit-address-family
```

Core2

15. On Core2, repeat this for the neighbor Core1.

```
ICX-Tx-Core2(config-bgp)# address-family ipv4 unicast
ICX-Tx-Core2(config-bgp-ipv4-uc)# neighbor 10.x.0.2 next-hop-self
ICX-Tx-Core2(config-bgp-ipv4-uc)# exit
ICX-Tx-Core2(config-bgp)#
```

16. On Core2, review the BGP routing table, check the next-hop IP for the routes.

```
ICX-Tx-Core2(config-bgp)# show bgp all path
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed
VRF : default
Local Router-ID 10.x.0.3

Address-family : IPv4 Unicast
------------------------------
    Network            Next Hop         Path
* i 192.0.2.0/24       10.255.101.11    64511

* i 198.19.101.0/24    10.255.101.11    64511 64513
```

```
* i 198.19.102.0/24    10.255.101.11    64511 64512 64513

* i 198.51.100.0/24    10.255.101.11    64511 64512

Total number of entries 4
```

Q: Did the next-hop change for the routes?

A: No, this change is only effective for the next update, so the session must be refreshed for these changes to be visible.

17. On Core2, clear the BGP session with Core1.

```
ICX-Tx-Core2(config-bgp)# do clear bgp 10.x.0.2
```

18. Review the BGP routing table, the next hop should now be the Loopback IP of the Core1. Also notice that the '>' best flag is now shown again for the routes.

```
ICX-Tx-Core2(config-bgp)# show bgp all path
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed
VRF : default
Local Router-ID 10.x.0.3

Address-family : IPv4 Unicast
------------------------------
    Network            Next Hop         Path
*>i 192.0.2.0/24       10.x.0.2        64511

*>i 198.19.101.0/24    10.x.0.2        64511 64513

*>i 198.19.102.0/24    10.x.0.2        64511 64512 64513

*>i 198.51.100.0/24    10.x.0.2        64511 64512
```

19. Review the IP routing table, are the routes visible now?

```
ICX-Tx-Core2(config-bgp)# show ip route 198.19.101.0

Displaying ipv4 routes selected for forwarding

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

198.19.101.0/24, vrf default
        via  10.x.2.2,  [200/0],  bgp
```

```
ICX-Tx-Core2(config-bgp)#
```

This demonstrates the need for next-hop-self between iBGP peers.

NOTE: Do not attempt to ping the remote hosts. The Core1 and Core2 are receiving the routes from the ISP, but they are not announcing any routes themselves, so there is no return path to make a connection possible at this point in the lab.

Task 4: Core2 eBGP Peering to ISP2

Diagram

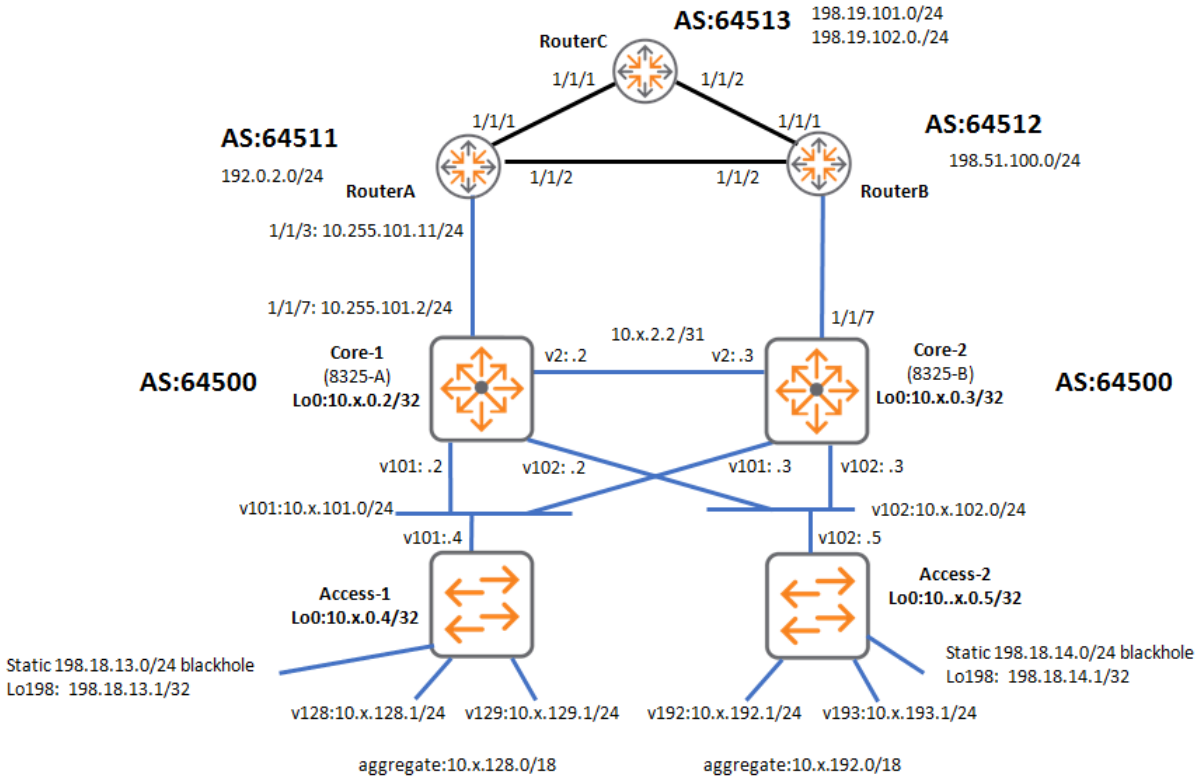

Objectives

In this task, Core2 will be configured with an eBGP peering session to RouterB (another ISP).

Once the peering is established, the routes will be reviewed.

This will also reveal that the customer now has become a transit AS, so traffic between RouterA and RouterB could be using the path via the customer network.

A solution using a route map will be configured to prevent this transit AS behavior.

Steps

Core2

1. On Core2, configure the routed uplink to RouterB.

```
ICX-Tx-Core2(config)# interface 1/1/7
ICX-Tx-Core2(config-if)# ip address 10.255.102.3/24
ICX-Tx-Core2(config-if)# no shutdown
ICX-Tx-Core2(config-if)# exit
```

2. Verify Core2 can reach RouterB on the new link.

```
ICX-Tx-Core2(config)# do ping 10.255.102.12
PING 10.255.102.12 (10.255.102.12) 100(128) bytes of data.
108 bytes from 10.255.102.12: icmp_seq=1 ttl=64 time=1.43 ms
108 bytes from 10.255.102.12: icmp_seq=2 ttl=64 time=1.23 ms
```

> **NOTE:** If the RouterB is not reachable, check the interface configuration again. If the configuration is correct, contact your instructor.

3. On Core2, attempt to configure RouterB (**10.255.102.12**) as a BGP neighbor for **AS 64512** by yourself. Only use the commands below in case you want to verify your commands.

```
ICX-Tx-Core2(config)# bfd
ICX-Tx-Core2(config)# router bgp 64500
ICX-Tx-Core2(config-bgp)# neighbor 10.255.102.12 remote-as 64512
ICX-Tx-Core2(config-bgp)# neighbor 10.255.102.12 fall-over bfd
ICX-Tx-Core2(config-bgp)# address-family ipv4 unicast
ICX-Tx-Core2(config-bgp-ipv4-uc)# neighbor 10.255.102.12 activate
ICX-Tx-Core2(config-bgp-ipv4-uc)# exit
```

Verify the Received Routes from the eBGP Peer

4. On Core2, verify the status of the eBGP peering. The state with 10.255.102.12 should be **'Established'.**

```
ICX-Tx-Core2(config-bgp)# show bgp all summary
VRF : default
BGP Summary
-----------
 Local AS               : 64500        BGP Router Identifier  : 10.x.0.3
 Peers                  : 2            Log Neighbor Changes   : No
 Cfg. Hold Time         : 180          Cfg. Keep Alive        : 60

Address-family : IPv4 Unicast
------------------------------
 Neighbor        Remote-AS MsgRcvd MsgSent   Up/Down Time State
AdminStatus
 10.x.0.2       64500       30      27      00h:19m:53s  Established   Up
 10.255.102.12   64512       16      18      00h:00m:30s  Established   Up

Address-family : IPv6 Unicast
------------------------------
```

```
Address-family : L2VPN EVPN
-----------------------------

ICX-Tx-Core2(config-bgp)#
```

5. Review that routes have been received from this eBGP peer.

```
ICX-Tx-Core2(config-bgp)# show bgp all paths
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed
VRF : default
Local Router-ID 10.x.0.3

Address-family : IPv4 Unicast
-----------------------------
    Network            Next Hop         Path
*>i 192.0.2.0/24       10.x.0.2        64511

* e 192.0.2.0/24       10.255.102.12    64512 64511

*>i 198.19.101.0/24    10.x.0.2        64511 64513

* e 198.19.101.0/24    10.255.102.12    64512 64511 64513

*>e 198.19.102.0/24    10.255.102.12    64512 64513

*>e 198.51.100.0/24    10.255.102.12    64512

Total number of entries 6

Address-family : IPv6 Unicast
-----------------------------
    Network            Next Hop         Path
Total number of entries 0
```

Transit AS problem: Verify the Routes That Are Advertised to the eBGP Peer

By default, BGP will advertise all routes that it has received from other BGP peers. The risk for a customer with BGP peering to two different ISPs, is that it becomes a transit AS between these ISPs.

It is best practice to apply a route filter to the eBGP peers to ensure that only locally originated routes can be announced to them.

First, review that Core2 is indeed advertising routes to the RouterB.

6. On Core2, no configuration has been performed to announce routes to the eBGP system. Check if any routes are being advertised to the eBGP peer.

```
ICX-Tx-Core2(config-bgp)# show bgp ipv4 unicast neighbors 10.255.102.12
advertised-routes
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
```

```
              i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default
Local Router-ID 10.x.0.3

    Network           Nexthop         Metric     LocPrf     Weight Path
*>e 192.0.2.0/24      10.255.102.3    0          0          0       64500 64511
i
*>e 198.19.101.0/24   10.255.102.3    0          0          0       64500 64511
64513 i
Total number of entries 2

ICX-Tx-Core2(config-bgp)#
```

This shows that the 2 routes that were received by RouterA (via Core1), are being advertised by Core2 to RouterB.

Core1

7. On Core1, the issue can also be verified. It is currently advertising routes that it received via Core2 from RouterB.

```
ICX-Tx-Core1(config-bgp)# show bgp ipv4 unicast neighbors 10.255.101.11
advertised-routes
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default
Local Router-ID 10.x.0.2

    Network           Nexthop         Metric     LocPrf     Weight Path
*>e 198.19.102.0/24   10.255.101.2    0          0          0       64500 64512
64513 i
*>e 198.51.100.0/24   10.255.101.2    0          0          0       64500 64512
i
Total number of entries 2

ICX-Tx-Core1(config)#
```

Solution: Route Policy for Locally Originated Networks

In the next section, a route policy will be created to ensure that only routes that have originated in the local AS can be advertised to the eBGP peers. This will be achieved by checking the AS-Path attribute.

All the routes that originate inside this AS, will have an empty AS-Path attribute, while any route that was received via an external AS (eBGP), will have at least the value of that external AS in the AS-Path.

To check the AS-Path of a route, a regular expression can be used. See https://www.regex101.com for regular expression examples.

These regex values are used:

- '^' indicates 'must begin with'
- '$' indicates 'must end with'

Combined, they result in '^$', which represents an empty string.

Since the outbound route filter will be used by both Core1 and Core2, the configuration will be performed on Core1, so the route-map will by synchronized by VSX (**'route-map**' is a separate synchronization feature of VSX).

8. On Core1, define a new AS path list that only allows an 'empty' AS path value.

```
ICX-Tx-Core1(config)# ip aspath-list bgp-local permit ^$
```

> **NOTE:** In the remote lab console web interface, it may be required to enter the **'^'** and **add a <SPACE>** so see the character.

9. Define a new route-map to select routes matching the new AS path list.

```
ICX-Tx-Core1(config)# route-map bgp-ebgp-out permit
ICX-Tx-Core1(config-route-map-bgp-ebgp-out-10)#  match aspath bgp-local
ICX-Tx-Core1(config-route-map-bgp-ebgp-out-10)#  exit
ICX-Tx-Core1(config)#
```

The outbound policy must be defined on each Core separately.

Core1

10. On Core1, apply the policy to the neighbor definition to RouterA.

```
ICX-Tx-Core1(config)# router bgp 64500
ICX-Tx-Core1(config-bgp)# address-family ipv4 unicast
ICX-Tx-Core1(config-bgp-ipv4-uc)# neighbor 10.255.101.11 route-map bgp-ebgp-out
out
ICX-Tx-Core1(config-bgp-ipv4-uc)# exit
```

Core2

11. On Core2, the route-map was synchronized by VSX. Verify that the route-map exists in the configuration.

```
ICX-Tx-Core2(config-bgp)# show route-map
Route-map: bgp-ebgp-out
   Seq 10, permit,
     Match :
       aspath-list                 : bgp-local
     Set :

Route-map: ospf_from_static
   Seq 10, permit,
```

```
    Match :
      ip prefix-list              : routerAext
    Set :
      metric                      : 1000
      metric-type                 : external_type1
```

12. On Core2, apply the policy to the neighbor definition to Router B. The switch should still be in the 'router bgp 64500' context.

```
ICX-Tx-Core2(config-bgp)# address-family ipv4 unicast
ICX-Tx-Core2(config-bgp-ipv4-uc)# neighbor 10.255.102.12 route-map bgp-ebgp-out
out
ICX-Tx-Core2(config-bgp-ipv4-uc)# exit
ICX-Tx-Core2(config-bgp)#
```

The update in the policy will be effective after the BGP sessions are restarted.

Core1

13. Reset the eBGP session on Core1.

```
ICX-Tx-Core1(config-bgp)# do clear bgp 10.255.101.11
```

Core2

14. Reset the eBGP session on Core2.

```
ICX-Tx-Core2(config-bgp)# do clear bgp 10.255.102.12
```

Wait a few moments so the eBGP session can be re-established.

Now the route-policy should be effective.

Core1

15. On Core1, check the outbound routes.

```
ICX-Tx-Core1(config-bgp)# show bgp ipv4 unicast neighbors 10.255.101.11
advertised-routes

Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default
Local Router-ID 10.x.0.2

    Network            Nexthop         Metric     LocPrf     Weight Path
Total number of entries 0

ICX-Tx-Core1(config-bgp)# exit
```

Core2

16. Verify this on Core2 as well.

```
ICX-Tx-Core2(config-bgp)# show bgp ipv4 unicast neighbors 10.255.102.12
advertised-routes

Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default
Local Router-ID 10.x.0.3

    Network            Nexthop         Metric     LocPrf     Weight Path
Total number of entries 0

ICX-Tx-Core2(config-bgp)# exit
```

This demonstrates how a route policy can prevent the customer from becoming a transit AS for other BGP autonomous systems.

Task 5: Announce Routes to eBGP Peers

Diagram

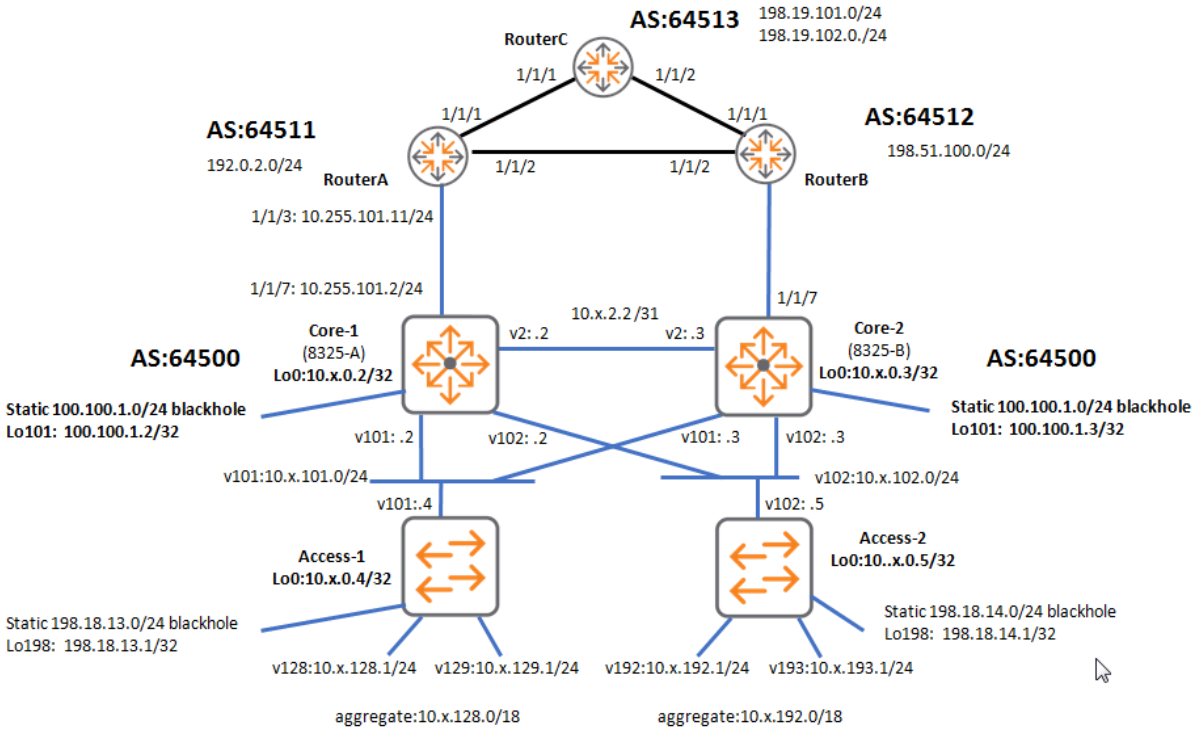

Objectives

In this task, the customer would like to advertise its local, 'public' subnet to the 2 ISPs.

First, a 'public' subnet will be defined on the Core1 and Core2 systems.

Next, Core1 will be configured to announce this subnet using eBGP to RouterA (ISP 64511) and the advertisement will be verified using RouterC.

The last step will be to repeat this configuration on Core2, the route is also advertised to RouterB (ISP 64512).

Steps

Define the Subnet on Core1 and Core2 for External Use

The Core1 and Core2 BGP AS 64500 will be hosting the **100.100.1.0/24** network.

Each Core switch will be configured with a loopback interface in this subnet for testing purposes.

This loopback will be included in the OSPF topology so Core1 and Core2 can reach each other's 100.100.1.0/24 interface and also forward data to the peer switch.

Core1

1. On Core1, define the new loopback with the public test IP.

```
ICX-Tx-Core1(config)# interface loopback 101
ICX-Tx-Core1(config-loopback-if)# ip address 100.100.1.2/32
ICX-Tx-Core1(config-loopback-if)# ip ospf 1 area 0
ICX-Tx-Core1(config-loopback-if)# exit
ICX-Tx-Core1(config)#
```

Core2

2. On Core2, define the new loopback IP as well.

```
ICX-Tx-Core2(config)# interface loopback 101
ICX-Tx-Core2(config-loopback-if)# ip address 100.100.1.3/32
ICX-Tx-Core2(config-loopback-if)# ip ospf 1 area 0
ICX-Tx-Core2(config-loopback-if)# exit
```

3. On Core2, verify the new loopback IP on Core1 can be reached via OSPF, while using the local loopback as source IP.

```
ICX-Tx-Core2(config)# do ping 100.100.1.2 source 100.100.1.3
PING 100.100.1.2 (100.100.1.2) from 100.100.1.3 : 100(128) bytes of data.
108 bytes from 100.100.1.2: icmp_seq=1 ttl=64 time=0.197 ms
108 bytes from 100.100.1.2: icmp_seq=2 ttl=64 time=0.169 ms
```

eBGP on Core1: Announce Local Subnet to RouterA

In these steps, Core1 will announce the 100.100.1.0/24 prefix to the RouterA.

Core1

4. On Core1, enter the BGP router context.

```
ICX-Tx-Core1(config)# router bgp 64500
ICX-Tx-Core1(config-bgp)# address-family ipv4 unicast
```

5. Use the 'network' command to inject the 100.100.1.0/24 route into the IPv4 BGP table.

```
ICX-Tx-Core1(config-bgp-ipv4-uc)# network 100.100.1.0/24
ICX-Tx-Core1(config-bgp-ipv4-uc)# exit
ICX-Tx-Core1(config-bgp)# exit
```

6. Verify if the new network is part of the IPv4 unicast routing table.

```
ICX-Tx-Core1(config)# show bgp ipv4 unicast paths
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed
VRF : default
Local Router-ID 10.x.0.2

    Network            Next Hop         Path
*>e 192.0.2.0/24       10.255.101.11    64511

*>e 198.19.101.0/24    10.255.101.11    64511 64513

*>i 198.19.102.0/24    10.x.0.3        64512 64513

* e 198.19.102.0/24    10.255.101.11    64511 64512 64513

*>i 198.51.100.0/24    10.x.0.3        64512

* e 198.51.100.0/24    10.255.101.11    64511 64512

Total number of entries 6
```

Q: Is the new route visible?

A: No, the new route is not injected yet. BGP requires the route to exist in the local routing table before it can be announced to BGP peers.

Since the local routing table only contains the /32 loopback routes, but not the /24 route, BGP rejects the route. To solve this, a dummy static route can be created. This should

only be done in case more specific subnets are still available in the routing table, such as the /32 subnets in this lab.

7. Define the dummy static route.

```
ICX-Tx-Core1(config)# ip route 100.100.1.0/24 blackhole
```

8. Next, verify if the route appears in the local BGP routing table.

```
ICX-Tx-Core1(config)# show bgp ipv4 unicast paths
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed
VRF : default
Local Router-ID 10.x.0.2

    Network             Next Hop         Path
*>  100.100.1.0/24      0.0.0.0
*>e 192.0.2.0/24        10.255.101.11    64511

*>e 198.19.101.0/24     10.255.101.11    64511 64513

*>i 198.19.102.0/24     10.x.0.3        64512 64513

* e 198.19.102.0/24     10.255.101.11    64511 64512 64513

*>i 198.51.100.0/24     10.x.0.3        64512

* e 198.51.100.0/24     10.255.101.11    64511 64512

Total number of entries 7
```

9. Verify if the route passes the outbound route-map to the RouterA.

```
ICX-Tx-Core1(config)# show bgp ipv4 unicast neighbors 10.255.101.11 advertised-
routes
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default
Local Router-ID 10.x.0.2

    Network             Nexthop          Metric     LocPrf     Weight Path
*>e 100.100.1.0/24      10.255.101.2     0          0          0       64500 i
Total number of entries 1

ICX-Tx-Core1(config)#
```

10. Verify connectivity using this IP to the RouterC (this router is connected behind RouterA and RouterB). If this router can be reached using a ping, it means RouterC knows about the 100.100.1.0/24 subnet as well.

RouterC has several loopback test IP addresses, one of them is 198.19.101.50.

```
ICX-Tx-Core1(config)# do ping 198.19.101.50 source 100.100.1.2
PING 198.19.101.50 (198.19.101.50) from 100.100.1.2 : 100(128) bytes of data.
108 bytes from 198.19.101.50: icmp_seq=1 ttl=62 time=3.82 ms
108 bytes from 198.19.101.50: icmp_seq=2 ttl=62 time=3.55 ms
108 bytes from 198.19.101.50: icmp_seq=3 ttl=62 time=3.37 ms
108 bytes from 198.19.101.50: icmp_seq=4 ttl=62 time=3.38 ms
```

Core2

11. On Core2, verify RouterC (198.19.101.50) can be reached when using the local public IP as the source IP (100.100.1.3). The reply confirms that all the hosts on the 100.100.1.0/24 subnet can be reached from the ISPs using a path via Core1.

```
ICX-Tx-Core2(config)# do ping 198.19.101.50 source 100.100.1.3
PING 198.19.101.50 (198.19.101.50) from 100.100.1.3 : 100(128) bytes of data.
108 bytes from 198.19.101.50: icmp_seq=1 ttl=62 time=3.67 ms
108 bytes from 198.19.101.50: icmp_seq=2 ttl=62 time=3.82 ms
108 bytes from 198.19.101.50: icmp_seq=3 ttl=62 time=3.46 ms
108 bytes from 198.19.101.50: icmp_seq=4 ttl=62 time=3.40 ms
108 bytes from 198.19.101.50: icmp_seq=5 ttl=62 time=3.95 ms

--- 198.19.101.50 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4005ms
rtt min/avg/max/mdev = 3.407/3.663/3.953/0.210 ms
```

12. Now repeat the configuration of the route advertisement on Core2.

```
ICX-Tx-Core2(config)# ip route 100.100.1.0/24 blackhole
ICX-Tx-Core2(config)# router bgp 64500
ICX-Tx-Core2(config-bgp)# address-family ipv4 unicast
ICX-Tx-Core2(config-bgp-ipv4-uc)# network 100.100.1.0/24
ICX-Tx-Core2(config-bgp-ipv4-uc)# exit
ICX-Tx-Core2(config-bgp)# exit
```

Route validation on RouterC

13. Open a connection to RouterC (**admin**/**aruba123**).

14. On RouterC, verify that there are now 2 paths for the 100.100.1.0/24 subnet.

```
ICX-RouterC-bgp# show bgp ipv4 unicast
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default
Local Router-ID 203.0.113.50

    Network           Nexthop         Metric   LocPrf     Weight Path
*>e 100.100.1.0/24    10.255.104.11   0        100        0      64511 64500 i
* e 100.100.1.0/24    10.255.105.12   0        100        0      64512 64500 i
*>e 192.0.2.0/24      10.255.104.11   0        100        0      64511 i
```

```
* e 192.0.2.0/24       10.255.105.12   0       100       0       64512 64511 i
*>  198.19.101.0/24    0.0.0.0         0       100       0       i
*>  198.19.102.0/24    0.0.0.0         0       100       0       i
* e 198.51.100.0/24    10.255.104.11   0       100       0       64511 64512 i
*>e 198.51.100.0/24    10.255.105.12   0       100       0       64512 i
*>  203.0.113.0/24     0.0.0.0         0       100       0       i
Total number of entries 9
```

Q: What is the AS path for each of the routes?

_____

A: There should be 1 entry with as-path '64511 64500' and another entry with '64512 64500'.

Q: Why is only 1 route selected as the best path ('>')?

_____

A: BGP default path selection will only select 1 best path when the AS-path length is the same for multiple paths.

15. Review the IP routing table for the 100.100.1.0/24 route. Notice that the next hop is in line with the BGP routing table shown in the previous step.

```
ICX-RouterC-bgp# show ip route 100.100.1.0

Displaying ipv4 routes selected for forwarding

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

100.100.1.0/24, vrf default
        via  10.255.104.11,  [20/0],  bgp
```

You've completed Lab 7!

Lab 08 IP IGMP Snooping- IP Multicast Snooping

Lab Diagram

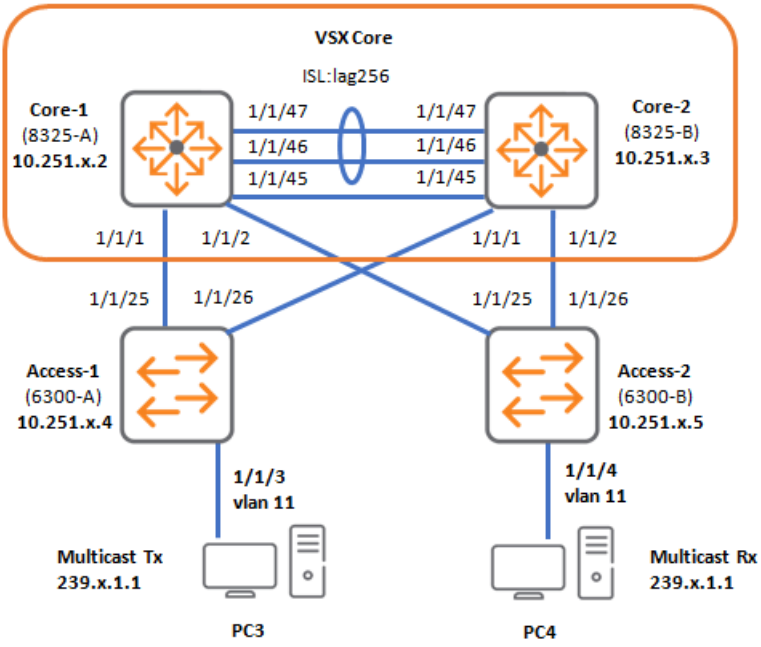

Overview

In this lab activity, IGMP Querier and IGMP snooping will be configured and demonstrated.

The Core VSX will be configured to perform the IGMP Querier function, since the core has the IP interface in the user VLAN.

The Access switches will be configured with IGMP snooping. This will ensure that the multicast traffic is only delivered to the ports with active subscribers.

The PC3, connected to Access1, will be transmitting some test multicasts, the PC4, connected to Access2, will be configured to register and listen for the test multicast traffic.

Objectives

- Configure IGMP Querier on the VSX system
- Configure IGMP Snooping on the Access switches
- Understand the need for IGMP configuration
- Verify the IGMP operation

Task 1: Prepare the Lab Start Configuration

Objectives

– This lab is built on the base VSX topology.

– Make sure to complete these steps to get the base VSX checkpoint configuration on the devices.

Steps (Required)

1. Open a console connection to the 6300A. Login using admin/aruba123.

```
ICX-Tx-Access1# copy checkpoint icx-lab04-vsx running-config
Configuration changes will take time to process, please be patient.
ICX-Tx-Access1#
```

2. Open a console connection to the 6300B. Login using admin/aruba123.

```
ICX-Tx-Access2# copy checkpoint icx-lab04-vsx running-config
Configuration changes will take time to process, please be patient.
ICX-Tx-Access2#
```

3. Open a console connection to the 8325A. Login using admin/aruba123.

```
ICX-Tx-Core1# copy checkpoint icx-lab04-vsx running-config
Configuration changes will take time to process, please be patient.
ICX-Tx-Core1#
```

4. Open a console connection to the 8325B. Login using admin/ aruba123.

```
ICX-Tx-Core2# copy checkpoint icx-lab04-vsx running-config
Configuration changes will take time to process, please be patient.
ICX-Tx-Core2#
```

Access1

5. On Access1, enter configuration mode, assign port 1/1/3 (connected to PC3) as access port to VLAN 11 and enable the port.

```
ICX-Tx-Access1(config)# interface 1/1/3
ICX-Tx-Access1(config-if)# vlan access 11
ICX-Tx-Access1(config-if)# no shutdown
ICX-Tx-Access1(config-if)# exit
```

PC3

6. Open a connection to PC3, reset the Lab NIC (disable/enable).
7. Verify the PC3 has an IP address from VLAN 11 (10.x.11.0/24).

Access2

8. On Access2, enter configuration mode, assign port 1/1/4 (connected to PC4) as access port to VLAN 11 and enable the port.

```
ICX-Tx-Access2(config)# interface 1/1/4
ICX-Tx-Access2(config-if)# vlan access 11
```

```
ICX-Tx-Access2(config-if)# no shutdown
ICX-Tx-Access2(config-if)# exit
```

PC4

1. Open a connection to PC4, reset the Lab NIC (disable/enable).
2. Verify the PC4 has an IP address from VLAN 11 (10.x.11.0/24).

Task 2: Setup the Multicast Sender and Receiver

Objectives

In this task, the default behavior of the switches will be demonstrated for intra-VLAN IP multicast traffic.

The PC3 (connected to Access1 1/1/3) will be transmitting test multicast traffic.

On the PC4 (connected to Access2 1/1/4), Wireshark will be used to verify that the traffic is flooded to all the ports on the network, even without any IGMP joins.

Steps

1. Open the PC3 (connected to Access1).
2. Start the UDP Multicast test tool from the desktop.
3. Under the 'Sender' section (top half of the screen), configure:
   – **Local Interface address**: Enter your local Lab NIC IP (in 10.x.11.0/24 subnet)

(do not confuse this with 'Local Multicast Interface Address')

NOTE: You may also drag and drop your 10.x.11.y IP from the 'Local interfaces' list to the 'Local Interface' field. See screenshot below.

– **Destination multicast address**: 239.x.1.1

– Click the **'Start Sender'** button to start.

NOTE: You may need to scroll down to see the 'Start Sender' button.

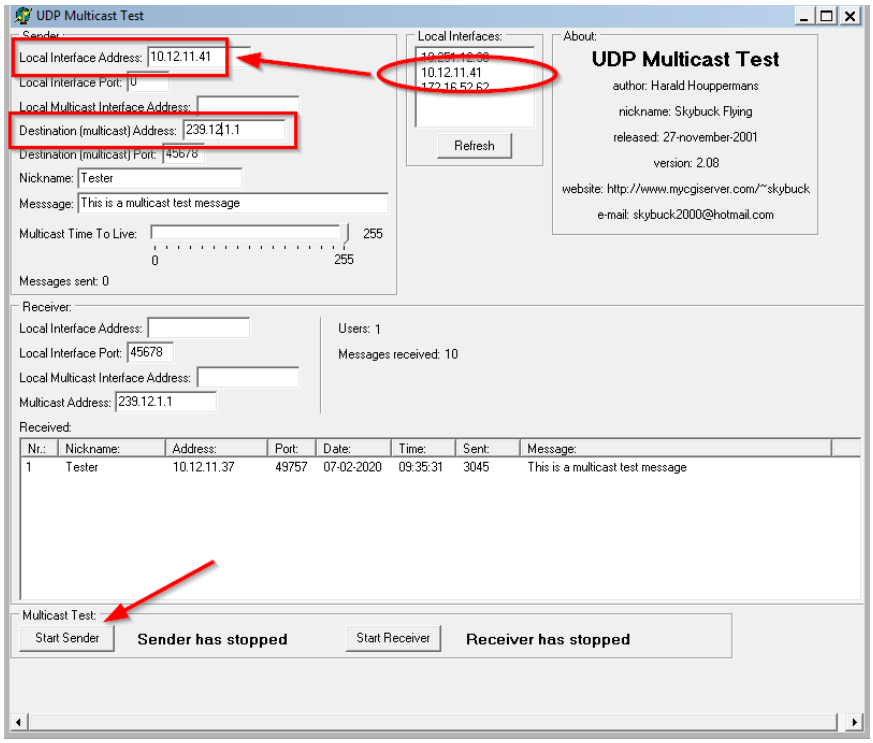

On PC3, locally verify the transmitted multicast traffic

4. On PC3, open Wireshark and select the 'Lab NIC' and start the capture (double-click).

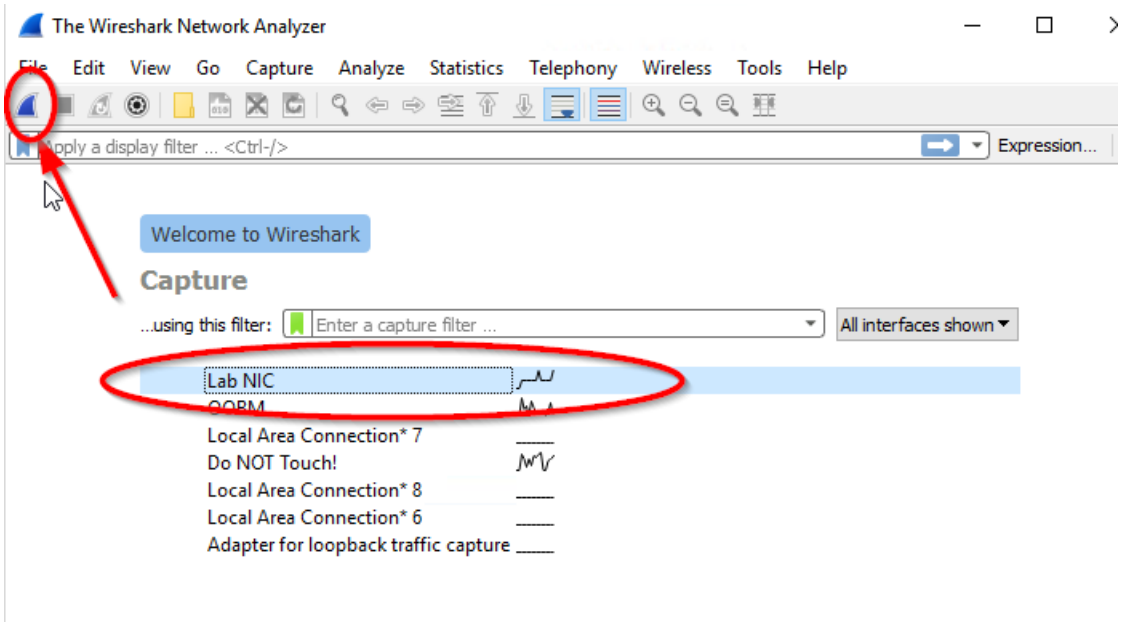

5. After about 5 seconds, stop the capture again.

6. Verify that packets are transmitted to the 239.x.1.1 multicast IP address.

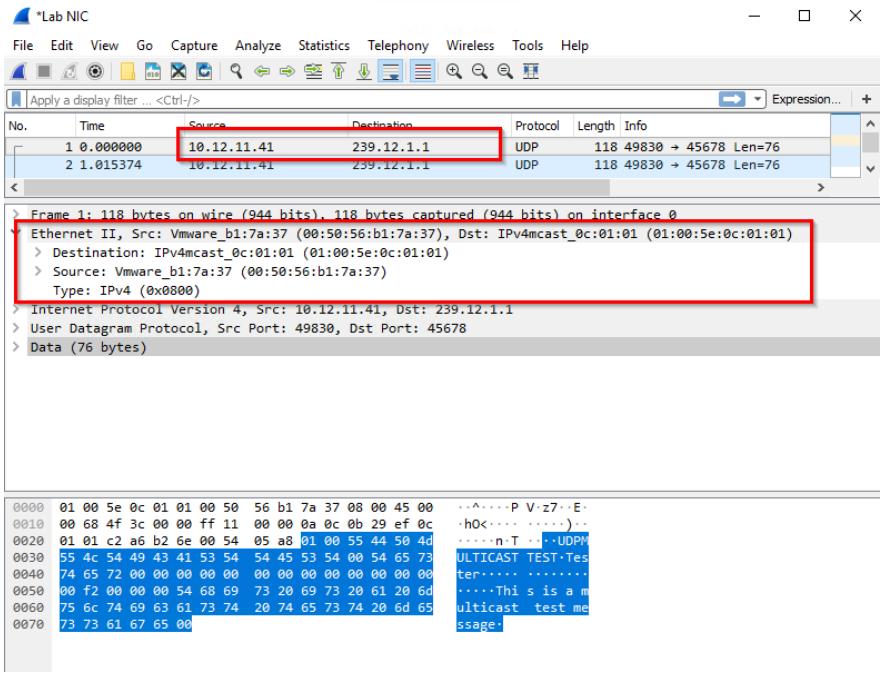

7. Open PC4 (connected to Access2), start Wireshark, and perform a packet capture on the 'Lab NIC' for about 5 seconds.
8. Verify that the multicast traffic is received on the 'Lab NIC'. If there is too much other traffic, apply this display filter to limit the display to IP multicast traffic.

ip.addr == 224.0.0.0/4

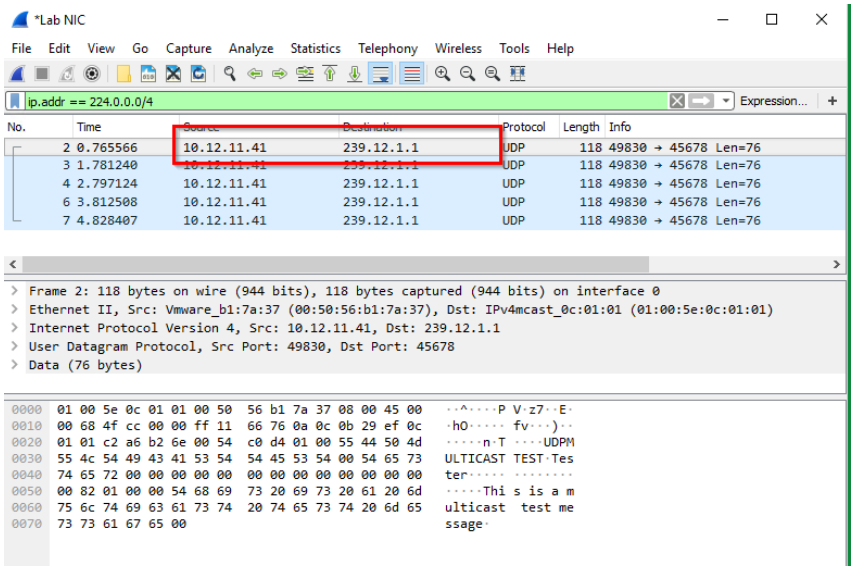

9. This demonstrates that multicast traffic is flooded by default.

10. On PC3, stop the multicast transmission, so you can start the next task without multicast traffic.

Task 3: Enable IGMP Querier and Snooping

Objectives

In this task, the network switches will be configured with IP IGMP snooping.

IP IGMP snooping will look for IP IGMP joins that are sent by the multicast listener devices.

To operate correctly, there should be an IP IGMP querier on the network. This function will be configured on the VSX Core devices.

Steps

Core1 - IGMP Querier and Snooping

1. Open a terminal connection to Core1, enter the configuration mode.
2. Review the default IP IGMP querier state.

```
ICX-Tx-Core1(config)# show ip igmp interface vlan11

IGMP is not enabled
```

3. Enable the querier function for VLAN 11 and VLAN 12.

```
ICX-Tx-Core1(config)# interface vlan11
ICX-Tx-Core1(config-if-vlan)# ip igmp enable
ICX-Tx-Core1(config-if-vlan)# exit
```

NOTE: The 'ip igmp querier' role is default on the Layer3 VLAN SVI interface, so it is only required to enable IGMP on the VLAN interface.

4. Repeat this for VLAN12.

```
ICX-Tx-Core1(config)# interface vlan12
ICX-Tx-Core1(config-if-vlan)# ip igmp enable
ICX-Tx-Core1(config-if-vlan)# exit
```

5. Verify that IP IGMP querier is now in the 'initial wait' state.

```
ICX-Tx-Core1(config)# show ip igmp interface vlan11

IGMP Configured Version  : 3
IGMP Operating Version   : 3
Querier State              : Initial Wait
Querier IP                 :
Querier Uptime             :
Querier Expiration Time    :
IGMP Snoop Enabled on VLAN : False
```

Q: Based on this output, what is the default IGMP snooping status on the VLAN?

_____

A: By default, IGMP snooping is disabled on the VLAN. Therefore, the IGMP querier and snooping functions are enabled separately. The IGMP Querier is enabled on the L3 VLAN Interface (interface vlan x), while the IGMP snooping is configured on the L2 VLAN context (vlan x).

6. Enable IGMP snooping on VLAN 11 and 12.

```
ICX-Tx-Core1(config)# vlan 11,12
ICX-Tx-Core1(config-vlan-<11,12>)# ip igmp snooping enable
ICX-Tx-Core1(config-vlan-<11,12>)# exit
```

7. Verify the updated state.

```
ICX-Tx-Core1(config)# show ip igmp interface vlan11

IGMP Configured Version  : 3
IGMP Operating Version   : 3
Querier State              : Initial Wait
Querier IP                 :
Querier Uptime             :
Querier Expiration Time    :
IGMP Snoop Enabled on VLAN : True

Active Group Address   Vers Mode Uptime    Expires
---------------------- ---- ---- --------- ---------
239.255.255.250        3    EXC  2m 55s    1m 25s
```

NOTE: The multicast group 239.255.255.250 may be observed at this point. This depends on the Windows client that may or may not have used this address already at this point.

Core2 - IGMP Querier and Snooping

8. Open a terminal connection to Core2, enter the configuration mode.
9. Enable the querier function for VLAN 11 and VLAN 12.

```
ICX-Tx-Core2(config)# interface vlan11
ICX-Tx-Core2(config-if-vlan)# ip igmp enable
ICX-Tx-Core2(config-if-vlan)# exit
```

```
ICX-Tx-Core2(config)# interface vlan12
ICX-Tx-Core2(config-if-vlan)# ip igmp enable
```

```
ICX-Tx-Core2(config-if-vlan)# exit
```

10. Enable IGMP snooping on the Layer2 vlans 11 and 12.

```
ICX-Tx-Core2(config)# vlan 11,12
ICX-Tx-Core2(config-vlan-<11,12>)# ip igmp snooping enable
ICX-Tx-Core2(config-vlan-<11,12>)# exit
```

11. Review the status of IGMP.

```
ICX-Tx-Core2(config)# show ip igmp interface vlan 11

IGMP Configured Version  : 3
IGMP Operating Version   : 3
Querier State              : Initial Wait
Querier IP                 :
Querier Uptime             :
Querier Expiration Time    :
IGMP Snoop Enabled on VLAN : True

Active Group Address   Vers Mode Uptime    Expires
---------------------- ---- ---- --------- ---------
239.255.255.250        3    EXC  2m 45s    1m 35s
```

NOTE: Depending on your lab pace, the switches may already have transitioned from 'Initial Wait' to 'Querier' or 'Non-Querier' at this point.

Enable IP IGMP Snooping on Access1 and Access2

Access1

12. Open a terminal connection to Access1, enter the configuration mode.
13. Enable IGMP snooping for VLAN 11 and VLAN 12.

```
ICX-Tx-Access1(config)# vlan 11,12
ICX-Tx-Access1(config-vlan-<11,12>)# ip igmp snooping enable
ICX-Tx-Access1(config-vlan-<11,12>)# exit
```

14. Verify the configuration status of IGMP snooping.

```
ICX-Tx-Access1(config)# show ip igmp snooping vlan 11
IGMP Snooping Protocol Info

Total VLANs with IGMP enabled             : 2
Current count of multicast groups joined  : 0

IGMP Drop Unknown Multicast               : Global
VLAN ID : 11
VLAN Name : VLAN11
IGMP Configured Version : 3
IGMP Operating Version : 3
```

```
Querier Address :
Querier Port :
Querier UpTime :
Querier Expiration Time :
```

Access2

15. Open a terminal connection to Access2, enter the configuration mode.
16. Enable IGMP snooping for VLAN 11 and VLAN 12.

```
ICX-Tx-Access2(config)# vlan 11,12
ICX-Tx-Access2(config-vlan-<11,12>)# ip igmp snooping enable
ICX-Tx-Access2(config-vlan-<11,12>)# exit
```

Verify the IGMP Querier status on the Core1.

It may take a few minutes before Core1 transitions to the IP Querier state. Make sure to wait in the lab until this has occurred.

Core1

17. On Core1, verify the state until the 'initial wait' has changed to 'querier'.

```
ICX-Tx-Core1(config)# show ip igmp interface vlan11

IGMP Configured Version  : 3
IGMP Operating Version   : 3
Querier State              : Querier
Querier IP [this switch]   : 10.x.11.2
Querier Uptime             : 7m 17s
Querier Expiration Time    : 0m 5s
IGMP Snoop Enabled on VLAN : True
```

Core2

18. Repeat this on Core2, verify it has transitioned to 'Non-Querier'.

```
ICX-Tx-Core2(config)# show ip igmp interface vlan 11

IGMP Configured Version  : 3
IGMP Operating Version   : 3
Querier State              : Non-Querier
Querier IP                 : 10.x.11.2
Querier Uptime             : 9m 17s
Querier Expiration Time    : 2m 20s
IGMP Snoop Enabled on VLAN : True
```

Now that the Core switch is sending out regular queries for IGMP membership, the Access switches IGMP snooping should be able to detect the 'querier' port.

Access1

19. On Access1, verify that the upstream IGMP querier was detected on the lag255 (this is the LAG to the VSX core).

```
ICX-Tx-Access1(config)# show ip igmp snooping vlan 11
IGMP Snooping Protocol Info

Total VLANs with IGMP enabled             : 2
Current count of multicast groups joined  : 1

IGMP Drop Unknown Multicast               : Global
VLAN ID : 11
VLAN Name : VLAN11
IGMP Configured Version : 3
IGMP Operating Version : 3
Querier Address : 10.x.11.2
Querier Port : lag255
Querier UpTime :5m 33s
Querier Expiration Time :2m 54s
```

Access2

20. Repeat this verification on Access2.

```
ICX-Tx-Access2(config)# show ip igmp snooping vlan 11
IGMP Snooping Protocol Info

Total VLANs with IGMP enabled             : 2
Current count of multicast groups joined  : 1

IGMP Drop Unknown Multicast               : Global
VLAN ID : 11
VLAN Name : VLAN11
IGMP Configured Version : 3
IGMP Operating Version : 3
Querier Address : 10.x.11.2
Querier Port : lag255
Querier UpTime :6m 40s
Querier Expiration Time :3m 53s
```

Task 4: Verify the IGMP Snooping Operation

Objectives

In this task, the multicast test traffic will be used again, but this time, the multicast traffic should be filtered by default due to the IGMP snooping configuration.

Once the lab has demonstrated that the traffic was filtered, the client device will send an IGMP join, using the multicast test application, and the multicast traffic should be forwarded to the port.

Any other ports that have not sent an IGMP join should not receive the multicast traffic for the test group.

Steps

Verify Multicast Traffic is No Longer Flooded

1. Open PC3 (connected to Access1) and start the multicast test traffic again by clicking the **'Start Sender'** button.
2. You may optionally re-start the Wireshark packet capture for a few seconds to confirm traffic is sent out.
3. Open PC4 (connected to Access2) and start the Wireshark packet capture for a few seconds and then stop again.
4. Confirm that **no** multicast test traffic was received.

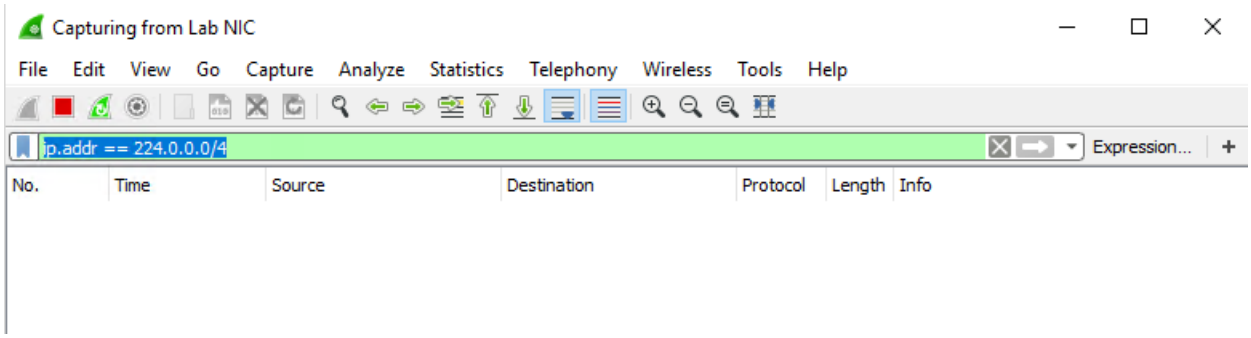

NOTE: Some IGMP traffic may be observed since the Core IGMP querier will send regular membership query reports.

Verify Multicast Traffic Is Still Delivered When Requested

5. On the PC4, connected to Access2, start the Wireshark capture again.
6. Start the UDP multicast test tool from the desktop
7. Configure the **Receiver** section (bottom half of the screen).

– **Local Multicast Interface address**: Enter the PC4 Lab NIC IP address

NOTE: The local address should be listed in the 'Local Interfaces' list. You may also drag and drop from this list into the field.

NOTE: Do not confuse this with the 'Local Interface Address' field.

– **Multicast address: 239.x.1.1**

– Click the **Start Receiver** button

NOTE: You may need to scroll down or resize the window to see the **Start Receiver** button.

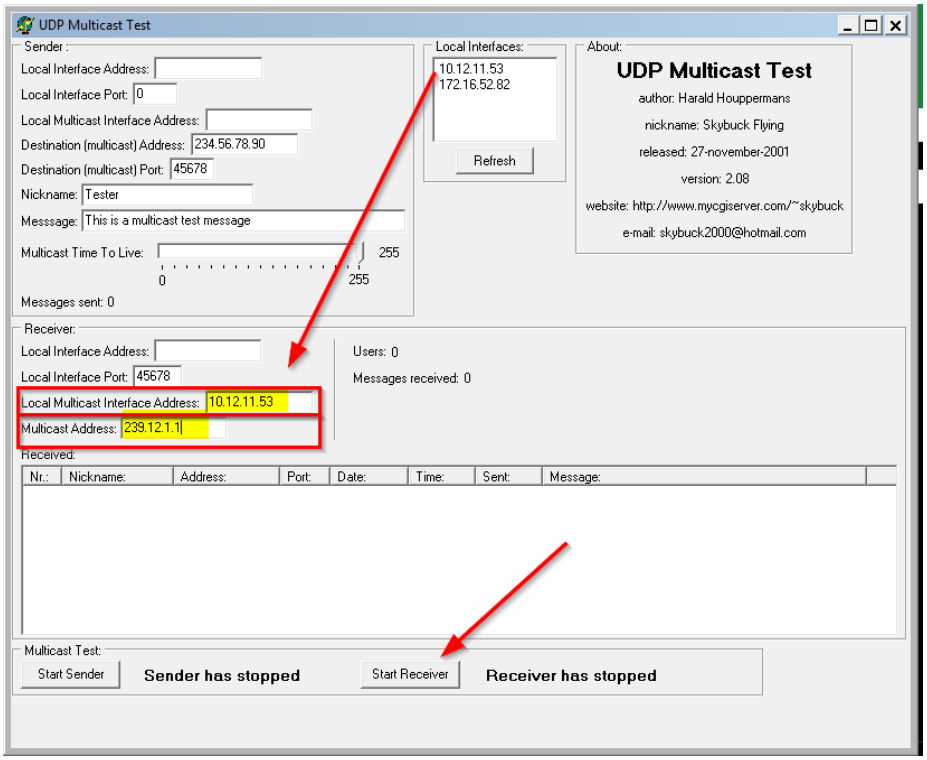

8. In the UDPMulticast test application, the **Received** section should now show the incoming test message.

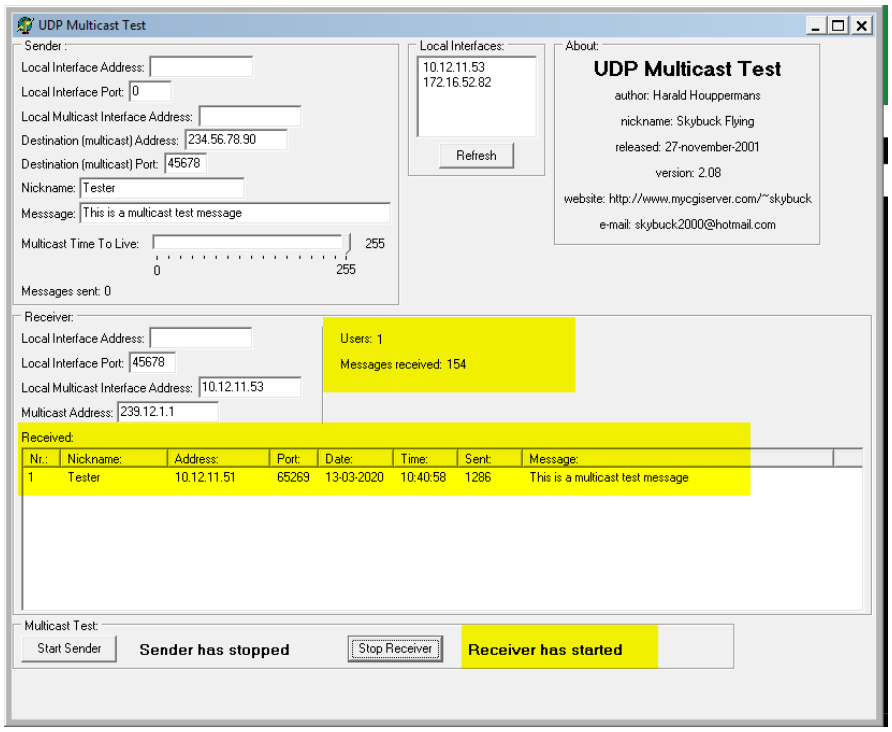

9. Stop the Wireshark trace and scroll back to the start of the trace.

> **TIP:** It may be easier to apply display filter **'ip.addr == 224.0.0.0/4'** (**not /24 !**) to see only the multicast traffic.

The trace should show that at the moment the multicast test client application starts the receiver, the IGMP join is sent by the PC.

Following the join message, the actual multicast traffic will now be received by this switch port.

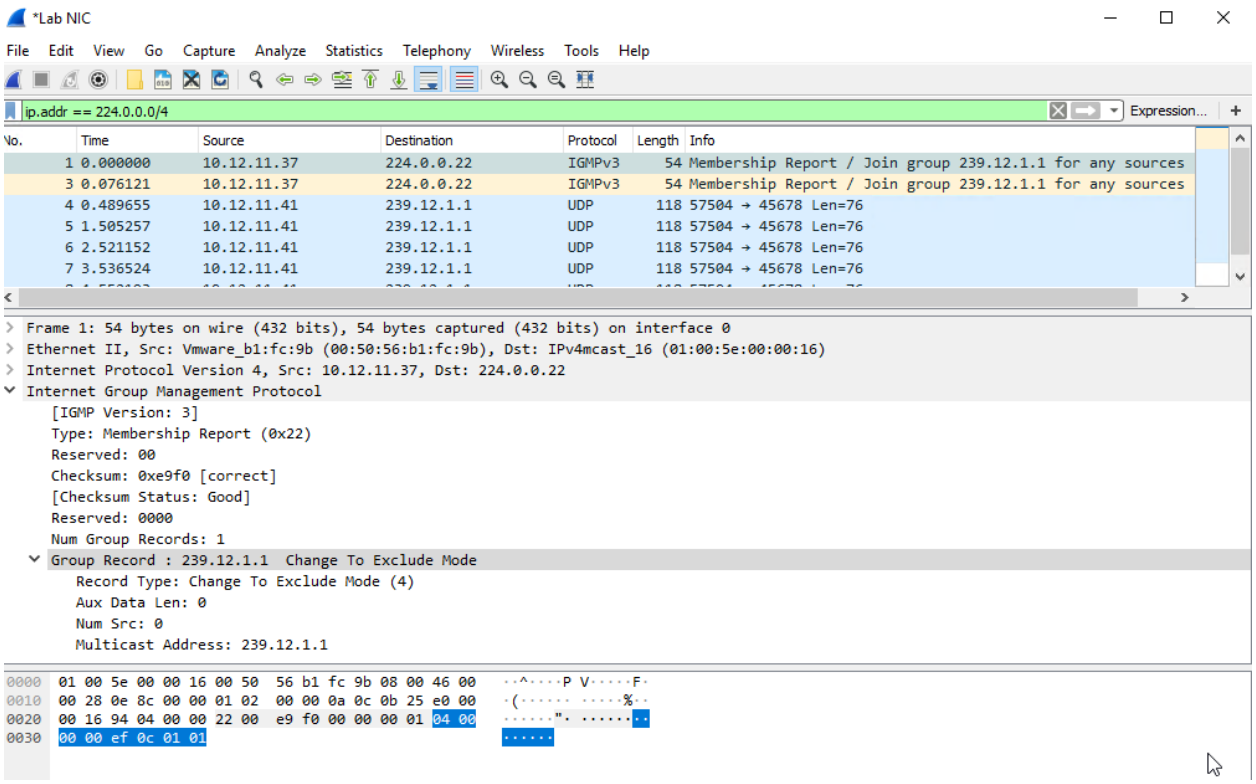

Verify the IGMP Snooping Status on the Switches

Access2

10. On Access2, review all the IGMP snooping groups.

```
ICX-Tx-Access2(config)# show ip igmp snooping groups

IGMP Group Address Information

VLAN ID Group Address   Expires   UpTime    Last Reporter   Type
------- --------------- --------- --------- --------------- ------
11      239.x.1.1       3m 34s    48m 32s   10.x.11.37      Filter
11      239.255.255.250 3m 36s    1h 16m    10.x.11.37      Filter
```

11. Next, review the IGMP group tables for VLAN 11.

```
ICX-Tx-Access2(config)# show ip igmp snooping vlan 11
IGMP Snooping Protocol Info

Total VLANs with IGMP enabled             : 2
Current count of multicast groups joined  : 2

IGMP Drop Unknown Multicast               : Global
VLAN ID : 11
```

```
VLAN Name : VLAN11
IGMP Configured Version : 3
IGMP Operating Version : 3
Querier Address : 10.x.11.2
Querier Port : lag256
Querier UpTime :23m 15s
Querier Expiration Time :4m 6s

Active Group Address   Tracking  Vers Mode Uptime    Expires
---------------------- ---------- ---- ---- --------- ----------
239.x.1.1              Filter     3    EXC  5m 46s    4m 16s
239.255.255.250        Filter     3    EXC  23m 11s   4m 18s
```

12. Review the details of the group 239.x.1.1 in this VLAN. This will show which port received the multicast request.

```
ICX-Tx-Access2(config)# show ip igmp snooping vlan 11 group 239.x.1.1

IGMP ports and group information for group 239.12.1.1

VLAN ID   : 11
VLAN Name : VLAN11

Group Address : 239.x.1.1
Last Reporter : 10.x.11.37
Group Type    : Filter

                                        V1        V2        Sources   Sources
Port      Vers Mode Uptime    Expires   Timer     Timer     Forwarded Blocked
--------- ---- ---- --------- --------- --------- --------- --------- --------
1/1/4     3    EXC  5m 58s    4m 4s                         0         0

ICX-Tx-Access2(config)#
```

13. On the Core1 and Core2, review the IGMP group table 239.x.1.1. Both systems should show a joined client on the lag2. This is the LAG to the Access2 switch where the receiver is connected.

Core1

```
ICX-Tx-Core2(config)# show ip igmp snooping vlan 11 group 239.x.1.1

IGMP ports and group information for group 239.x.1.1

VLAN ID   : 11
VLAN Name : VLAN11

Group Address : 239.x.1.1
Last Reporter : 10.x.11.37
Group Type    : Filter
```

```
                                          V1        V2        Sources   Sources
Port      Vers Mode Uptime    Expires     Timer     Timer     Forwarded Blocked
--------- ---- ---- --------- --------- --------- --------- --------- --------
lag2      3    EXC  8m 21s    3m 51s                        0         0
```

Core2

```
ICX-Tx-Core2(config)# show ip igmp snooping vlan 11 group 239.x.1.1

IGMP ports and group information for group 239.12.1.1

VLAN ID   : 11
VLAN Name : VLAN11

Group Address : 239.x.1.1
Last Reporter : 10.x.11.37
Group Type    : Filter

                                          V1        V2        Sources   Sources
Port      Vers Mode Uptime    Expires     Timer     Timer     Forwarded Blocked
--------- ---- ---- --------- --------- --------- --------- --------- --------
lag2      3    EXC  8m 52s    3m 19s                        0         0
```

Access1

14. On the Access1, review the IGMP group tables.

```
ICX-Tx-Access1(config)# show ip igmp snooping vlan 11
IGMP Snooping Protocol Info

Total VLANs with IGMP enabled             : 2
Current count of multicast groups joined  : 1

IGMP Drop Unknown Multicast               : Global
VLAN ID : 11
VLAN Name : VLAN11
IGMP Configured Version : 3
IGMP Operating Version : 3
Querier Address : 10.x.11.2
Querier Port : lag255
Querier UpTime :28m 17s
Querier Expiration Time :3m 50s

Active Group Address   Tracking  Vers Mode Uptime    Expires
---------------------- ---------- ---- ---- --------- ----------
239.255.255.250        Filter    3    EXC  28m 14s   3m 28s
```

Q: Why is the 239.x.1.1 group address not listed here?

_____

A: Multicast traffic is always forwarded to the IGMP querier port, so there is no need to filter the traffic on this port.

This demonstrates the filtering and forwarding of the multicast traffic.

Optional Task 5: Verify IGMP Snooping Fast-leave

This task is **optional** and can be done if time permits. Check with your instructor.

Objectives

In this task, the IGMP Fast leave feature will be verified. When a client signals that it no longer needs the multicast stream using an IGMP leave message, the IGMP querier will check if there are any other receivers active on the network.

If no other receivers respond, the multicast can be filtered immediately.

This will be demonstrated using the PC4 (connected to Access2).

Steps

On the PC4 (connected to Access2), you will stop the receiver.

Verify using Wireshark that:

- the IGMP leave message is sent
- the querier checks if any other receivers need the multicast
- the multicast transmission stops

1. Open the connection to PC4 (connected to Access2).
2. Start a new Wireshark packet capture for the 'Lab NIC' and leave it running. You should observe the incoming multicast to 239.x.1.1.

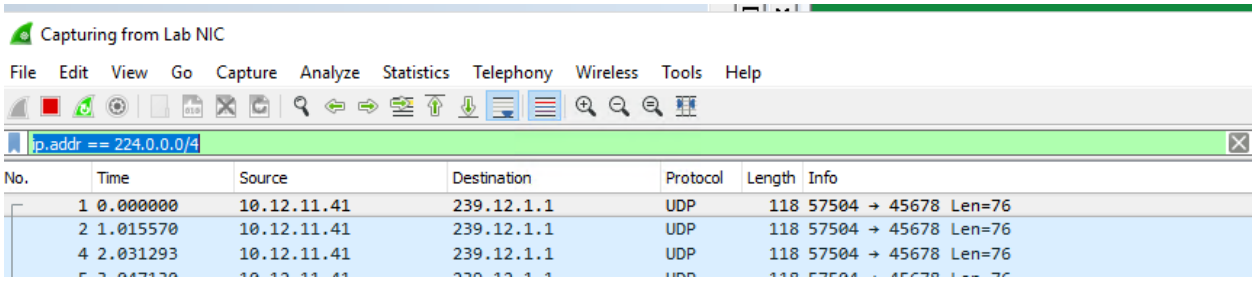

3. In the multicast test application, stop the receiver.

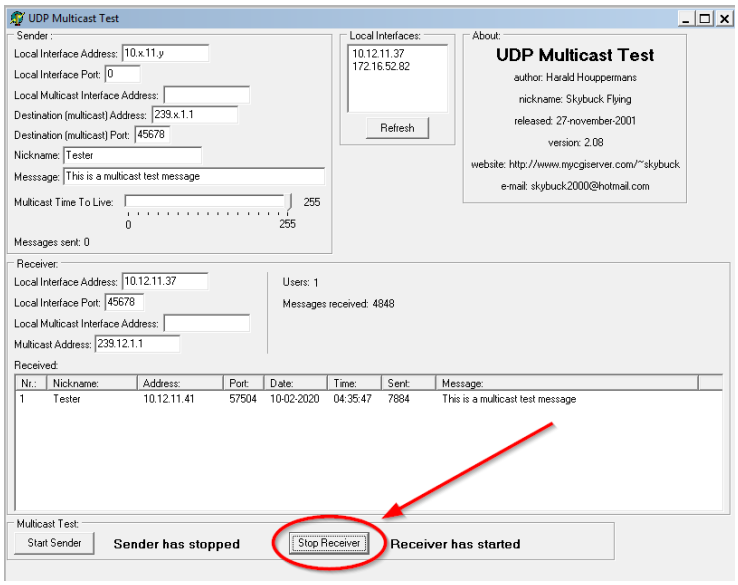

4. Return to the Wireshark session and stop the packet capture after about 5 seconds.

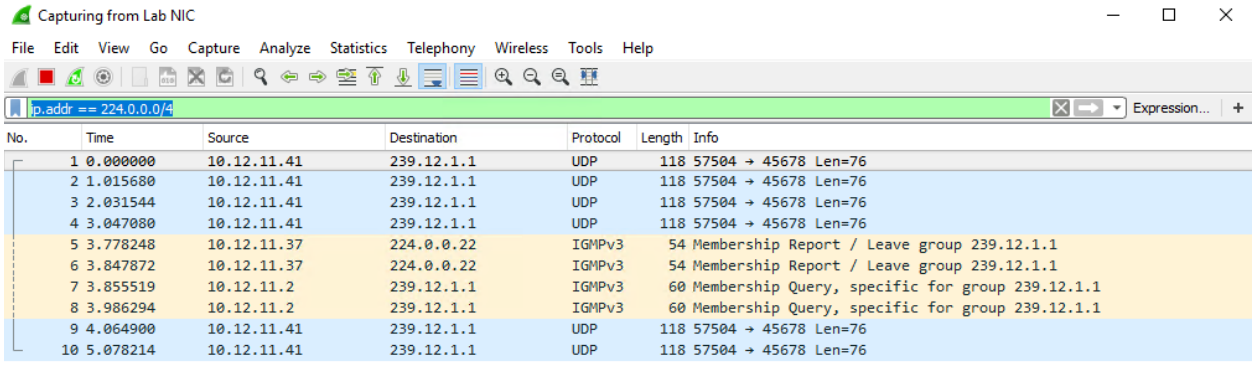

The trace should demonstrate that the client is sending an IGMP 'leave group' to the 224.0.0.22 address (the IGMPv3 multicast address).

Next, the current IGMP querier will check if any other memberships are still active.

Since no other clients respond, the snooping will also terminate the multicast forwarding, so after 1-2 additional seconds, there should be no more multicast traffic for the 239.x.1.1 destination.

5. On all the switches, verify that the 239.x.1.1 is no longer listed.

```
show ip igmp snooping groups
```

6. You can easily repeat the tests by restarting the multicast receiver client again. This demonstrates the fast-leave operation.

You have completed Lab 8!

Lab 09 IP PIM Sparse Mode- Multicast Routing

Lab Diagram

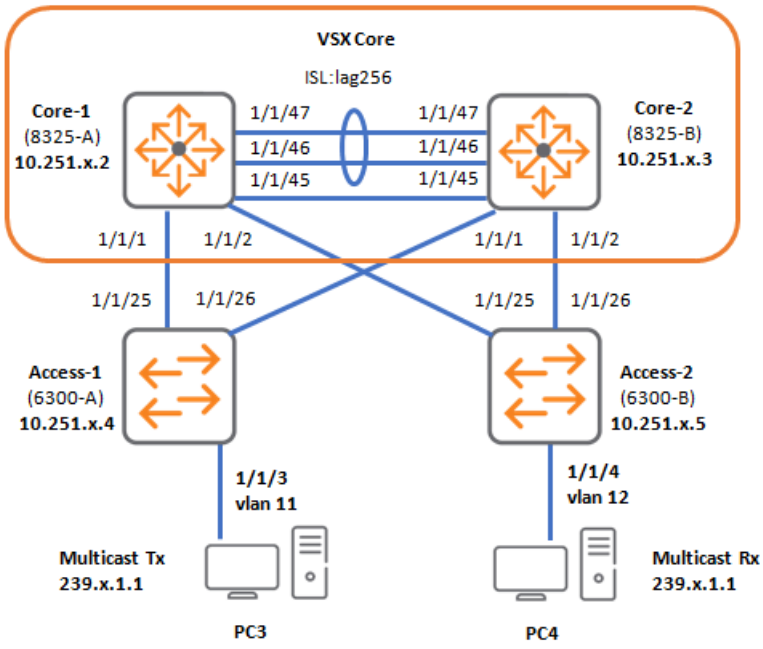

Overview

In this lab activity, multicast routing will be configured using PIM Sparse Mode on the VSX Core. Multicast routing also requires an IGMP Querier, this function has been enabled in the previous lab activity (Configuring IGMP). The result will be that the multicast that is transmitted by PC3 in VLAN11 can be received by PC4 in VLAN 12.

Objectives

- Understand multicast routing
- Apply PIM Sparse Mode global and interface level configuration
- Verify the multicast routing table

Task 1: Prepare and Review the Lab Setup

Objectives

This lab requires the completion of the 'Configuring IGMP' lab activity.

PC3 should be in VLAN 11, in this task; PC4 will be moved to VLAN 12.

Steps

Access2

1. Access the terminal of Access2 and enter configuration mode.
2. Assign the port of PC4 to VLAN 12.

```
ICX-Tx-Access2(config)# interface 1/1/4
ICX-Tx-Access2(config-if)# vlan access 12
ICX-Tx-Access2(config-if)# exit
ICX-Tx-Access2(config)#
```

3. Open PC4 (connected to Access2). Renew the IP address by bouncing (disable/enable) the 'Lab NIC' port.

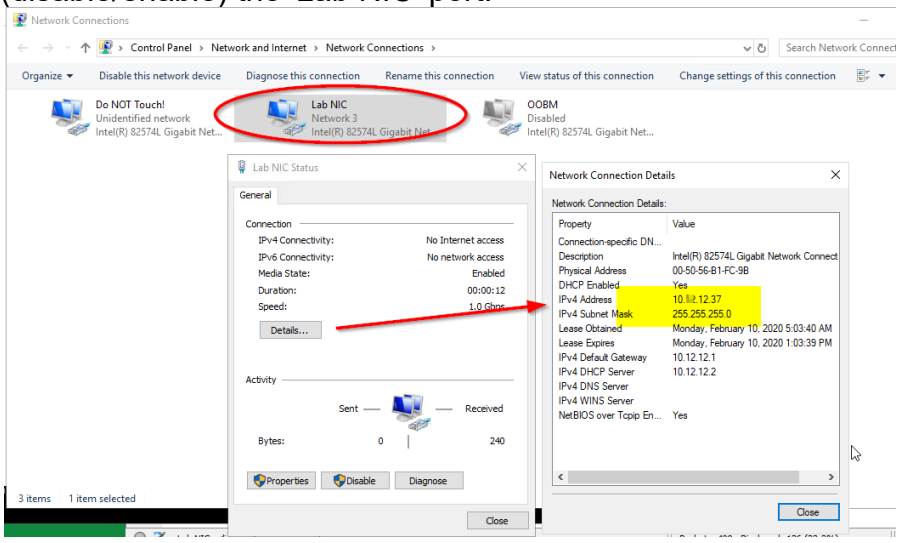

4. The client should now have an IP address from the VLAN 12 range.

Verify IP IGMP Querier on the Core switches

PIM multicast routing relies on IGMP join packets from the client devices to learn which multicasts are requested by the clients.

Therefore, IP IGMP querier function must be enabled and functional on the receiving client interfaces before PIM routing can operate.

Core1

5. On Core1, verify the IP IGMP Querier status for vlan11 and vlan12 interfaces.

```
ICX-Tx-Core1(config)# show ip igmp interface vlan11

IGMP Configured Version  : 3
IGMP Operating Version   : 3
Querier State              : Querier
Querier IP [this switch]   : 10.x.11.2
Querier Uptime             : 3h 38m
...
```

```
ICX-Tx-Core1(config)# show ip igmp interface vlan12

IGMP Configured Version  : 3
IGMP Operating Version   : 3
Querier State              : Querier
Querier IP [this switch]   : 10.x.12.2
Querier Uptime             : 5m 24s
Querier Expiration Time    : 1m 58s
```

Core2

6. Repeat the verification on Core2.

NOTE: This can also be done via Core1 by appending 'vsx-peer' to the commands.

```
ICX-Tx-Core2(config)# show ip igmp interface vlan11

IGMP Configured Version  : 3
IGMP Operating Version   : 3
Querier State              : Non-Querier
Querier IP                 : 10.x.11.2
Querier Uptime             : 26m 30s
Querier Expiration Time    : 2m 57s
```

```
ICX-Tx-Core2(config)# show ip igmp interface vlan12

IGMP Configured Version  : 3
IGMP Operating Version   : 3
Querier State              : Non-Querier
Querier IP                 : 10.x.12.2
Querier Uptime             : 5m 27s
Querier Expiration Time    : 3m 32s
```

Task 2: Configure PIM Sparse

Objectives

In this task, the PIM Sparse Mode (PIM-SM) Candidate Bootstrap Router (C-BSR) is configured.

Once the BSR has been elected, the Candidate Rendez-vous points (RP) will be elected and distributed by the BSR.

On the RP, the administrator can specify that the RP should be the handling point for a specific range of multicast addresses.

Steps

Core1 PIM-SM configuration

1. Open a terminal to Core1, enter the configuration mode.
2. Access the router PIM context.

```
ICX-Tx-Core1(config)# router pim
```

3. For multicast on a VSX cluster, the best practice is to configure PIM Dual-DR under the PIM router command. With the 'active-active' command, the proxy-DR will also learn the multicast routes and this will allow fast recovery time if the actual DR fails.

```
ICX-Tx-Core1(config-pim)# active-active
```

> **NOTE:** In a VSX cluster, the VSX member with the highest IP address will be the forwarder of the multicast traffic. In this lab setup, the Core2 will become the forwarder. Do not configure PIM interface level DR priority when active-active is enabled, since the active-active feature will over-rule the manually configured interface priority.

4. Attempt to enable the C-BSR function using the VLAN1 interface as the source address.

> **NOTE:** In a typical deployment, a loopback interface would be used for this function. This lab uses the VLAN 1 IP interface since it is already configured. If a loopback interface is used, make sure it is routable in the network.

```
ICX-Tx-Core1(config-pim)# bsr-candidate source-ip-interface vlan1
PIM SM should be configured on one or more interface.
```

Q: Why did the switch return an error?

A: While the VLAN1 IP interface exists, the PIM protocol must first be enabled on the interface before it can be used in the PIM election process.

5. Enable PIM on the VLAN1 interface.

```
ICX-Tx-Core1(config-pim)# exit
ICX-Tx-Core1(config)# interface vlan1
ICX-Tx-Core1(config-if-vlan)# ip pim-sparse enable
ICX-Tx-Core1(config-if-vlan)# exit
ICX-Tx-Core1(config)#
```

6. Re-try to enable the C-BSR function using the loopback interface and enable the PIM routing protocol.

```
ICX-Tx-Core1(config)# router pim
ICX-Tx-Core1(config-pim)# bsr-candidate source-ip-interface vlan1
ICX-Tx-Core1(config-pim)# enable
```

7. Assign the Core1 BSR priority 250.

```
ICX-Tx-Core1(config-pim)# bsr-candidate priority 250
ICX-Tx-Core1(config-pim)# exit
```

8. Review the local BSR candidate information.

```
ICX-Tx-Core1(config)# show ip pim bsr local

Status and Counters - PIM-SM Local Candidate-BSR Information

VRF                       : default
C-BSR Admin Status        : This system is a Candidate-BSR
C-BSR Address             : 10.x.1.2
C-BSR Priority            : 250
C-BSR Hash Mask Length    : 30
C-BSR Message Interval    : 60
C-BSR Source IP Interface : vlan1
```

9. Enable PIM SM on interfaces VLAN 11 and VLAN 12.

```
ICX-Tx-Core1(config)# interface vlan11
ICX-Tx-Core1(config-if-vlan)# ip pim-sparse enable
ICX-Tx-Core1(config-if-vlan)# exit
```

```
ICX-Tx-Core1(config)# interface vlan12
ICX-Tx-Core1(config-if-vlan)# ip pim-sparse enable
ICX-Tx-Core1(config-if-vlan)# exit
```

Core2 PIM-SM Configuration

10. On Core2, configure PIM-SM on the interfaces and make it C-BSR.

```
ICX-Tx-Core2(config)# interface vlan1
ICX-Tx-Core2(config-if-vlan)# ip pim-sparse enable
ICX-Tx-Core2(config-if-vlan)# exit
```

```
ICX-Tx-Core2(config)# interface vlan11
ICX-Tx-Core2(config-if-vlan)# ip pim-sparse enable
ICX-Tx-Core2(config-if-vlan)# exit
```

```
ICX-Tx-Core2(config)# interface vlan12
ICX-Tx-Core2(config-if-vlan)# ip pim-sparse enable
ICX-Tx-Core2(config-if-vlan)# exit
```

```
ICX-Tx-Core2(config)# router pim
ICX-Tx-Core2(config-pim)# active-active
ICX-Tx-Core2(config-pim)# bsr-candidate source-ip-interface vlan1
ICX-Tx-Core2(config-pim)# enable
ICX-Tx-Core2(config-pim)# exit
```

Verify the BSR Election

Core1

11. On the Core1, review the elected BSR information.

```
ICX-Tx-Core1(config)# show ip pim bsr elected

Status and Counters - PIM-SM Elected Bootstrap Router Information

VRF                       : default
E-BSR Address             : 10.x.1.2
E-BSR Priority            : 250
E-BSR Hash Mask Length    : 30
E-BSR Up Time             : 1 mins 57 secs
Next Bootstrap Message    : 13 secs
```

Core2

12. On Core2, review the elected BSR information.

```
ICX-Tx-Core2(config)# show ip pim bsr elected

Status and Counters - PIM-SM Elected Bootstrap Router Information
```

```
VRF                         : default
E-BSR Address               : 10.x.1.2
E-BSR Priority              : 250
E-BSR Hash Mask Length      : 30
E-BSR Up Time               : 2 mins 40 secs
Next Bootstrap Message      : 1 mins 40 secs
```

NOTE: It may take up to 2 minutes before Core2 shows 10.x.1.2 as the E-BSR address.

PIM-SM Rendez-vous Point (RP)

Core1

13. On Core1, review the default RP candidate information and active RP sets.

```
ICX-Tx-Core1(config)# show ip pim rp-candidate
ICX-Tx-Core1(config)# show ip pim rp-set
```

This shows that there is no default RP on the system.

14. Make Core1 a candidate RP.

```
ICX-Tx-Core1(config)# router pim
ICX-Tx-Core1(config-pim)# rp-candidate source-ip-interface vlan1
```

Core2

15. Make Core2 a candidate RP.

```
ICX-Tx-Core2(config)# router pim
ICX-Tx-Core2(config-pim)#  rp-candidate source-ip-interface vlan1
```

Core1 - Verify the RP Set

16. On Core1, review the RP-set.

```
ICX-Tx-Core1(config-pim)# show ip pim rp-set

VRF: default

Status and Counters - PIM-SM Learned RP-Set Information
Group Address    Group Mask       RP Address       Hold Time  Expire Time
---------------  ---------------  ---------------  ---------  -----------
224.0.0.0        240.0.0.0        10.x.1.3         150        149
224.0.0.0        240.0.0.0        10.x.1.2         150        149
```

NOTE: It may take a minute for the Core2 (10.x.1.3) to appear in the list. You do not need to wait for it.

RP for Multicast Subset Group

By default, an RP will be RP for all possible multicast entries (224.0.0.0/4). In a large network with multiple sites or regions, the administrator may want to have a regional RP for the local multicast traffic. In this case, the administrator should use structural IP Addressing for the multicast applications.

In this lab, each table will be assumed a 'region', so the local RPs will be configured to be the preferred RPs for the local multicast ranges (239.x.0.0/16).

Core1

17. On Core1, configure the 239.x.0.0/16 group-prefix for the RP.

```
ICX-Tx-Core1(config-pim)# rp-candidate group-prefix 239.x.0.0/16
ICX-Tx-Core1(config-pim)# exit
```

Core2

18. Repeat this on Core2.

```
ICX-Tx-Core2(config-pim)# rp-candidate group-prefix 239.x.0.0/16
ICX-Tx-Core2(config-pim)# exit
```

19. On Core2, review the resulting local candidate RP information.

```
ICX-Tx-Core2(config)# show ip pim rp-candidate
Status and Counters- PIM-SM Candidate-RP Information

VRF                       : default
C-RP Admin Status         : This system is a Candidate-RP
C-RP Address              : 10.x.1.3
C-RP Hold Time            : 150
C-RP Advertise Period     : 60
C-RP Priority             : 192
C-RP Source IP Interface  : vlan1

Group Address    Group Mask
---------------  ---------------
224.0.0.0        240.0.0.0
Group Address    Group Mask
---------------  ---------------
239.12.0.0       255.255.0.0
```

Verify RP-Set and group-prefixes

20. Verify that both Core1 (10.x.12.2) and Core2 (10.x.12.3) are now listed in the RP-set for the 239.x.0.0/16 group prefix.

```
ICX-Tx-Core2(config)# show ip pim rp-set

VRF: default

Status and Counters - PIM-SM Learned RP-Set Information
Group Address    Group Mask       RP Address       Hold Time  Expire Time
---------------  ---------------  ---------------  ---------  -----------
```

```
224.0.0.0         240.0.0.0         10.x.1.3          150         131
224.0.0.0         240.0.0.0         10.x.1.2          150         131
239.x.0.0         255.255.0.0       10.x.1.2          150         131
239.x.0.0         255.255.0.0       10.x.1.3          150         131
```

NOTE: It may take up to a minute before the updated RP-SET is available.

Task 3: Verify Multicast Forwarding

Objectives

In this task, the multicast routing will be tested between the hosts in VLAN 11 and VLAN 12.

- PC3 (connected to Access1) belongs to VLAN 11 and will be transmitting the multicast test traffic.
- PC4 (connected to Access2) belongs to VLAN 12 and will request the test multicast traffic.

Once the multicast traffic arrives the status of the multicast group information will be reviewed on the Core switches.

Steps

Test Multicast Transmitter and Registration

1. On PC3 (connected to Access1), configure the UDP Multicast Test tool as '**Sender**'. (The tool can be found on the desktop).

Local Interface Address **10.x.11.y** (Your PC as shown in 'Local Interfaces' list)

Destination multicast address **239.x.1.1**

2. Click on '**Start Sender**' to make the multicast Sender is active. It will now start to send traffic to the 239.x.1.1 destination address, using the local interface 10.x.11.y.

NOTE: You may need to scroll down to see the 'Start Sender' button.

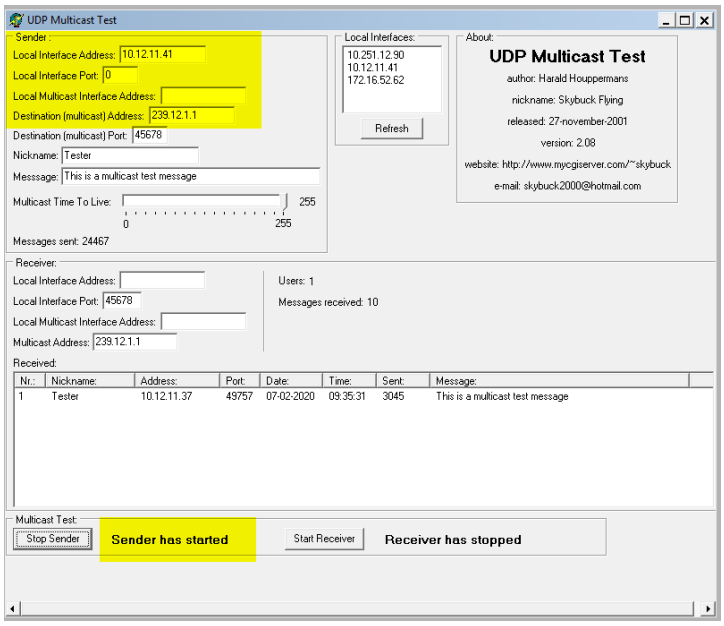

Core1

3. On Core1, review the brief output of the current multicast routing table. The 239.x.1.1 multicast should have been learned by the data-plane on VLAN 11. This also shows the source address of the transmitter (10.x.11.y).

```
ICX-Tx-Core1(config)# show ip mroute brief
IP Multicast Route Entries

VRF : default
Total number of entries : 1

Group Address      Source Address     Neighbor           Interface
-------------      --------------     --------           ---------
239.x.1.1          10.x.11.y          10.x.11.3         vlan11
```

4. Review the details of the 239.x.1.1 route. This will show that the multicast was registered with the PIM-SM protocol.

```
ICX-Tx-Core1(config)# show ip mroute 239.x.1.1

IP Multicast Route Entries

VRF : default

Group Address                : 239.x.1.1
Source Address               : 10.x.11.y
Neighbor                     : 10.x.11.3
Incoming interface           : vlan11
Multicast Routing Protocol   : PIM-SM
Unicast Routing Protocol     : connected
Metric                       : 0
Metric  Pref                 : 0
```

Q: What is the Incoming interface?

_____

A: The multicast traffic is received by the Core switch on the VLAN 11 interface.

Q: Are there any outgoing interfaces listed?

_____

A: Since there are no registered subscribers, there are no outgoing interfaces at this moment.

Test the Multicast Routing Forwarding

5. On PC4 (connected to Access2), start a wireshark trace on the 'Lab NIC'.

6. Next enable the UDP Multicast test tool as **'Receiver'** (lower).
   - **Local Multicast interface address: 10.x.12.y**

     Make sure this is the updated VLAN12 ip address, you can click on **'Refresh'** button to see the updated **'Local Interfaces'**.
   - **Multicast address: 239.x.1.1**
7. Scroll down and click on **'Start Receiver'**. Verify that the test messages are received in the tool.
8. Return to the Wireshark, stop the trace. The test tool has sent the IGMP join request. The Core switch has received the IGMP join and has initiated the multicast routing forwarding process.

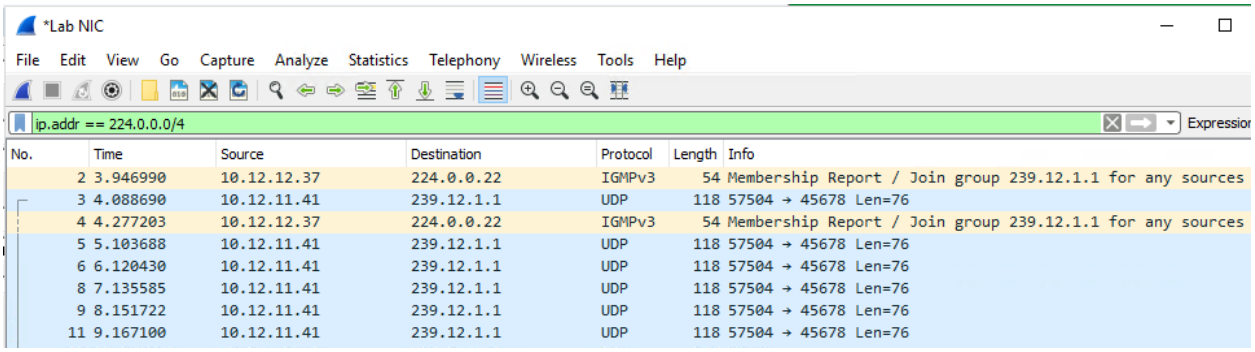

Core1

9. On Core1, review that the host in VLAN 12 has joined using IGMP.

```
ICX-Tx-Core1(config)# show ip igmp group 239.x.1.1

IGMP group information for group 239.12.1.1

Interface Name    : vlan12
VRF Name          : default

Group Address     : 239.12.1.1
Last Reporter     : 10.x.12.y

                                  V1         V2         Sources    Sources
Vers Mode Uptime     Expires      Timer      Timer      Forwarded  Blocked
---- ---- --------- --------- ---------- ---------- ---------- --------
3    EXC  4m 13s    3m 7s
```

10. On Core1, review the multicast routing table. Notice the downstream interface, this was initiated based on the IGMP join. The Proxy DR is used by VSX to synchronize the active multicast routes between the 2 VSX nodes.

```
ICX-Tx-Core1(config)# show ip mroute 239.x.1.1

IP Multicast Route Entries
```

```
VRF : default

Group Address                 : 239.x.1.1
Source Address                : 10.x.11.y
Neighbor                      : 10.x.11.3
Incoming interface            : vlan11
Multicast Routing Protocol    : PIM-SM
Unicast Routing Protocol      : connected
Metric                        : 0
Metric  Pref                  : 0
Downstream Interface
Interface      State               By_Proxy_Dr
---------      -----               -----------
vlan12         forwarding          false
```

11. And check the vsx-peer as well (output from Core2).

```
ICX-Tx-Core1(config)# show ip mroute 239.x.1.1 vsx-peer

IP Multicast Route Entries

VRF : default

Group Address                 : 239.x.1.1
Source Address                : 10.x.11.y
Neighbor                      : 10.x.11.2
Incoming interface            : vlan11
Multicast Routing Protocol    : PIM-SM
Unicast Routing Protocol      : connected
Metric                        : 0
Metric  Pref                  : 0
Downstream Interface
Interface      State               By_Proxy_Dr
---------      -----               -----------
vlan12         forwarding          true

Group Address                 : 239.12.1.1
Source Address                : 10.x.11.y
Neighbor                      : 10.x.12.2
Incoming interface            : vlan12
Multicast Routing Protocol    : PIM-SM
Unicast Routing Protocol      : connected
Metric                        : 0
Metric  Pref                  : 0
```

You've completed Lab 9!

Lab 10: 802.1X and User Roles- 802.1X Authentication and User Roles on AOS-CX

Lab diagram

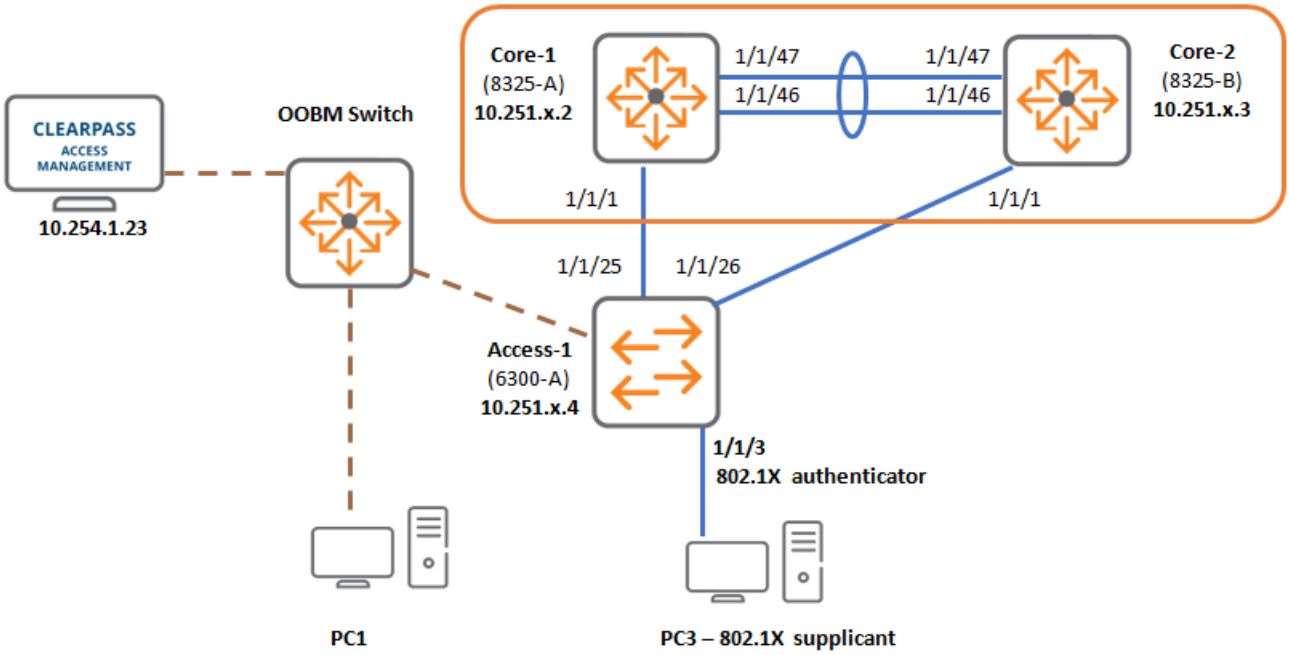

Overview

In this lab activity, 802.1X will be configured on an access switch. This will require the configuration of a RADIUS server and will include RADIUS change of authorization (CoA) support.

For the first 802.1X authentication part of the lab, the RADIUS server will provide standard IETF attribute value pairs (AVP) to assign the user to a VLAN.

Next, the lab will introduce Aruba user-roles; therefore, the access instructions will be stored in the switch configuration and the RADIUS server will use the Aruba Vendor-Specific Attribute (VSA) called *Aruba-User-Role*.

Objectives

- Configure RADIUS server, authentication groups and dynamic authorization (CoA)
- Configure RADIUS tracking
- Configure 802.1X on a switch port

- Configure user roles with policies and classes

Task 1: Prepare the Lab Start Configuration

Objectives

- This lab is built on the base VSX topology.
- Make sure to complete these steps to get the base VSX checkpoint configuration on the devices.

Steps (Required)

1. Open a console connection to the 6300A. Login using **admin** and a password of **aruba123**.

```
ICX-Tx-Access1# copy checkpoint icx-lab04-vsx running-config
Configuration changes will take time to process, please be patient.
ICX-Tx-Access1#
```

2. Open a console connection to the 6300B. Login using **admin** and a password of **aruba123**.

```
ICX-Tx-Access2# copy checkpoint icx-lab04-vsx running-config
Configuration changes will take time to process, please be patient.
ICX-Tx-Access2#
```

3. Open a console connection to the 8325A. Login using **admin** and a password of **aruba123**.

```
ICX-Tx-Core1# copy checkpoint icx-lab04-vsx running-config
Configuration changes will take time to process, please be patient.
ICX-Tx-Core1#
```

4. Open a console connection to the 8325B. Login using **admin** and a password of **aruba123**.

```
ICX-Tx-Core2# copy checkpoint icx-lab04-vsx running-config
Configuration changes will take time to process, please be patient.
ICX-Tx-Core2#
```

Task 2: RADIUS Server Setup

Diagram

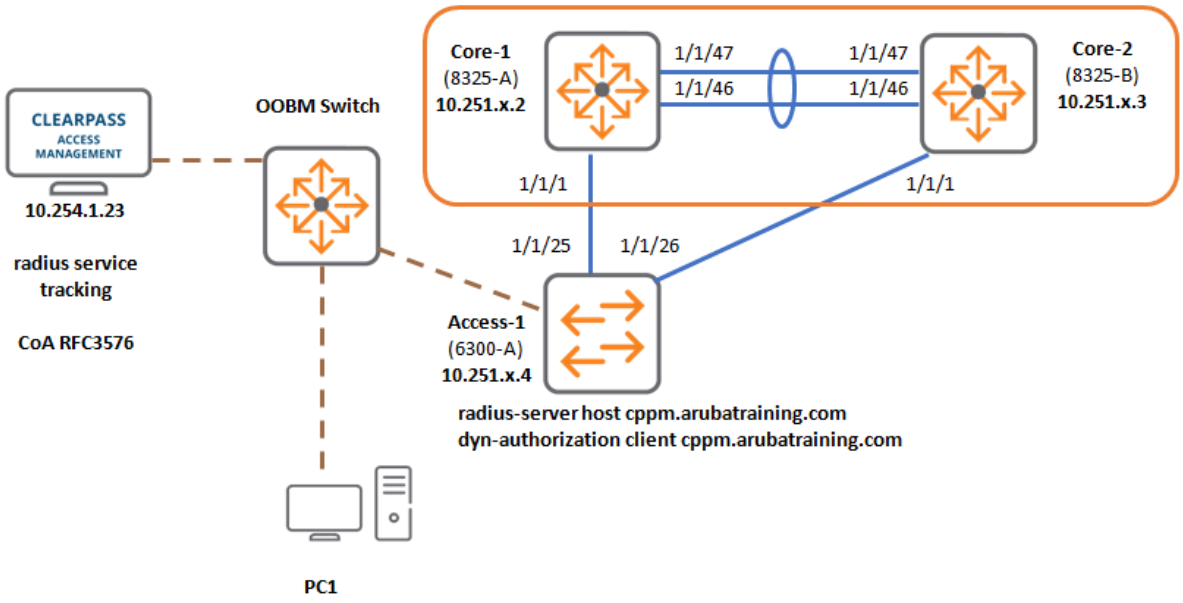

Objectives

In this task, the Access1 switch will be configured to use the ClearPass RADIUS server. The training lab is using the OOBM network to access the RADIUS server since that is readily available. In real deployments, the RADIUS server would typically be reachable over the in-band network.

AOS-CX supports the use of server groups.

RADIUS accounting and interim accounting will be enabled.

CoA (RFC3576 support) will be enabled and verified.

RADIUS server tracking will be enabled to verify that the RADIUS server is operational.

Steps

Access1

1. Open a session to Access1 and enter the configuration mode.

In a later lab, the ClearPass Downloadable User Roles will be configured. Since this feature will require a HTTPS connection between the switch and the ClearPass host, this HTTPS connection will be validated by the switch; therefore, the CPPM HTTPS certification subject name must match with the DNS name.

To prepare for this later activity, the RADIUS server will be defined based on the hostname and not based on an IP address. If the customer has an entry for the ClearPass host in the DNS system, that can be used as well.

In this lab setup, the hostname will be locally defined on the switch, similar to the 'hosts' file on many Operating Systems.

2. Define a new hostname for the cppm.arubatraining.com in the mgmt VRF.

```
ICX-Tx-Access1(config)# ip dns host cppm.arubatraining.com 10.254.1.23 vrf mgmt
```

3. Verify that the new entry works with a test ping.

```
ICX-Tx-Access1(config)# do ping cppm.arubatraining.com vrf mgmt
PING cppm.arubatraining.com (10.253.1.23) 100(128) bytes of data.
108 bytes from 10.253.1.23: icmp_seq=1 ttl=63 time=0.727 ms
...
```

4. Define a new RADIUS host for the ClearPass server.

NOTE: Make sure to use the correct VRF commands, otherwise the RADIUS server will be defined in the default VRF, but it will not be reachable.
Typically, the RADIUS would be configured in an in-band VRF, such as the 'default' VRF. The 'mgmt' VRF is convenient in the training lab environment.

NOTE: Make sure to verify the key when entering the command.

```
ICX-Tx-Access1(config)# radius-server host cppm.arubatraining.com key plaintext
aruba123 vrf mgmt
```

5. Define a new server group, add the previously defined CPPM host. Again, make sure to enter the correct VRF.

```
ICX-Tx-Access1(config)# aaa group server radius cppm
ICX-Tx-Access1(config-sg)# server cppm.arubatraining.com vrf mgmt
ICX-Tx-Access1(config-sg)# exit
```

6. Enable RADIUS accounting to this server group, enable interim accounting, set the interim accounting to 5 minutes. This will ensure that the switch updates the radius server about the connected devices every 5 minutes and will ensure that ClearPass has a view over the currently connected devices in the network.

```
ICX-Tx-Access1(config)# aaa accounting port-access start-stop interim 5 group
cppm
```

7. Enable the Change of Authorization (CoA) processing on the switch. This is also known as RFC 3576 support or dynamic authorization. This allows the RADIUS server (ClearPass) to send a message to the switch to request an update in the authorization or a re-authentication. Basically, ClearPass can have a 'triggered' re-authentication based on authentication events. This could be required in a later lab when using the Captive Portal authentication.

```
ICX-Tx-Access1(config)# radius dyn-authorization client cppm.arubatraining.com
secret-key plaintext aruba123 vrf mgmt replay-protection disable
```

```
ICX-Tx-Access1(config)# radius dyn-authorization enable
```

NOTE: Replay protection disable indicates that the timestamp of the CoA packet will not be inspected by the switch. In a real deployment, the replay protection should be enabled. The default allowed time difference between the RADIUS host and the switch is 300 seconds. Since in the lab, the time of the switch and CPPM may not be in sync, the replay-protection is disabled.

NOTE: Since ClearPass will send a RADIUS CoA message to the switch, it is important that the switch is configured to accept these RADIUS messages. Typically, it is the switch that sends the RADIUS Request packets to the RADIUS server, and the RADIUS server responds to these request packets.

In case of the CoA, it is the RADIUS server that initiates the packet, so it is considered to be the CoA client, while the switch receives the CoA packet, so it is considered to be the CoA server.

8. Enable radius tracking credentials. RADIUS tracking is a feature that allows the switch to send radius test requests on a regular basis, so it can detect if the RADIUS server is still responding. For this feature, dedicated credentials and ClearPass service can be configured, so it does not interfere with any regular production credentials or ClearPass services. The tracking test is performed every 5 minutes by default, this can also be changed if needed.

```
ICX-Tx-Access1(config)# radius-server tracking user-name icx-radius-track
password plaintext aruba123
```

9. The tracking must be enabled per radius-server host.

NOTE: Make sure to enter the correct hostname and VRF context, otherwise a new radius-server is defined in the configuration.

```
ICX-Tx-Access1(config)# radius-server host cppm.arubatraining.com vrf mgmt
tracking enable
```

10. Access PC1 and open a web browser connection to the ClearPass host (https://10.254.1.23/tips). Login with these credentials:
    - Username**: icx-adminX** (replace X with your table number)

      e.g.: *icx-admin12* for Table 12, *icx-admin1* for Table1
    - Password: **aruba123**

11. Navigate to **Monitoring > Live Monitoring > Access Tracker**. At the right-top, select **'Edit'** to change the columns.

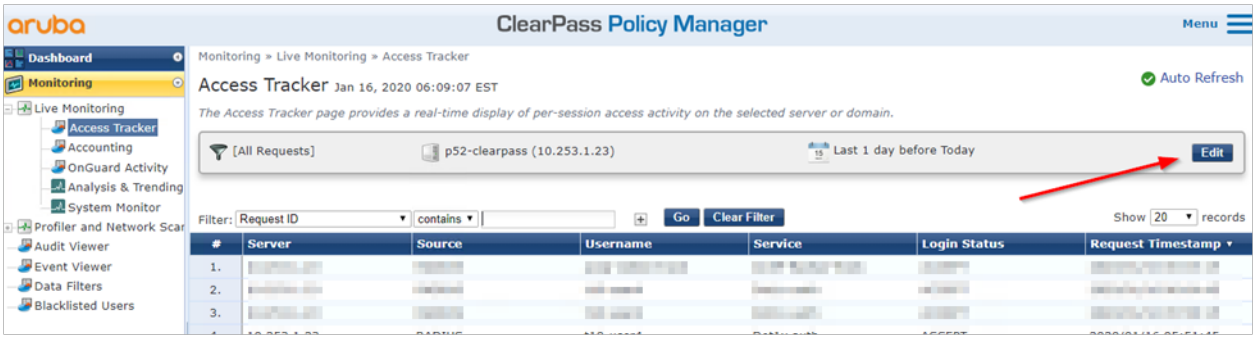

12. Add the **'NAS IP Address'** and the **'Host MAC Address'** columns to the View and arrange them in a logical order (see image for an example).

NOTE: In some remote lab environments, these column settings may have been applied previously. If the columns **'NAS IP Address'** and **'Host MAC Address'** have been configured already, move to the next step.

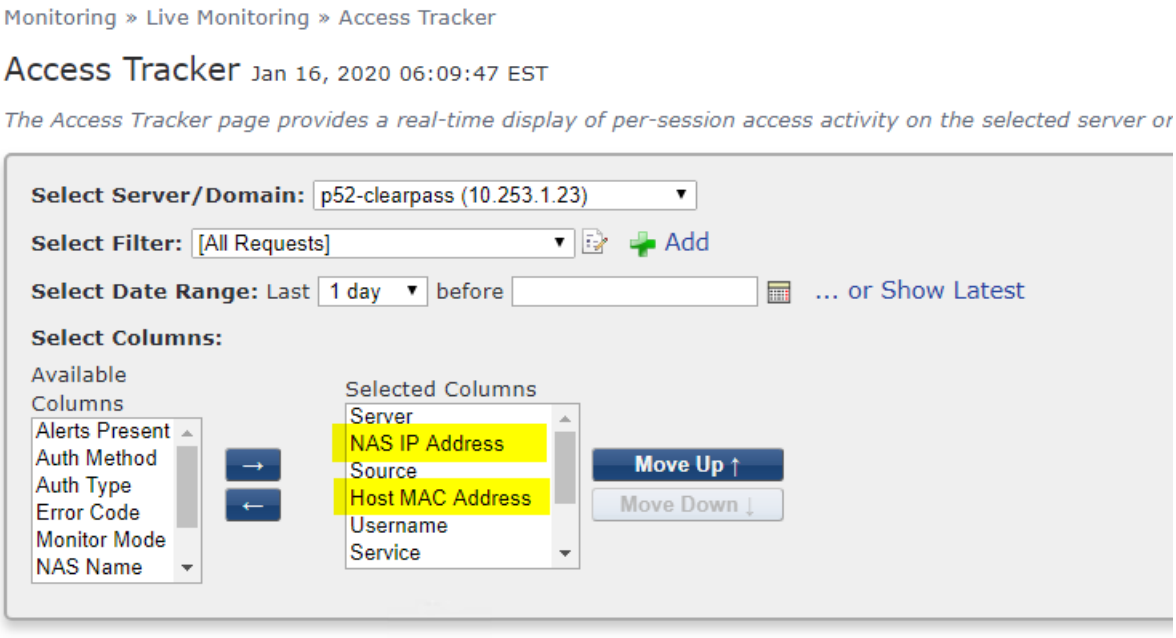

13. **Optional step**: Add **'Enforcement Profiles'** as the last column at the end and click **Save**.

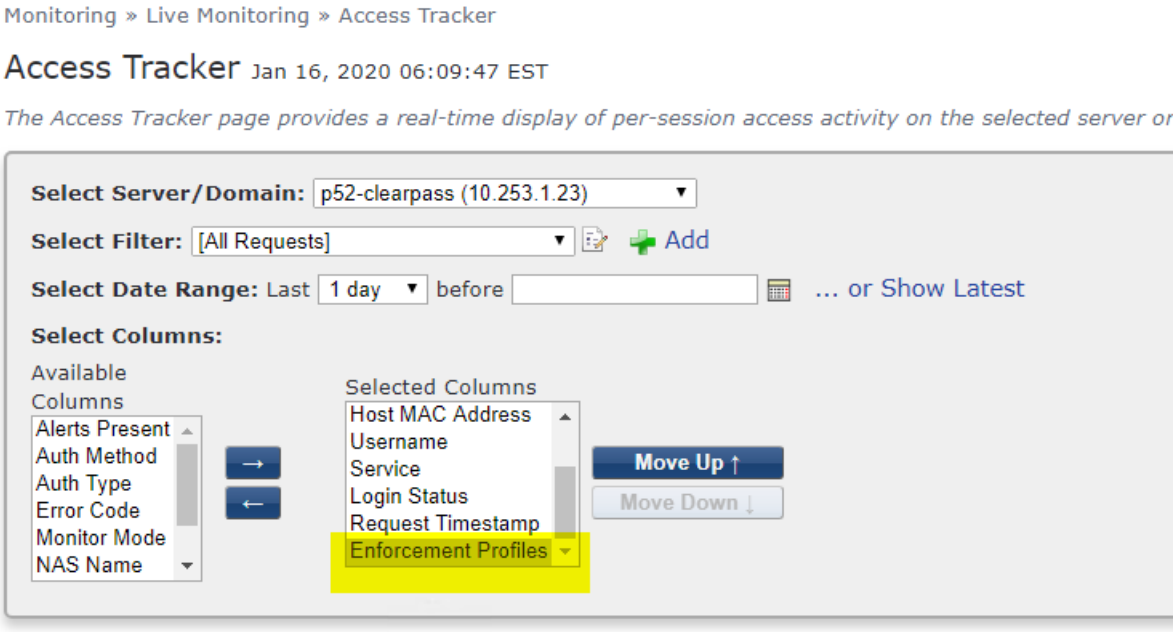

14. The switch from your table can be recognized based on the NAS IP Address in the view now (10.251.X.4). There should be an 'ACCEPT' message displayed for your Access1 switch.

Filter Access Tracker

15. To make sure you only see the authentication requests of your own table, you can apply a display filter based on your NAS-IP of your table (so 10.251.x (replace x with your table #) is sufficient to differentiate from the other tables).

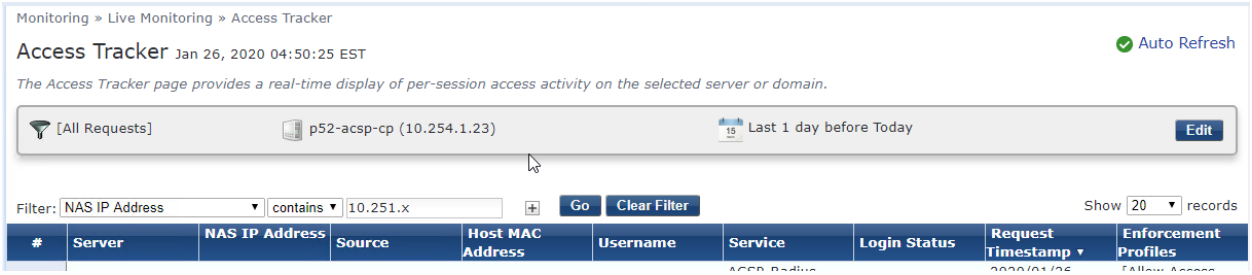

If there is no request from your switch, apply these troubleshooting steps:

- Verify the IP DNS entry on the switch.
- Verify the RADIUS server entry on the switch is using the correct name.
- Verify on the switch if a ping to this name works.
- Verify the RADIUS secret on the switch.

If there is a Reject message, apply these troubleshooting steps:

- Verify the RADIUS tracking username and password on the switch.

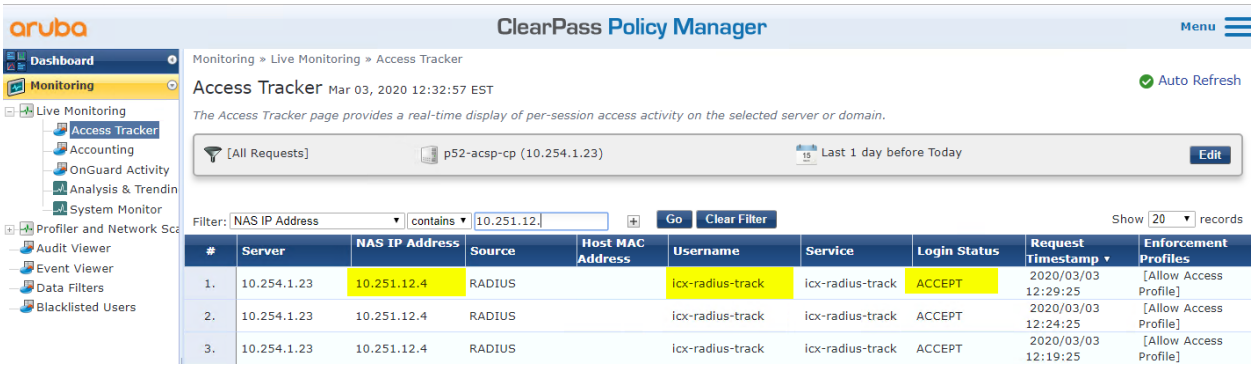

16. Click on your switch’s entry in ClearPass and review the 'Input' and 'Output' tabs

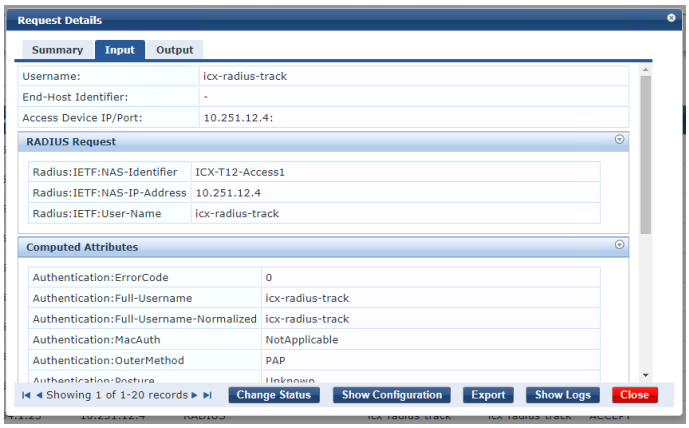

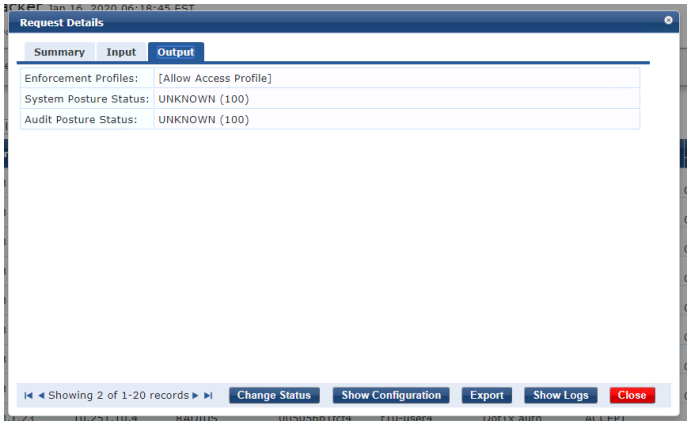

17. On the Access1 switch, check the RADIUS tracking status.

```
ICX-Tx-Access1(config)# show radius-server detail
******* Global RADIUS Configuration *******

Shared-Secret: None
Timeout: 5
Auth-Type: pap
Retries: 1
Tracking Time Interval (seconds): 300
Tracking Retries: 1
Tracking User-name: icx-radius-track
Tracking Password:
AQBapUmeqwuSjUoetq4KWXbTnUyBILPjxzok4qzRZeSXsBIzCQAAACMxHv6lEnY7jA==
Number of Servers: 1

****** RADIUS Server Information ******
Server-Name               : cppm.arubatraining.com
Auth-Port                 : 1812
Accounting-Port           : 1813
VRF                       : mgmt
Shared-Secret             :
AQBapVWcNJavUClNBQenFaJwwRrR+nWcJUvsQlHUbuaiOvlDCAAAAMCnYwT2Ful+
Timeout (default)         : 5
Retries (default)         : 1
Auth-Type (default)       : pap
Server-Group              : cppm
```

```
Group-Priority           : 1
Tracking                 : enabled
Reachability-Status      : reachable
ClearPass-Username       :
ClearPass-Password       : None

ICX-Tx-Access1(config)#
```

18. Review the RADIUS server statistics.

```
ICX-Tx-Access1(config)# show radius-server statistics authentication
 Server Name     : cppm.arubatraining.com
 Auth-Port       : 1812
 Accounting-Port : 1813
 VRF             : mgmt

  Authentication Statistics
  -------------------------
    Round Trip Time           : 3
    Pending Requests          : 0
    Timeouts                  : 0
    Bad Authenticators        : 0
    Packets Dropped           : 0
    Access Requests           : 3
    Access challenge          : 0
    Access Accepts            : 3
    Access Rejects            : 0
    Access Response Malformed : 0
    Access Retransmits        : 0
    Tracking Requests         : 3
    Tracking Responses        : 3
    Unknown Response Code     : 0
```

In case that the switch does not receive a RADIUS response from the RADIUS server, the switch will consider it 'unreachable' and it will use another RADIUS server, if defined.

> **NOTE:** The switch does not need a RADIUS Accept packet for the tracking to be successful. Any RADIUS response, both Accept and Reject, results in the tracking feature to mark the RADIUS server as reachable.

Here is an example output for failed tracking, this is for reference only. In this example, the switch did not receive any response from the RADIUS server.

```
ICX-Tx-Access1(config)# show radius-server statistics authentication
 Server Name     : cppm.arubatraining.com
 Auth-Port       : 1812
 Accounting-Port : 1813
 VRF             : mgmt
```

```
  Authentication Statistics
  -------------------------
    Round Trip Time           : 0
    Pending Requests          : 0
    Timeouts                  : 4
    Bad Authenticators        : 0
    Packets Dropped           : 0
    Access Requests           : 10
    Access challenge          : 0
    Access Accepts            : 4
    Access Rejects            : 2
    Access Response Malformed : 0
    Access Retransmits        : 3
    Tracking Requests         : 7
    Tracking Responses        : 6
    Unknown Response Code     : 0
```

```
ICX-Tx-Access1(config)# show radius-server detail
******* Global RADIUS Configuration *******

Shared-Secret: None
Timeout: 5
Auth-Type: pap
Retries: 1
Tracking Time Interval (seconds): 300
Tracking Retries: 1
Tracking User-name: icx-radius-track
Tracking Password:
AQBapdsoTpeGwAuE0Yi7Kbh7A+z2T0kIGViZ+ESpcz43A1ApCwAAAL77YGoIkfiI9BIC
Number of Servers: 1

****** RADIUS Server Information ******
Server-Name               : cppm.arubatraining.com
Auth-Port                 : 1812
Accounting-Port           : 1813
VRF                       : mgmt
Shared-Secret             :
AQBapRkC49b8TnigvMULoSuIJ2ndGHK9Tnn6RCU6I0bG3Ft3AwAAAPOclQ==
Timeout (default)         : 5
Retries (default)         : 1
Auth-Type (default)       : pap
Server-Group              : cppm
Group-Priority            : 1
Tracking                  : enabled
Reachability-Status       : unreachable
ClearPass-Username        :
ClearPass-Password        : None

ICX-Tx-Access1(config)#
```

Task 3: Basic 802.1X Authentication with a Single User

Diagram

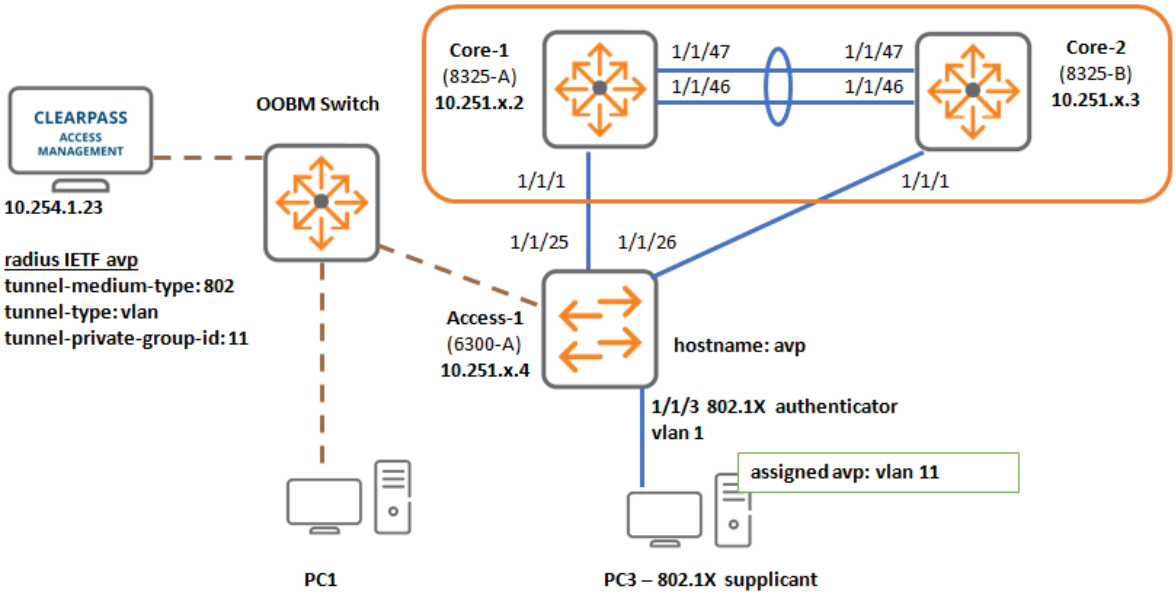

Objectives

In this task, the PC3 (connected to Access1) will be performing 802.1X authentication.

The purpose of this lab is to demonstrate the basic 802.1X authentication process and some of the RADIUS attributes that a RADIUS server can return to an AOS-CX switch in order to control the client sessions. This task will not be using the role-based configuration yet; that will be done in an upcoming task.

In order for ClearPass to process this authentication request differently from the upcoming authentication with the user-roles, the hostname of the switch will be changed. The hostname of the switch is included in the RADIUS request to the RADIUS server as the 'NAS Identifier'. A service on the ClearPass server has been pre-configured to look for the string 'avp' in the 'NAS Identifier'.

For this task, the ClearPass server will need to return some standard RADIUS Attribute Value Pairs (AVP); therefore, the hostname should also contain the string 'avp' to match this service.

Steps

1. Change the hostname of the Access1 switch for this lab task. Replace x with your table number

Example for table 12: .ICX-T12-Access1-avp, for table1: ICX-T1-Access1-avp

> **NOTE:** Make sure the string '**avp**' is present in your hostname.

```
ICX-Tx-Access1(config)# hostname ICX-Tx-Access1-avp
```

NOTE: The hostname is used as the NAS-identifier in the access-request to ClearPass. ClearPass has been pre-configured to look for the value '**avp**' in the NAS-identifier to return standard IETF **Attribute Value Pairs** with VLAN instructions. This lab depends on these instructions.

2. Enable 802.1X authentication on the switch and ensure it is using the previously defined RADIUS server group.

```
ICX-Tx-Access1-avp(config)# aaa authentication port-access dot1x authenticator
radius server-group cppm
ICX-Tx-Access1-avp(config)# aaa authentication port-access dot1x authenticator
enable
```

3. Assign port 1/1/3 to VLAN 1. In this lab, it is simply used as the base landing VLAN in case the RADIUS server does not assign a VLAN. There is no DHCP server or no other resources in this VLAN in this example.

```
ICX-Tx-Access1-avp(config)# interface 1/1/3
ICX-Tx-Access1-avp(config-if)# vlan access 1
```

4. Enter the 802.1X authenticator context and enable 802.1X on the port.

```
ICX-Tx-Access1-avp(config-if)# aaa authentication port-access dot1x authenticator
ICX-Tx-Access1-avp(config-if-dot1x-auth)# enable
ICX-Tx-Access1-avp(config-if-dot1x-auth)# exit
ICX-Tx-Access1-avp(config-if)# exit
ICX-Tx-Access1-avp(config)#
```

5. On PC3 (connected to the Access1 port 1/1/3), ensure that the Windows Service '**Wired AutoConfig**' is started. This service is not started by default. Open a command prompt with administrator rights, enter **'net start dot3svc'**

TIP: In the **ICX-Files** folder on the **desktop**, a script can be executed to start or stop this service with administrator credentials
lab 10 - dot1x - supplicant - services - start.cmd
lab 10 - dot1x - supplicant - services - stop.cmd

Configure the Wired 'Lab NIC' with 802.1X Authentication

6. Open **Network Connections** (Click **Start** > **Settings** > **Network & Internet** > **Ethernet** > **Change adapter options**).

7. Under network connections, open the Properties of the **'Lab NIC'**. Click the **'Authentication'** tab and select **'Enable IEEE 802.1X authentication'.**

> **NOTE:** If the **'Authentication'** tab is not available, the 'Wired AutoConfig' service has not been started on the client PC. Check the previous steps to ensure the service is started.

8. Open **'Settings'** for the 'PEAP' method and **uncheck** 'Verify the server certificate'.

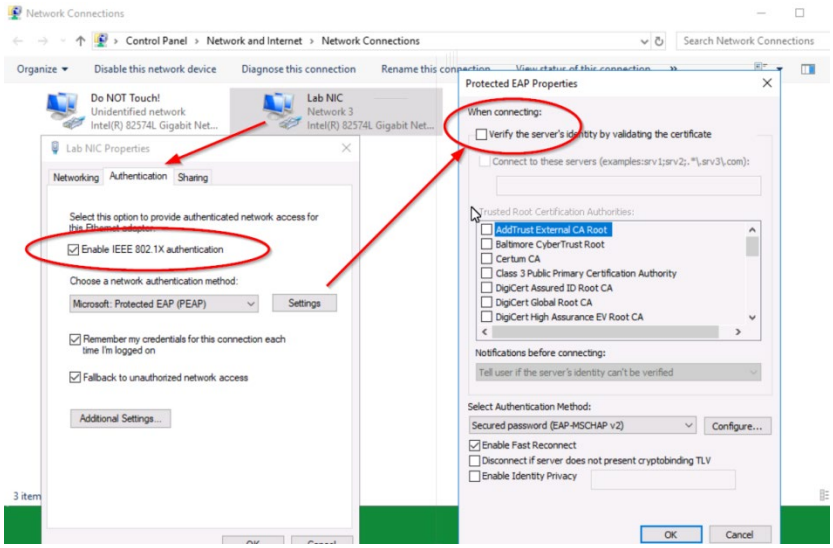

9. Click **'Configure'** next to the 'EAP-MSCHAP v2' method. Make sure the option 'Automatically use my Windows logon name' is **unchecked.** Click **'OK'** to close the window.

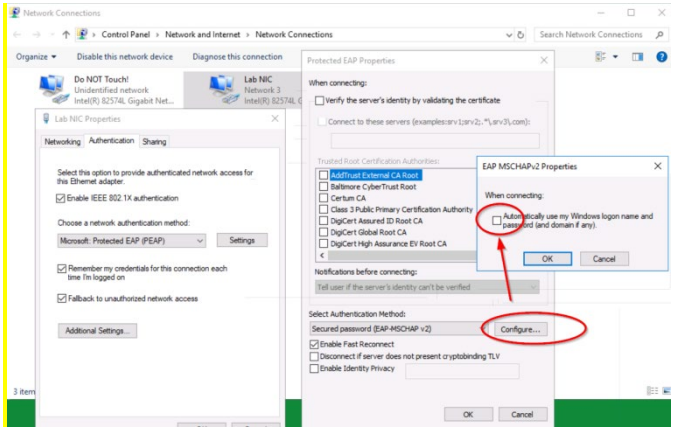

10. Click **'OK'** to close the 'Protected EAP Properties',
11. Click **'Additional Settings'**, check '**Specify authentication mode'** and select **'User authentication**'.

12. Click **'Save credentials'** to enter the credentials.

    username: **icx-avp1**

    password: **aruba123**

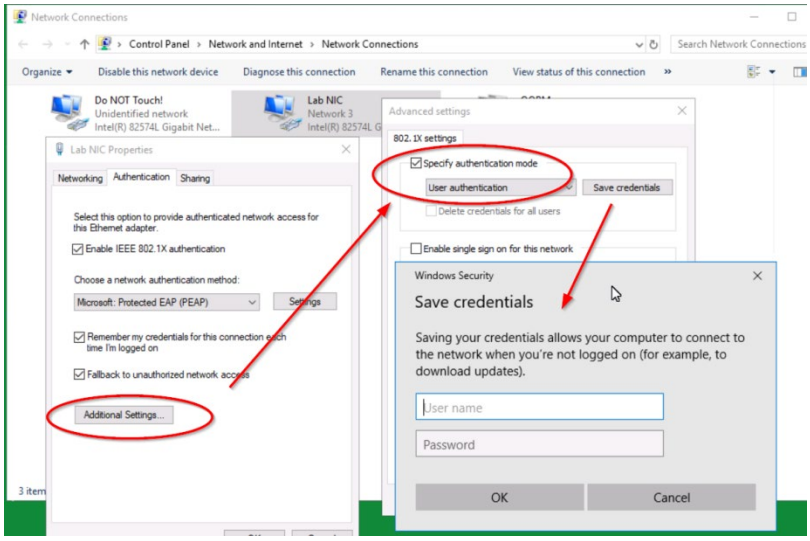

13. Click **'OK'** to save the credentials. Click **'OK'** to close the 'Advanced settings'. Click **'OK'** to close the 'Lab NIC Properties'.

14. Using PC1 (the OOBM mgmt PC), check Access Tracker in ClearPass.

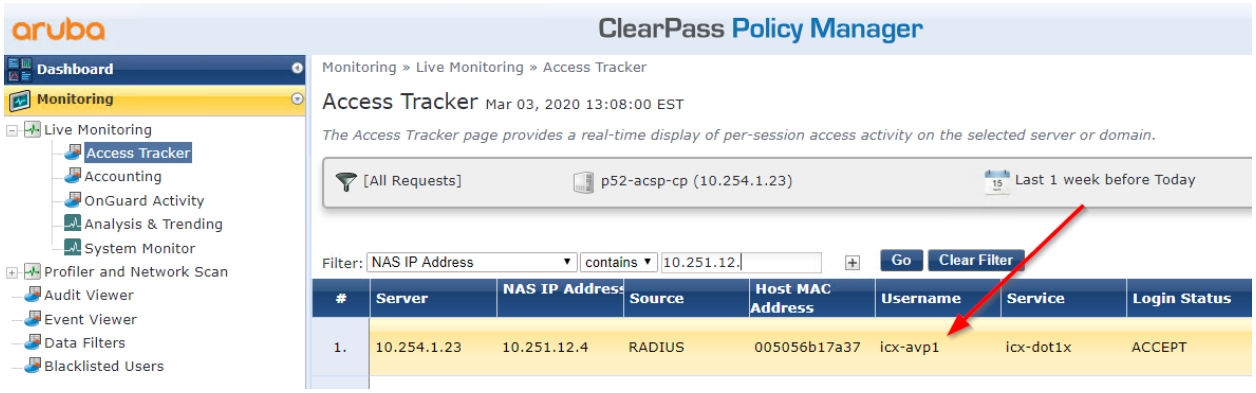

15. Open your authentication entry (make sure to check it has your NAS IP), click **'Input'** and review the attributes that were sent by the switch in the access-request.

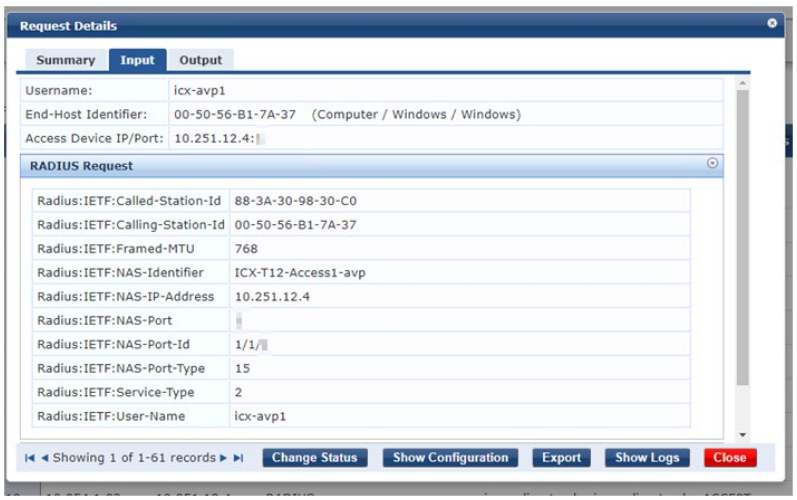

16. Next review the **'Output'** tab. This shows that ClearPass is returning the IETF VLAN attributes and a session timeout of 180 seconds to the switch

> **NOTE:** The value of 180 seconds was used in the lab environment to demonstrate the re-authentication. In a real deployment this value would be controlled by the customer security policy, this may be several hours.

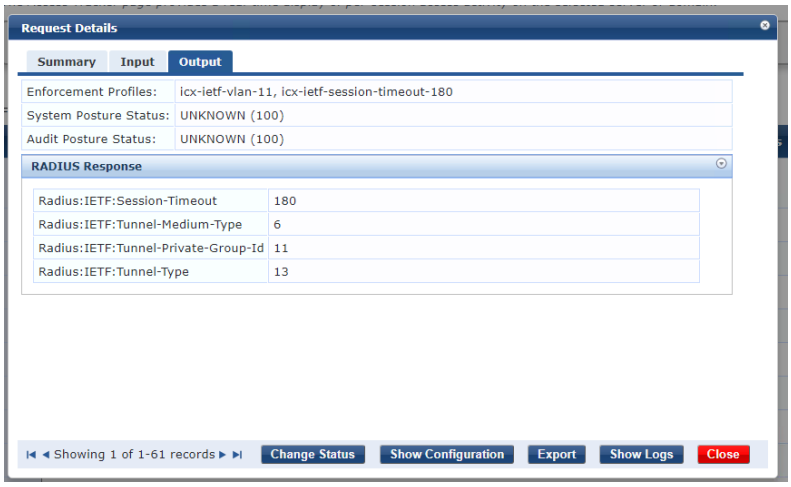

17. On the switch, check the authenticated clients. Take note of the MAC address of the authenticated client.

```
ICX-Tx-Access1-avp(config)# show aaa authentication port-access dot1x
authenticator interface all client-status

Client 00:50:56:b1:7a:37, icx-avp1, 1/1/3
==========================================

  Authentication Details
  ----------------------
    Status                        : Authenticated
```

```
    Type                          : Pass-Through
    EAP-Method                    : PEAP
    Auth Failure reason           :
    Time Since Last State Change  : 66s

  Authentication Statistics
  -------------------------
    Authentication                          : 9
    Authentication Timeout                  : 0
    EAP-Start While Authenticating          : 2
    EAP-Logoff While Authenticating         : 0
    Successful Authentication               : 1
    Failed Authentication                   : 7
    Re-Authentication                       : 0
    Successful Re-Authentication            : 0
    Failed Re-Authentication                : 0
    EAP-Start When Authenticated            : 0
    EAP-Logoff When Authenticated           : 0
    Re-Auths When Authenticated             : 0
    Cached Re-Authentication                : 0
```

18. Review the mac-address table. The client MAC address should have been assigned to VLAN 11 by port-access-security. Remember that the default VLAN of the interface 1/1/3 has been set to VLAN 1.

```
ICX-Tx-Access1-avp(config)# show mac-address-table
MAC age-time            : 300 seconds
Number of MAC addresses : 10

MAC Address          VLAN     Type                       Port
-------------------------------------------------------------
20:4c:03:5f:98:02    1        dynamic                    lag256
88:3a:30:97:b6:00    1        dynamic                    lag256
00:50:56:b1:7a:37    11       port-access-security       1/1/3
02:02:00:00:12:00    11       dynamic                    lag256
90:20:c2:bc:17:00    11       dynamic                    lag256
90:20:c2:bc:97:00    11       dynamic                    lag256
90:20:c2:bc:97:00    12       dynamic                    lag256
02:02:00:00:12:00    12       dynamic                    lag256
88:3a:30:97:b6:00    12       dynamic                    lag256
90:20:c2:bc:17:00    12       dynamic                    lag256
ICX-Tx-Access1-avp(config-if)#
```

19. Review the switch log for entries of the **'port-accessd'** process. After a few minutes(180seconds), the user will be re-authenticated based on the RADIUS session-timeout. These are the options used:

    -r show log in reverse order

    -n 4 last 4 lines of the log

-d filter on daemon (process)

```
ICX-Tx-Access1-avp(config)# show logging -r -n 4 -d port-accessd
-------------------------------------------------
Event logs from current boot
-------------------------------------------------
2020-03-03T18:01:27.948757+00:00 ICX-T12-Access1-avp port-accessd[2929]:
Event|10503|LOG_INFO|MSTR|1|Port 1/1/3 is unblocked by port-access
2020-03-03T18:01:27.775739+00:00 ICX-T12-Access1-avp port-accessd[2929]:
Event|10502|LOG_INFO|MSTR|1|Port 1/1/3 is blocked by port-access
2020-03-03T17:58:28.654705+00:00 ICX-T12-Access1-avp port-accessd[2929]:
Event|10503|LOG_INFO|MSTR|1|Port 1/1/3 is unblocked by port-access
2020-03-03T17:44:29.107796+00:00 ICX-T12-Access1-avp port-accessd[2929]:
Event|10502|LOG_INFO|MSTR|1|Port 1/1/3 is blocked by
```

20. Review the client accounting information. Notice that the session time will never exceed 180 seconds due to the re-authentication interval.

```
ICX-Tx-Access1-avp(config)# show aaa accounting port-access interface all client-
status

Port Access Client Status Details

Client 00:50:56:b1:7a:37, icx-avp1
============================
  Session Details
  ---------------
    Port                : 1/1/3
    Session Time        : 45s

  Accounting Details
  ----------------------
    Accounting Session ID  : 1579185236543
    Input Packets          : 1197
    Input Octets           : 186098
    Output Packets         : 809
    Output Octets          : 82067
    Input Gigaword         : 0
    Output Gigaword        : 0


ICX-Tx-Access1-avp(config)#
```

AOS-CX Internal User Roles

AOS-CX will internally always apply the authorization settings as a role. This is a significant change compared with the ArubaOS-Switch devices, where the switch was either in 'radius' mode **or** in 'user-role' mode. With those switches it was not possible to

have some devices authenticated and authorized with RADIUS attributes, while other devices would be using a user-role.

In AOS-CX, this is solved by converting any RADIUS authorization attributes into a temporary user role of type 'radius'. Therefore, this role can co-exist with any local defined roles or with any ClearPass downloadable user roles.

21. On the Access1 switch, review the port-access roles, note the type of the role. This shows which radius attributes were returned and converted into this temporary user role.

```
ICX-Tx-Access1-avp(config)# show port-access role

Role Information:

Name  : RADIUS_496422333
Type  : radius
-------------------------------------------
    Reauthentication Period             :
    Authentication Mode                 :
    Session Timeout                     : 180 secs
    Client Inactivity Timeout           :
    Description                         :
    Gateway Zone                        :
    UBT Gateway Role                    :
    Access VLAN                         : 11
    Native VLAN                         :
    Allowed Trunk VLANs                 :
    MTU                                 :
    QOS Trust Mode                      :
    PoE Priority                        :
    Captive Portal Profile              :
    Policy                              :
```

22. Next review the authenticated users on the interfaces. Notice the difference with the previous command client-status:

To show 802.1X authenticated clients

show aaa authentication port-access **dot1x** authenticator interface all client-status

To show the role authorization

show aaa authentication port-access **interface** all client-status

```
ICX-Tx-Access1-avp(config)# show aaa authentication port-access interface all
client-status

Port Access Client Status Details

Client 00:50:56:b1:7a:37, icx-avp1
==========================
  Session Details
  ---------------
```

```
    Port         : 1/1/3
    Session Time : 123s

  Authentication Details
  ----------------------
    Status          : dot1x Authenticated
    Auth Precedence : dot1x - Authenticated, mac-auth - Not attempted

  Authorization Details
  ----------------------
    Role   : RADIUS_496422333
    Status : Applied
```

Task 4: Change of Authorization Verification

Diagram

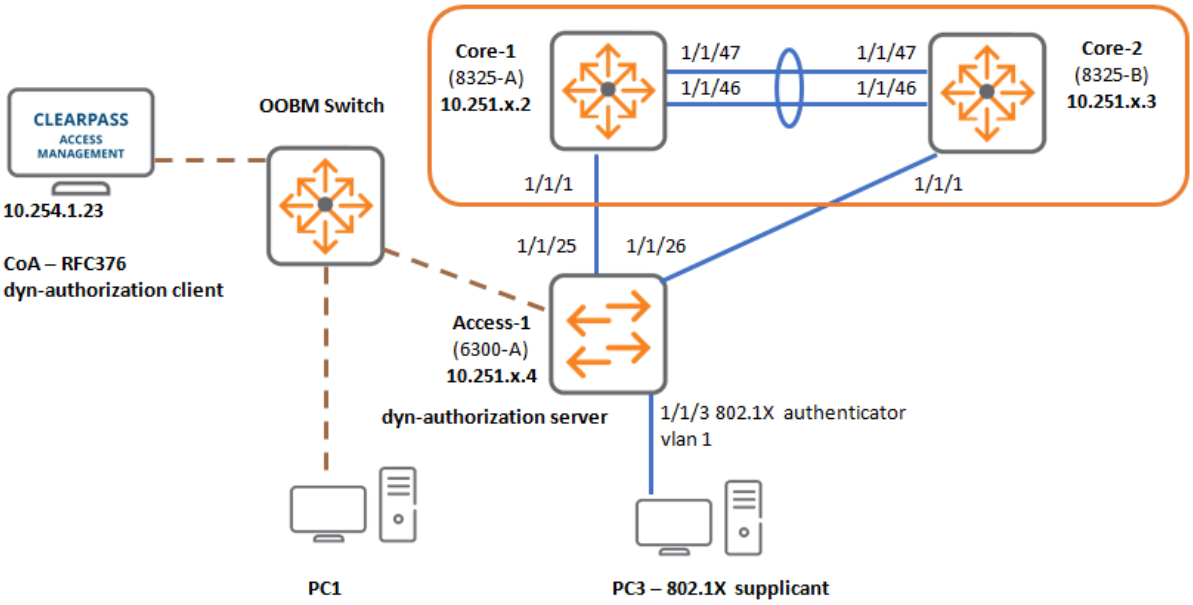

Objectives

Verify that a disconnect message from the ClearPass RADIUS host is processed correctly by the switch.

This allows the ClearPass server to dynamically trigger re-authentication on the access device when the access security policy or the access conditions would have changed.

Steps

1. First review the current online time of the 802.1X authenticated user. When the CoA disconnect message is sent in the next steps, you should see that this timer will be reset, since the CoA will trigger a new authentication.

```
ICX-Tx-Access1-avp(config-if)# show aaa authentication port-access interface all
client-status

Port Access Client Status Details

Client 00:50:56:b1:7a:37, icx-avp1
============================
  Session Details
  ---------------
    Port         : 1/1/3
    Session Time : 136s
```

```
Authentication Details
----------------------
  Status          : dot1x Authenticated
  Auth Precedence : dot1x - Authenticated, mac-auth - Not attempted

Authorization Details
---------------------
  Role   : RADIUS_496422333
  Status : Applied
```

2. On the ClearPass system, navigate to **Monitoring > Live Monitoring > Access Tracker**.
3. Open the latest authentication event from your own Access switch (use the NAS IP to find your own authentication session).
4. On the **'Summary'** page, use the **'Change Status'** to get access to the CoA options.

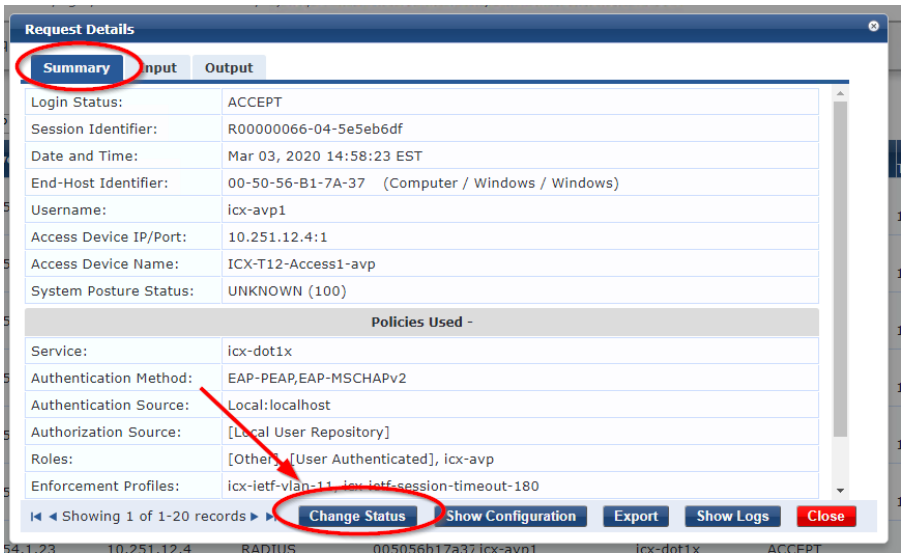

5. Use **'Submit'** to test the CoA.

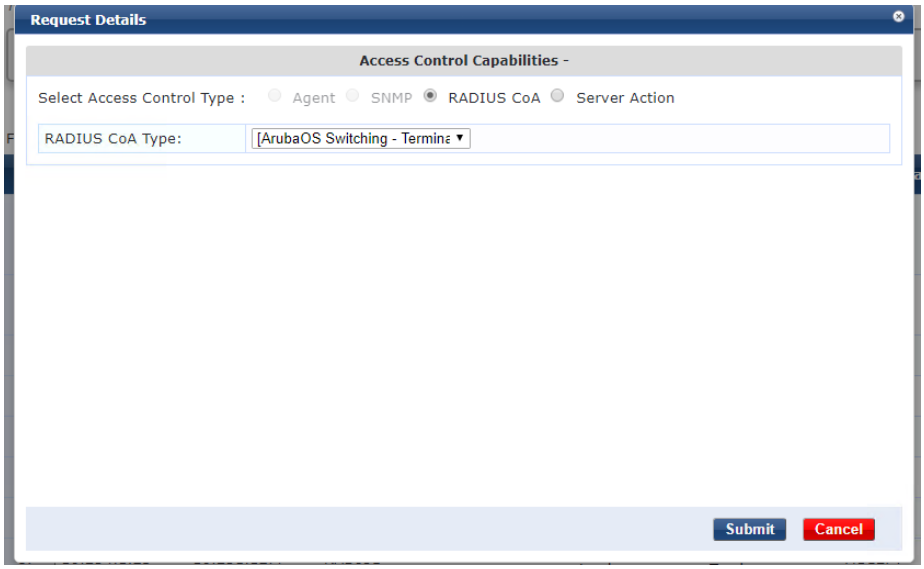

6. After a few seconds, a 'successful' message should be displayed.

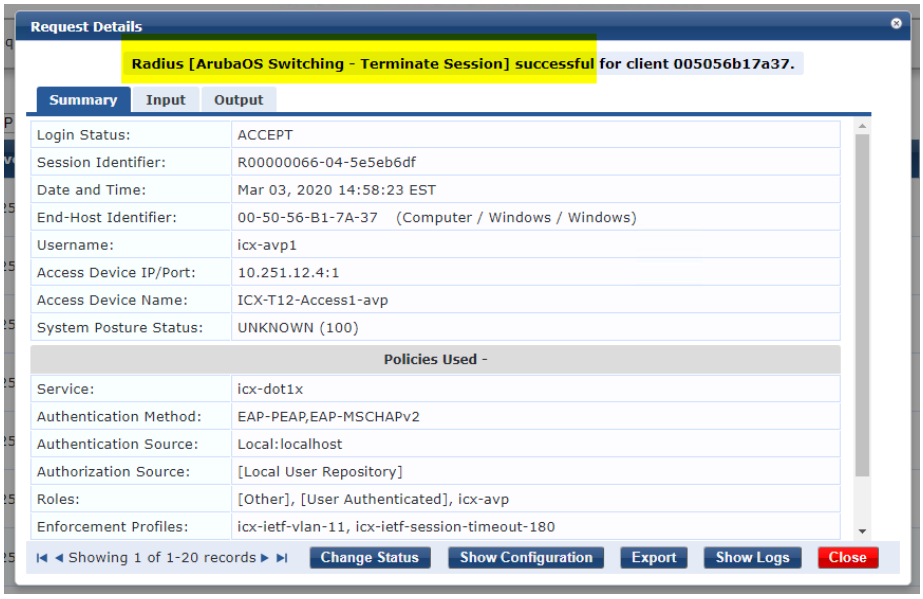

NOTE: If the CoA is not successful, you should check the dynamic authorization configuration on the switch. This was done in Task2.

7. In the Access tracker list, a new authentication entry should show up a few seconds later.

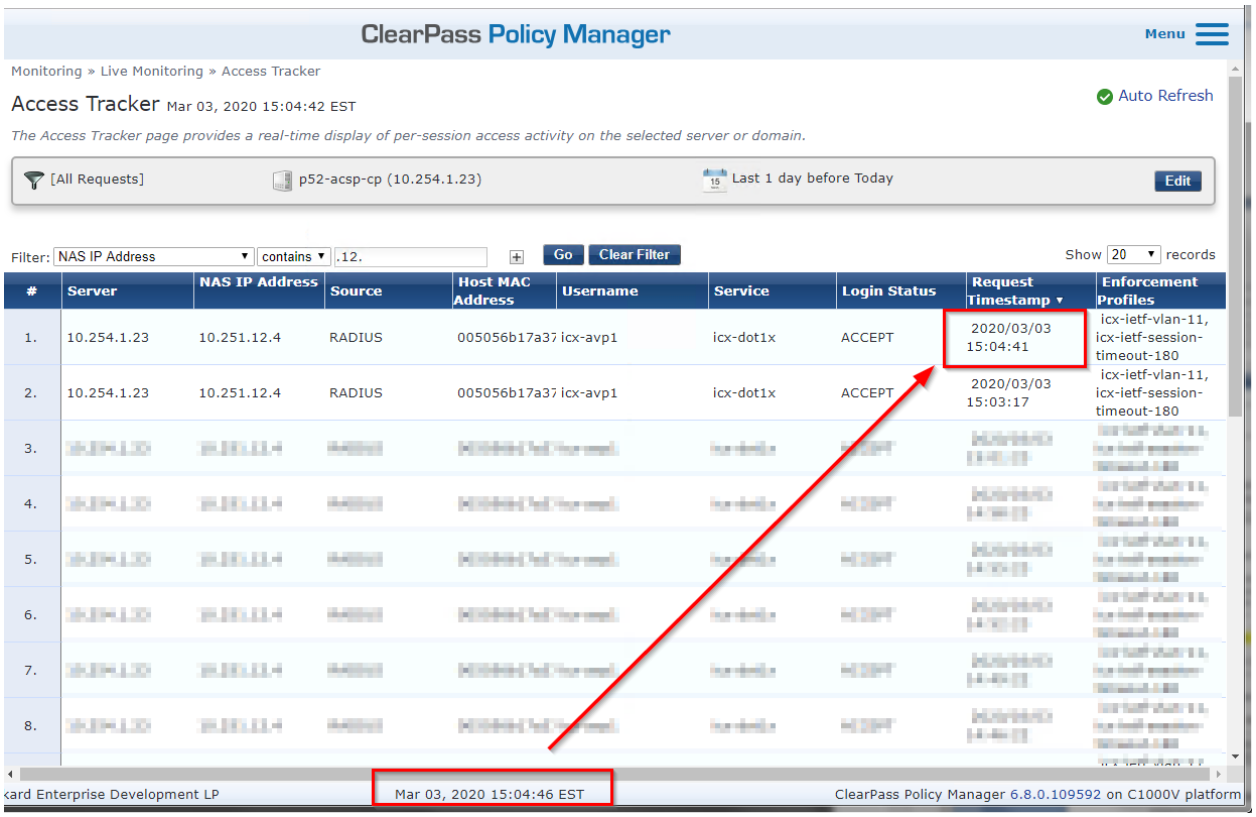

8. On the Access1 switch, review the 'online time' for the authenticated user. The timer should have been reset.

```
ICX-Tx-Access1-avp(config-if)# show aaa authentication port-access interface all
client-status

Port Access Client Status Details

Client 00:50:56:b1:7a:37, icx-avp1
==============================
  Session Details
  ---------------
    Port         : 1/1/3
    Session Time : 10s

  Authentication Details
  ----------------------
    Status          : dot1x Authenticated
    Auth Precedence : dot1x - Authenticated, mac-auth - Not attempted

  Authorization Details
  ---------------------
    Role   : RADIUS_496422333
    Status : Applied
```

9. Review the statistics of 'dynamic authorization'.

```
ICX-Tx-Access1-avp(config)# show radius dyn-authorization
Status and Counters - RADIUS Dynamic Authorization Information

  RADIUS Dynamic Authorization                  : Enabled
  RADIUS Dynamic Authorization UDP Port         : 3799
  Invalid Client Addresses in CoA Requests      : 0
  Invalid Client Addresses in Disconnect Requests: 0

Dynamic Authorization Client Information
========================================

IP Address          : cppm.arubatraining.com
VRF                 : mgmt
Replay Protection   : Disabled
Time Window         : 300
Disconnect Requests : 1
Disconnect ACKs     : 1
Disconnect NAKs     : 0
CoA Requests        : 0
CoA ACKs            : 0
CoA NAKs            : 0
Shared-Secret       :
AQBapVWcNJavUClNBQenFaJwwRrR+nWcJUvsQlHUbuaiOvlDCAAAAMCnYwT2Ful+
ICX-Tx-Access1-avp(config)#
```

NOTE: The actual statistics numbers may be different in your setup.

10. And the details for the CPPM client.

```
ICX-Tx-Access1-avp(config)# show radius dyn-authorization client
cppm.arubatraining.com vrf mgmt
Status and Counters - RADIUS Dynamic Authorization Client Information

  VRF Name                : mgmt
  Authorization Client    : cppm.arubatraining.com
  Unknown Packets         : 0
  Message-Type                 Disconnect       CoA
  ------------------------------------------------------------
  Total Requests               1                0
  Authorize Only Requests      0                0
  Malformed Requests           0                0
  Bad Authenticator Requests   0                0
  Dropped Requests             0                0
  Total ACK Responses          2                0
  Total NAK Responses          0                0
  Session Not Found Responses  0                0
  User Sessions Modified       1                0
ICX-Tx-Access1-avp(config)#
```

Task 5: Aruba User Role Based Access

Diagram

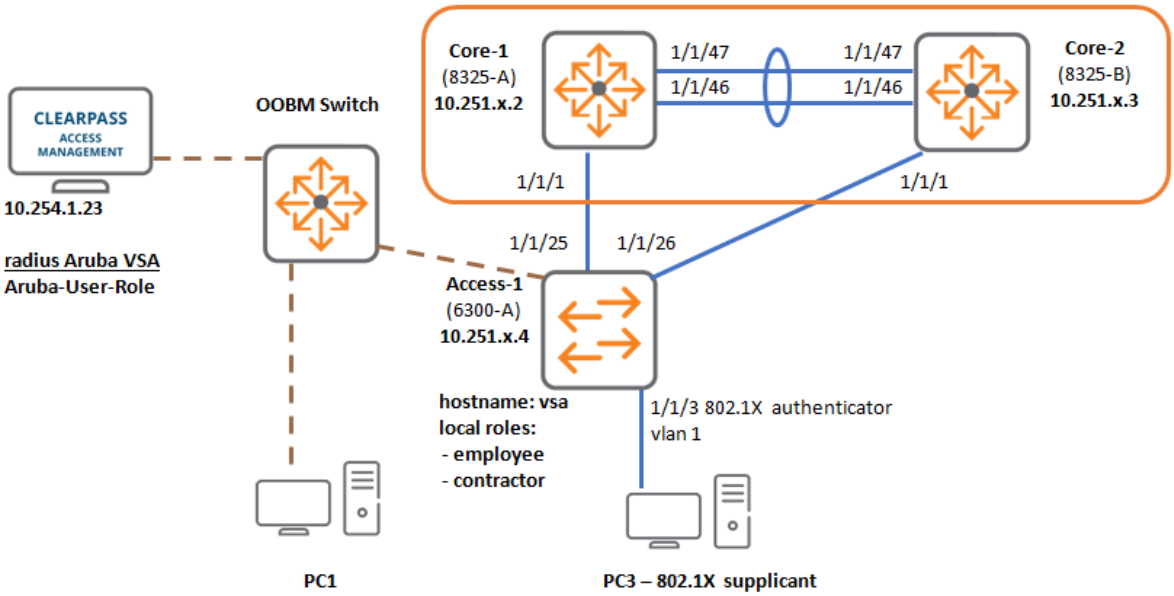

Objectives

In this task, you will explore how Aruba User Roles will simplify how settings can be assigned to an authenticated user or device. On the switch, user roles will be defined for an employee and contractor, each role will have different network access.

The ClearPass server will return this role to the switch when the users complete successful authentication. In order for ClearPass to return the correct Aruba VSA (Vendor-Specific-Attribute) 'Aruba-User-Role', you will change the hostname of the switch, so a different policy will be applied by the ClearPass server for this task.

To define the roles, you will first define classes and policies. These can control access to the network.

- The role employee will be assigned to VLAN 11, with an unrestricted ACL.
- The role contractor will be assigned to the same VLAN 11, but with a restricted ACL.

Steps

Define Traffic Classes

Each user role will be configured with unique access restrictions to the network. To control network access, the administrator first needs to describe the traffic classes.

These are global objects that can be used in multiple policies.

1. Open a terminal session to Access1 and enter the configuration mode.

2. Define a new class for 'any' traffic.

```
ICX-Tx-Access1-avp(config)# class ip any
ICX-Tx-Access1-avp(config-class-ip)# match any any any
ICX-Tx-Access1-avp(config-class-ip)# exit
```

3. Define a new class for 'servers'. This is for lab purpose only, so two IP addresses of the Core switches are used. The class is using the /255.0.255.255 to handle the unique Table x in each table IP address.

```
ICX-Tx-Access1-avp(config)# class ip servers
ICX-Tx-Access1-avp(config-class-ip)# match any any 10.0.1.2/255.0.255.255
ICX-Tx-Access1-avp(config-class-ip)# match any any 10.0.1.3/255.0.255.255
```

4. Notice how the system automatically adds line numbers for each entered line.

```
ICX-Tx-Access1-avp(config-class-ip)# show running  current-context
class ip servers
    10 match any any 10.0.1.2/255.0.255.255
    20 match any any 10.0.1.3/255.0.255.255
```

5. The sequencing makes it easy to remove a single line, or to insert a line between two existing lines. Remove the second line again and verify 1 server is listed.

```
ICX-Tx-Access1-avp(config-class-ip)# no 20
ICX-Tx-Access1-avp(config-class-ip)# show running current-context
class ip servers
    10 match any any 10.0.1.2/255.0.255.255
ICX-Tx-Access1-avp(config-class-ip)# exit
```

Define Policy for Employee

These are not the actual roles yet. A policy assigns an action to a traffic class.

6. Define the employee policy. Employees are granted full access to the network.

```
ICX-Tx-Access1-avp(config)# port-access policy employee
ICX-Tx-Access1-avp(config-pa-policy)# class ip any
```

7. By default, when a class is added to a policy, it gets an action 'permit', but many more actions are possible, such as QOS actions. Review the available actions using the '?' option, but no action must be configured in this step (due to the default 'permit').

```
ICX-Tx-Access1-avp(config-pa-policy)# class ip any action ?
  cir              Specify the committed information rate
  drop             Drop all traffic
  dscp             Specify a Differentiated Services Code Point value
  ip-precedence    Specify the IP precedence
  local-priority   Specify a local priority value
  pcp              Specify the Priority Code Point (PCP) value
```

```
  redirect        Specify a redirect destination
  <cr>
ICX-Tx-Access1-avp(config-pa-policy)# exit
```

8. Define the contractor policy. Contractors will not be allowed access to the class 'servers', so this action is set to 'drop'. Any other traffic is permitted.

The order of the policy is important for the traffic processing. Just like with the classes, the policy object will automatically add line numbers for each entry.

```
ICX-Tx-Access1-avp(config)# port-access policy contractor
ICX-Tx-Access1-avp(config-pa-policy)#  class ip servers action drop
ICX-Tx-Access1-avp(config-pa-policy)#  class ip any
```

9. Review the current context configuration.

```
ICX-Tx-Access1-avp(config-pa-policy)# show run current-context
port-access policy contractor
    10 class ip servers action drop
    20 class ip any
ICX-Tx-Access1-avp(config-pa-policy)#
```

10. The policy also allows the administrator to add a comment to each line.

```
ICX-Tx-Access1-avp(config-pa-policy)# 10 comment No server access for Contractors
ICX-Tx-Access1-avp(config-pa-policy)# show run cur
port-access policy contractor
    10 class ip servers action drop
    10 comment No server access for Contractors
    20 class ip any
ICX-Tx-Access1-avp(config-pa-policy)# exit
```

11. Review the policy that was just created for the contractor.

```
ICX-Tx-Access1-avp(config)# show port-access policy contractor

Access Policy Details:
======================

Policy Name   : contractor
Policy Type   : Local
Policy Status : Applied

SEQUENCE    CLASS                        TYPE ACTION
----------- ---------------------------- ---- ----------------------------------
10          servers                      ipv4 drop
20          any                          ipv4 permit
```

Define the User Roles for Employee

12. Define the user role for employee. In this step, the previously defined access policy will be bound to the role. The role also contains the L2 VLAN assignment.

```
ICX-Tx-Access1-avp(config)# port-access role employee
ICX-Tx-Access1-avp(config-pa-role)# associate policy employee
ICX-Tx-Access1-avp(config-pa-role)# vlan access 11
ICX-Tx-Access1-avp(config-pa-role)# exit
```

13. Define the user role for contractor.

```
ICX-Tx-Access1-avp(config)# port-access role contractor
ICX-Tx-Access1-avp(config-pa-role)# associate policy contractor
ICX-Tx-Access1-avp(config-pa-role)# vlan access 11
ICX-Tx-Access1-avp(config-pa-role)# exit
ICX-Tx-Access1-avp(config)#
```

14. Review the defined objects. Notice that these roles are 'local' roles, while in the previous section, the **'radius'** role was automatically defined based on the RADIUS attributes. In a later lab, role definitions will be downloaded from ClearPass, these are not **'local'** roles, but will be **'clearpass'** roles.

```
ICX-Tx-Access1-avp(config)# show port-access role

Role Information:

Name  : RADIUS_496422333
Type  : radius
---------------------------------------------
    Reauthentication Period             :
    Authentication Mode                 :
    Session Timeout                     : 180 secs
    Client Inactivity Timeout           :
    Description                         :
    Gateway Zone                        :
    UBT Gateway Role                    :
    Access VLAN                         : 11
    Native VLAN                         :
    Allowed Trunk VLANs                 :
    MTU                                 :
    QOS Trust Mode                      :
    PoE Priority                        :
    Captive Portal Profile              :
    Policy                              :

Name  : contractor
Type  : local
---------------------------------------------
    Reauthentication Period             :
    Authentication Mode                 :
```

```
    Session Timeout                     :
    Client Inactivity Timeout           :
    Description                         :
    Gateway Zone                        :
    UBT Gateway Role                    :
    Access VLAN                         : 11
    Native VLAN                         :
    Allowed Trunk VLANs                 :
    MTU                                 :
    QOS Trust Mode                      :
    PoE Priority                        :
    Captive Portal Profile              :
    Policy                              : contractor

Name  : employee
Type  : local
-----------------------------------------------
    Reauthentication Period             :
    Authentication Mode                 :
    Session Timeout                     :
    Client Inactivity Timeout           :
    Description                         :
    Gateway Zone                        :
    UBT Gateway Role                    :
    Access VLAN                         : 11
    Native VLAN                         :
    Allowed Trunk VLANs                 :
    MTU                                 :
    QOS Trust Mode                      :
    PoE Priority                        :
    Captive Portal Profile              :
    Policy                              : employee
```

Test Access for User-Role Employee

Diagram

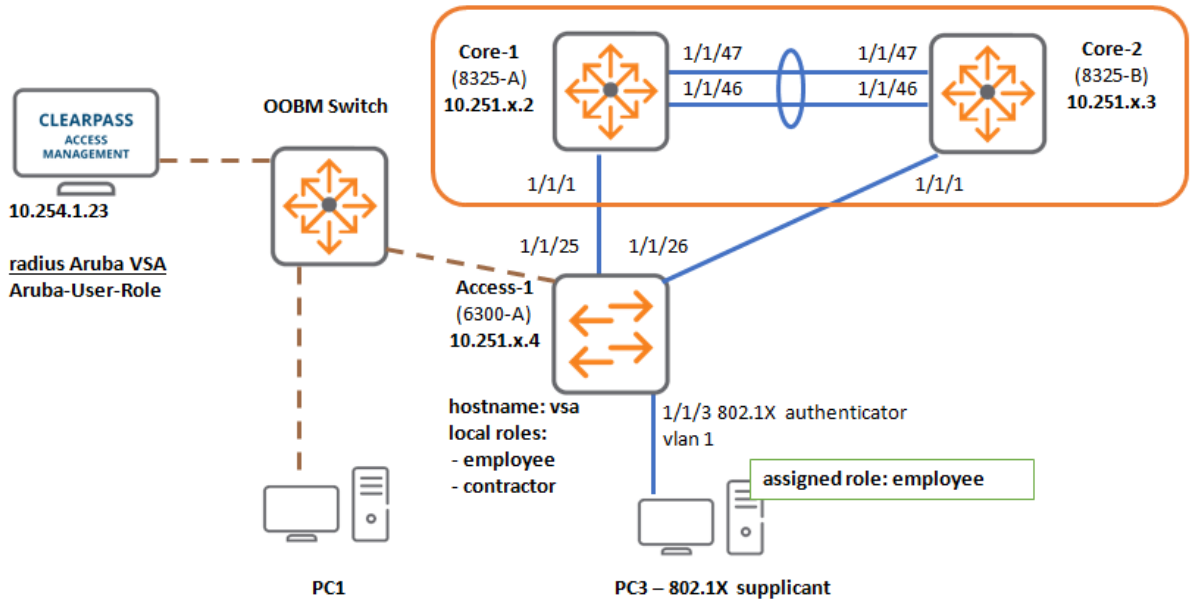

15. Change hostname string from **'avp'** to **'vsa'.** Also, replace **Tx** with your table information; for example, T1 would be Table 1 and T12 would be Table 12.

```
ICX-Tx-Access1-avp(config)# hostname ICX-Tx-Access1-vsa
ICX-Tx-Access1-vsa(config)#
```

> **NOTE:** The hostname is used as the NAS-identifier. ClearPass has been pre-configured to look for the value '**vsa**' in the NAS-identifier to return the **Aruba-User-Role** VSA attributes that are needed for this lab.

16. On PC3 (connected to Access1 port 1/1/3), change the 802.1X authentication credentials to:
    - Username: **icx-employee**
    - Password: **aruba123**

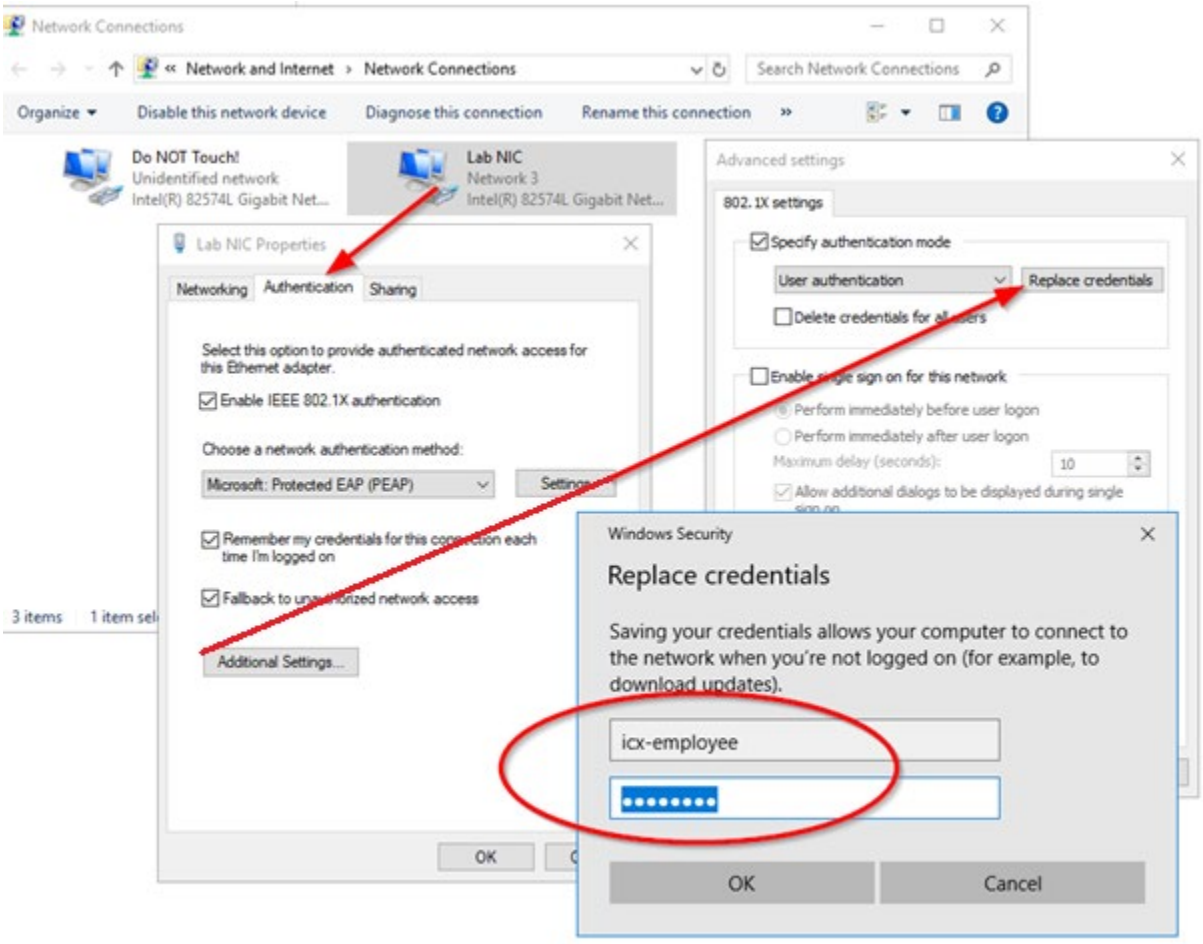

17. On the ClearPass system, navigate to Access tracker. Open your most recent authentication entry (it should have a username of ‘icx-employee’. Make sure to check the NAS-IP so you are not looking at the entry of another table. Apply the filter 'NAS IP Address' contains '10.251.x' if you only want to see your own authentication events.

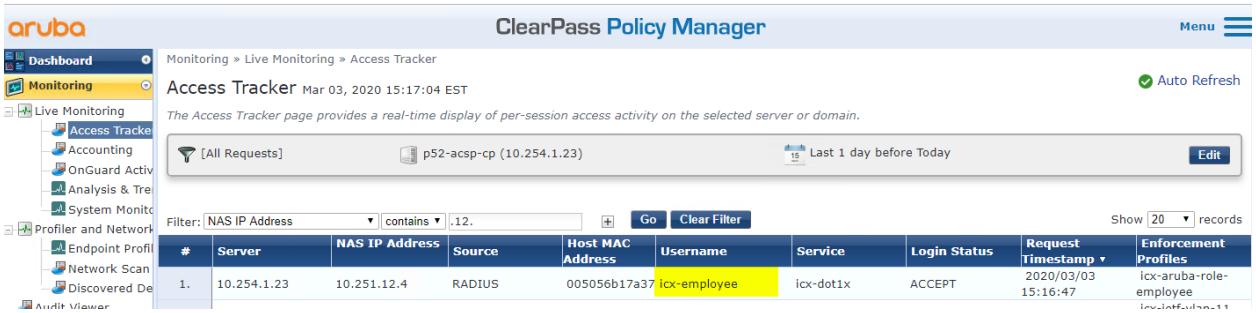

18. Check the '**Output**' of the entry and expand the “**RADIUS Response**” section. Notice that AOS-CX is now also using the Aruba RADIUS dictionary, so to push

a user role, the 'Aruba-User-Role' attribute is used, just like the Aruba Mobility Controller.

NOTE: The ArubaOS-Switch platform was using the 'Hewlett-Packard-Enterprise' RADIUS dictionary, so customers migrating to AOS-CX should update the RADIUS policy to reflect this change.

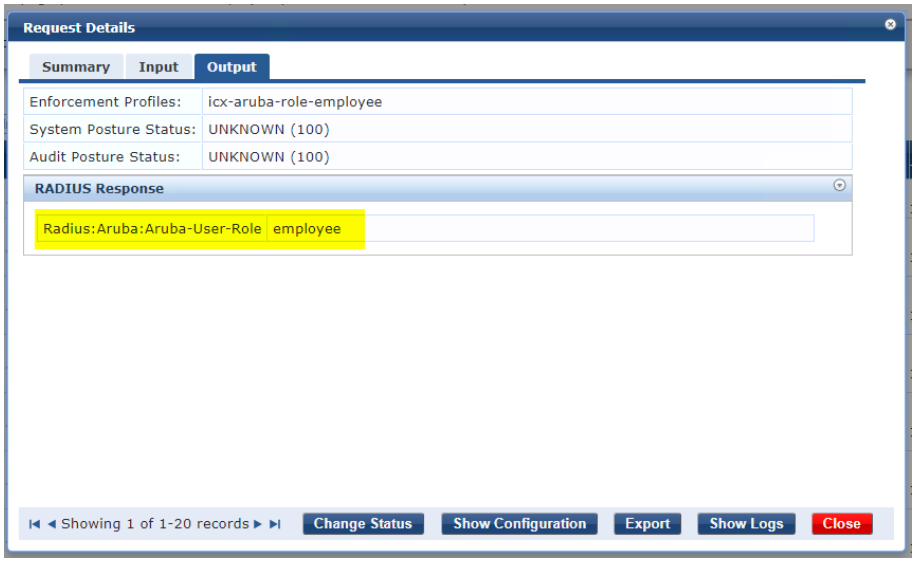

19. On the switch review the authenticated 802.1X users.

```
ICX-Tx-Access1-vsa(config)# show aaa authentication port-access dot1x
authenticator interface all client-status

Client 00:50:56:b1:7a:37, icx-employee, 1/1/3
==========================================

  Authentication Details
  ----------------------
    Status                        : Authenticated
    Type                          : Pass-Through
    EAP-Method                    : PEAP
    Auth Failure reason           :
    Time Since Last State Change  : 906s
```

20. Review the user role authorization, the user should have been assigned the role 'employee' based on the RADIUS Aruba-User-Role assignment by ClearPass.

```
ICX-Tx-Access1-avp(config)# show aaa authentication port-access interface all
client-status

Port Access Client Status Details
```

```
Client 00:50:56:b1:7a:37, icx-employee
==============================
  Session Details
  ---------------
    Port         : 1/1/3
    Session Time : 1073s

  Authentication Details
  ----------------------
    Status          : dot1x Authenticated
    Auth Precedence : dot1x - Authenticated, mac-auth - Not attempted

  Authorization Details
  ---------------------
    Role   : employee
    Status : Applied
```

21. On PC3, verify that the employee role can access the 'server' host. Ping should be successful.

```
ping 10.x.1.2
```

Optional: Test Access for User-Role Contractor

You may test the contractor access, but these are optional steps.

Diagram

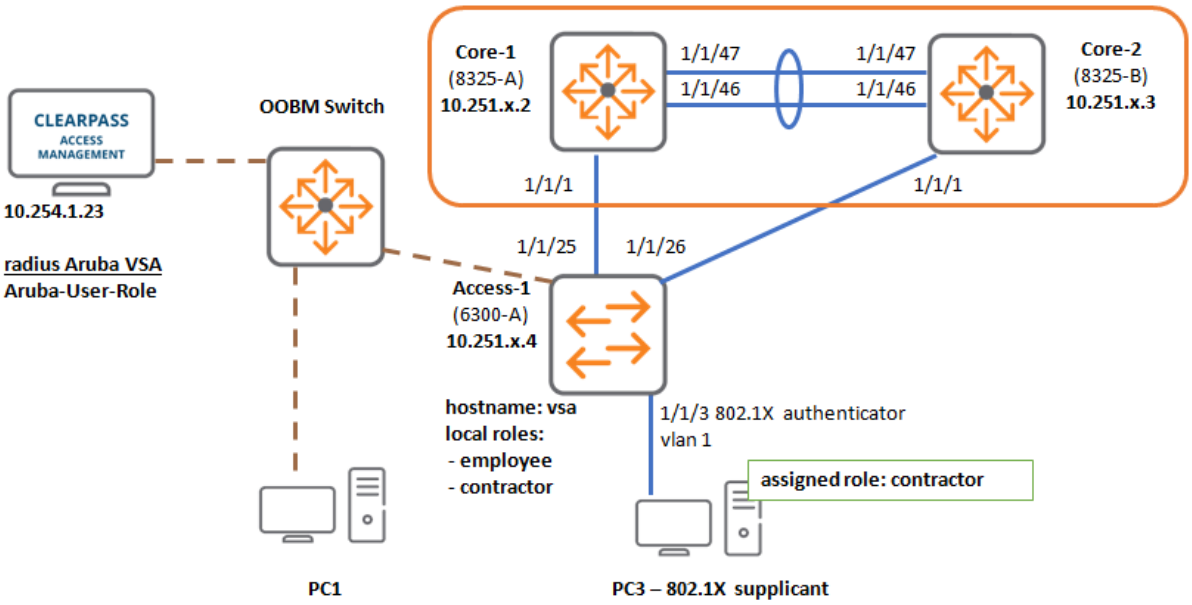

22. **Optional:** On PC3 (connected to Access1 1/1/3), change the 802.1X credentials to:
    - Username: **icx-contractor**
    - Password: **aruba123**

23. ClearPass has been configured to return the Aruba-User-Role **'contractor'** for this user. Check the Output tab's RADIUS Response for the Access Tracker record.

24. **Optional:** On the switch, review the authentication result.

```
ICX-Tx-Access1-avp(config)# show aaa authentication port-access dot1x
authenticator interface all client-status

Client 00:50:56:b1:7a:37, icx-contractor, 1/1/3
==========================================

  Authentication Details
  ----------------------
    Status                        : Authenticated
    Type                          : Pass-Through
    EAP-Method                    : PEAP
```

```
    Auth Failure reason           :
    Time Since Last State Change  : 6s
```

25. **Optional:** Review the role authorization.

```
ICX-Tx-Access1-avp(config)# show aaa authentication port-access interface all
client-status

Port Access Client Status Details

Client 00:50:56:b1:7a:37, icx-contractor
============================
  Session Details
  ---------------
    Port         : 1/1/3
    Session Time : 1340s

  Authentication Details
  ----------------------
    Status          : dot1x Authenticated
    Auth Precedence : dot1x - Authenticated, mac-auth - Not attempted

  Authorization Details
  ---------------------
    Role   : contractor
    Status : Applied
```

26. **Optional:** On the client PC, verify the contractor role cannot access the 'server' host. Ping should **not** be successful.

```
ping 10.x.1.2
```

Task 6: Unknown Role Assignment

Diagram

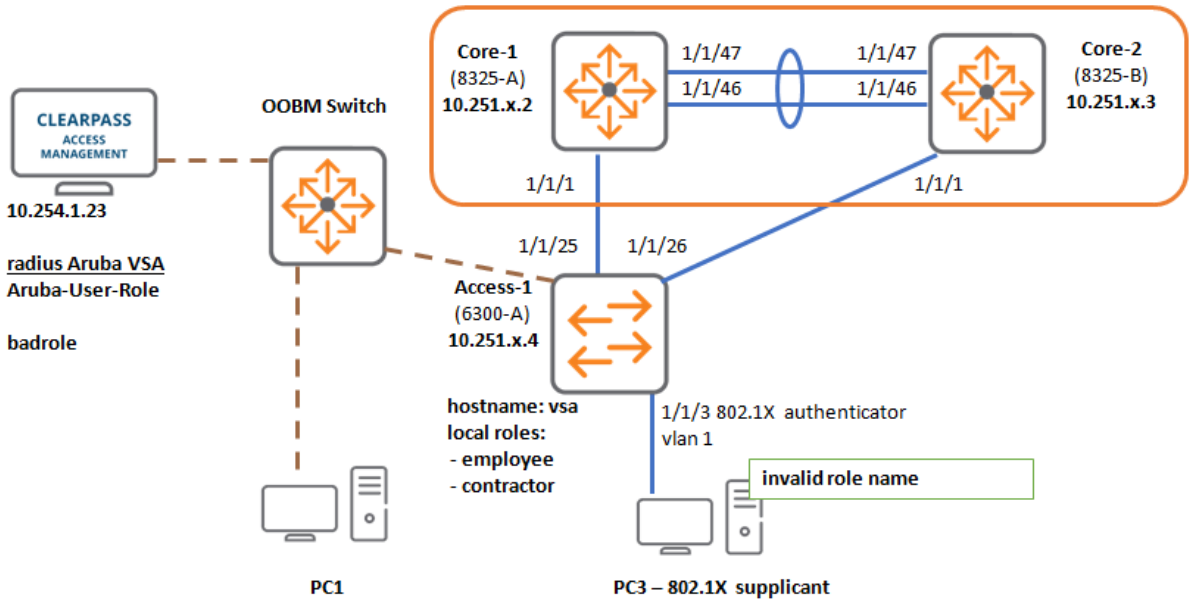

Objectives

Test result of configuration inconsistency between the RADIUS host and the switch configuration. From PC3, you'll log in with user **'icx-badrole'** and verify the switch status.

Steps

On PC3, Login with a user that has an incorrect role assignment.

1. On the PC3, change the 802.1X authentication credentials to:
   - Username: **icx-badrole**
   - Password: **aruba123**

Verify the result on the switch.

2. On the Access1 switch, verify the user has authenticated successfully.

```
ICX-Tx-Access1-avp(config)# show aaa authentication port-access dot1x
authenticator interface all client-status

Client 00:50:56:b1:7a:37, icx-badrole, 1/1/3
==========================================

  Authentication Details
```

```
  ----------------------
    Status                          : Authenticated
    Type                            : Pass-Through
    EAP-Method                      : PEAP
    Auth Failure reason             :
    Time Since Last State Change    : 39s
```

3. Review the role assignment.

```
ICX-Tx-Access1-avp(config)# show aaa authentication port-access interface all
client-status

Port Access Client Status Details

Client 00:50:56:b1:7a:37, icx-badrole
============================
  Session Details
  ---------------
    Port         : 1/1/3
    Session Time : 1887s

  Authentication Details
  ----------------------
    Status          : dot1x Authenticated
    Auth Precedence : dot1x - Authenticated, mac-auth - Not attempted

  Authorization Details
  ---------------------
    Role   :
    Status : Invalid
```

4. Disable port 1/1/3 so that it does not interfere with the next lab.

```
ICX-Tx-Access1-avp(config)# interface 1/1/3
ICX-Tx-Access1-avp(config-if)# shutdown
ICX-Tx-Access1-avp(config-if)# exit
```

You've completed Lab 10!

Lab 11: MAC-Based Authentication- MAC-Based Authentication and User Roles

Lab Diagram

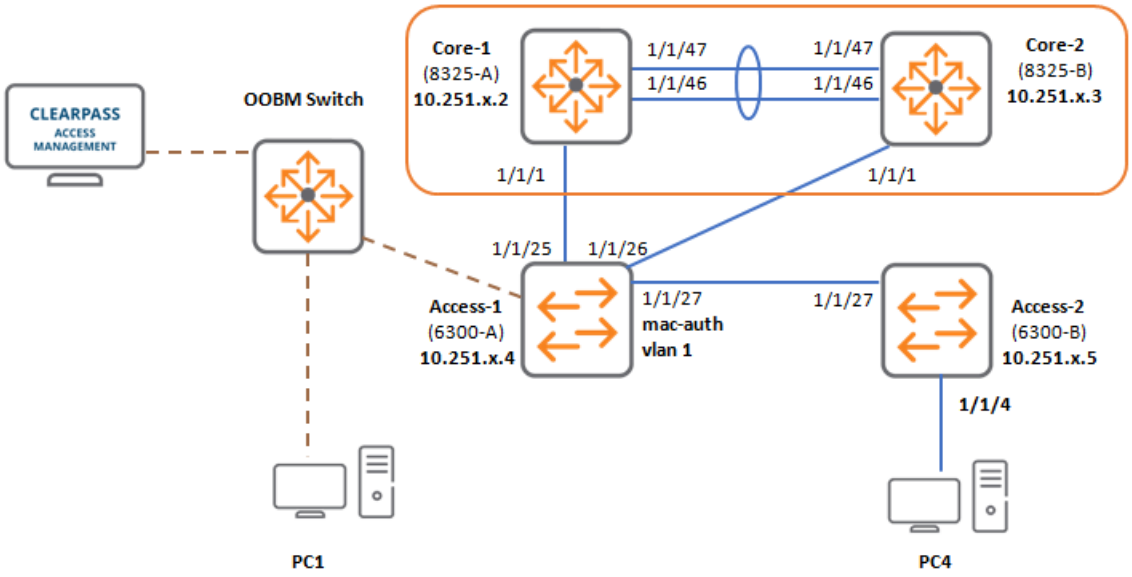

Overview

In this lab activity, the MAC authentication feature will be configured. While MAC authentication is not a secure authentication method, it will be required in most networks, since there will typically be some devices that cannot perform 802.1X authentication.

By using MAC authentication, some control can be applied to these devices.

Just like with the 802.1X authentication, it is possible to use user roles to combine authentication settings into a logical object group. The lab will also explore the interaction of 802.1X and mac authentication on the same switch port, and the difference between client-based and device-based authentication will be demonstrated.

The last part of the lab will show how LLDP device profiles can be used to push the correct configuration to a switch port when, for instance, an Aruba AP is connected. While this is technically not MAC authentication, the device profile feature is also based on information received by a peer device via LLDP.

Objectives

- Configure MAC authentication and review single and multiple users on a port

- Configure user roles with MAC authentication
- Review the interaction between 802.1X and MAC authentication on the same port
- Understand client-mode and device-mode authentication
- Understand the LLDP device profile feature

Task 1: MAC Authentication with a Single Device on a Port

Diagram

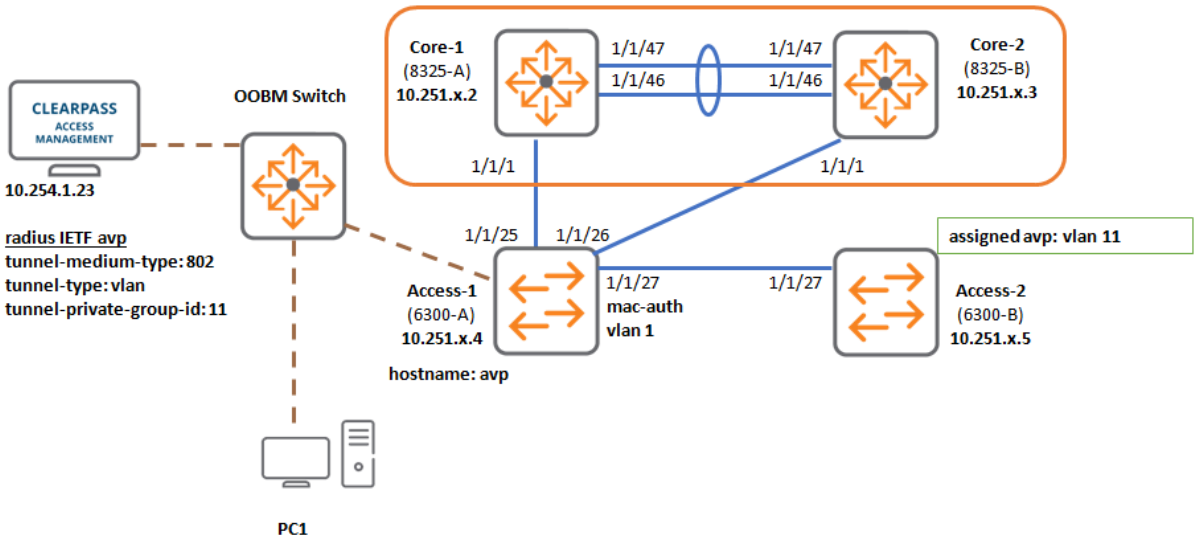

Objectives

Setup MAC authentication on the port connected to **Access2**. This will allow you to use Access2 as a 'connected device', and by using the PC that is connected to Access2, it will be possible to verify the operation with multiple clients on the same physical port.

Steps

Disconnect the Access2 switch from the VSX Core

Access2

1. Open a terminal session to Access2. This switch will act as an endpoint during the authentication labs.
2. On Access2, disable the uplink ports to the VSX Core1 and Core2.

```
ICX-Tx-Access2# configure terminal
ICX-Tx-Access2(config)# interface 1/1/25,1/1/26
ICX-Tx-Access2(config-if-<1/1/25,1/1/26>)# shutdown
ICX-Tx-Access2(config-if-<1/1/25,1/1/26>)# exit
ICX-Tx-Access2(config)#
```

3. On Access2, enable the port 1/1/27 that connects to Access1. This port is still at its default configuration state: it is an access port in VLAN 1.

```
ICX-Tx-Access2(config)# interface 1/1/27
ICX-Tx-Access2(config-if)# no shutdown
ICX-Tx-Access2(config-if)# exit
```

Enable MAC Authentication on Access1

Access1

4. Open a terminal session to Access1 and enter configuration mode.
5. Change the hostname to include 'avp'. Change “Tx” to “T” followed by your table number; for example, Table 1 would be “T1” and Table 12 would be T12.

```
ICX-Tx-Access1-vsa(config)# hostname ICX-Tx-Access1-avp
ICX-Tx-Access1-avp(config)#
```

> **NOTE:** The hostname is used as the NAS-identifier in the access-request to ClearPass. ClearPass has been pre-configured to look for the value '**avp**' in the NAS-identifier to return standard IETF **Attribute Value Pairs** with VLAN instructions that are needed for this lab.

6. Enable MAC authentication on the switch and ensure it is using the previously defined RADIUS server group.

```
ICX-Tx-Access1-avp(config)# aaa authentication port-access mac-auth radius
server-group cppm
ICX-Tx-Access1-avp(config)# aaa authentication port-access mac-auth enable
```

7. Assign port 1/1/27 to VLAN 1. In this lab, it is simply used as the base landing VLAN in case the RADIUS server does not assign a VLAN. There is no DHCP server or no other resources in this VLAN in this example.

```
ICX-Tx-Access1-avp(config)# interface 1/1/27
ICX-Tx-Access1-avp(config-if)# vlan access 1
```

8. Enter the MAC-auth context and enable MAC-auth on the port.

```
ICX-Tx-Access1-avp(config-if)# aaa authentication port-access mac-auth
ICX-Tx-Access1-avp(config-if-macauth)# enable
ICX-Tx-Access1-avp(config-if-macauth)# exit
ICX-Tx-Access1-avp(config-if)#
```

9. Enable the port.

```
ICX-Tx-Access1-avp(config-if)# no shutdown
ICX-Tx-Access1-avp(config-if)# exit
ICX-Tx-Access1-avp(config)#
```

10. Verify that the MAC address of the Access2 switch is now authenticated. The ClearPass system has been configured to assign the MAC OUI range of the Aruba Switch to VLAN 11.

NOTE: It make take a few moments for Access2 to generate some traffic, repeat below command until the authentication shows up. This should occur within about 30 seconds.

```
ICX-Tx-Access1-avp(config)# show aaa authentication port-access mac-auth
interface 1/1/27 client-status

Port Access Client Status Details

Client  88:3a:30:97:b6:00, 883a3097b600, 1/1/27
==========================================

  Authentication Details
  ----------------------
    Status                                  : Authenticated
    Auth-Method                             : chap
    Auth Failure reason                     :
    Time Since Last State Change            : 23s

  Authentication Statistics
  -------------------------
    Authentication                : 1
    Authentication Timeout        : 0
    Successful Authentication     : 1
    Failed Authentication         : 0
    Re-Authentication             : 0
    Successful Re-Authentication  : 0
    Failed Re-Authentication      : 0
    Re-Auths When Authenticated   : 0
    Cached Re-Authentication      : 0
```

11. Confirm the MAC-address on the port is now dynamically learned on VLAN 11 by reviewing the mac-address-table.

```
ICX-Tx-Access1-avp(config)# show mac-address-table
MAC age-time            : 300 seconds
Number of MAC addresses : 9

MAC Address          VLAN     Type                        Port
-------------------------------------------------------------
20:4c:03:5f:98:02    1        dynamic                     lag256
88:3a:30:97:b6:00    11       port-access-security        1/1/27
02:02:00:00:12:00    11       dynamic                     lag256
90:20:c2:bc:17:00    11       dynamic                     lag256
90:20:c2:bc:97:00    11       dynamic                     lag256
<...output omitted...>
```

12. For this initial MAC authentication lab, ClearPass returns standard RADIUS IETF VLAN assignment attributes. These are dynamically assigned into a role on the switch. Review the authorization of this MAC address:

```
ICX-Tx-Access1-avp(config)# show aaa authentication port-access interface 1/1/27
client-status

Port Access Client Status Details
Client 88:3a:30:97:b6:00, 883a3097b600
============================
  Session Details
  ---------------
    Port         : 1/1/27
    Session Time : 7s

  Authentication Details
  ----------------------
    Status          : mac-auth Authenticated
    Auth Precedence : dot1x - Not attempted, mac-auth - Authenticated

  Authorization Details
  ---------------------
    Role   : RADIUS_1881452276
    Status : Applied
```

13. Review the role details to discover the assigned RADIUS attributes.

```
ICX-Tx-Access1-avp(config)# show port-access role radius

Role Information:

Name  : RADIUS_1881452276
Type  : radius
------------------------------------------------
    Reauthentication Period             :
    Authentication Mode                 :
    Session Timeout                     :
    Client Inactivity Timeout           :
    Description                         :
    Gateway Zone                        :
    UBT Gateway Role                    :
    Access VLAN                         : 11
    Native VLAN                         :
    Allowed Trunk VLANs                 :
    MTU                                 :
    QOS Trust Mode                      :
    PoE Priority                        :
    Captive Portal Profile              :
    Policy                              :
```

Optional Task 2: Verify Access with Two Devices Connected on Same Port

This task is **optional** and can be done if time permits. Check with your instructor.

Diagram

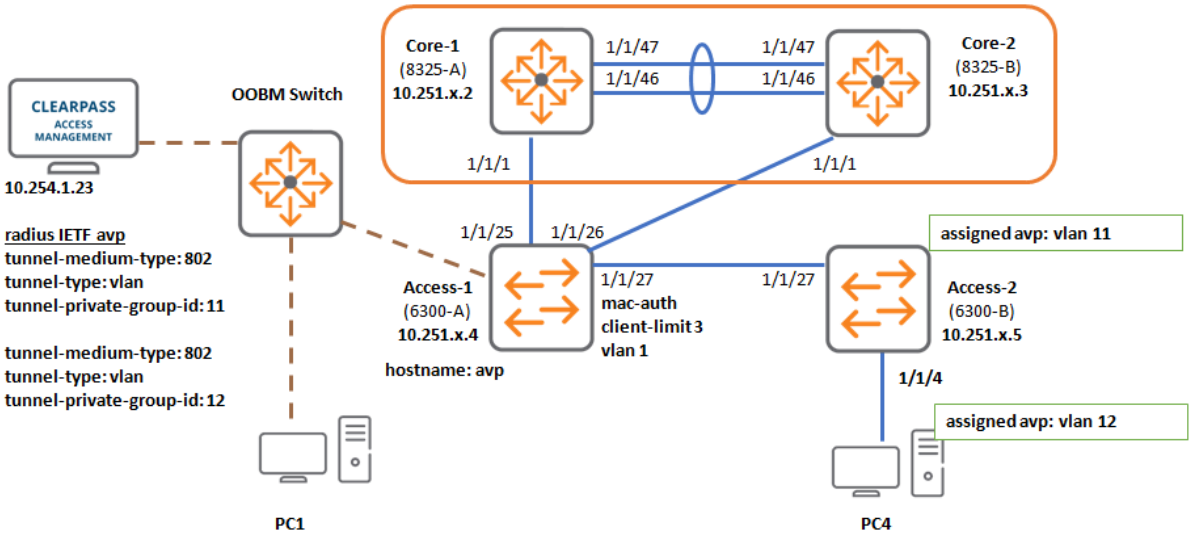

Objectives

In this task, the PC that is connected to Access2 will be authenticated as well. This will result in two untagged devices that are authenticated on the same physical port, but they are still assigned to different VLANs. This is also known as MAC-based VLAN (MBV), since the switch assigns the traffic to the correct VLAN based on the authenticated client MAC address.

Initially two clients (PC4 and Access2) will be connected to the same port 1/1/27 on Access1. Only one of these two clients will be able to come online, since there is a default client limit of *one*. Next, this client limit will be increased so both clients can come online at the same time.

Steps

Access2

1. On the Access2 switch, enable port 1/1/4 and assign it to VLAN 1. This is the port that connects to PC4.

```
ICX-Tx-Access2(config)# int 1/1/4
ICX-Tx-Access2(config-if)# vlan access 1
ICX-Tx-Access2(config-if)# no shutdown
ICX-Tx-Access2(config-if)# exit
```

2. Open a connection to **PC4** connected to Access2.

3. Open the Network Connections, look for the 'Lab NIC'.

> **IMPORTANT:** Make sure you select the **'Lab NIC'** and not the **'DO NOT TOUCH'** interface.

4. Bounce the **'Lab NIC'** interface (disable/enable). This will trigger the DHCP client to send out a request.

Access2

5. Verify on Access2 that a MAC address is learned on the port 1/1/4. The MAC will probably start with **00:50:56**, that is the range used by VMware to generate MAC addresses for the VMs.

```
ICX-Tx-Access2(config)# show mac-address-table
MAC age-time            : 300 seconds
Number of MAC addresses : 5

MAC Address          VLAN     Type                      Port
-------------------------------------------------------------
90:20:c2:bc:97:00    1        dynamic                   1/1/27
90:20:c2:bc:17:00    1        dynamic                   1/1/27
88:3a:30:98:30:c0    1        dynamic                   1/1/27
00:50:56:b1:fc:9b    1        dynamic                   1/1/4
02:02:00:00:12:00    1        dynamic                   1/1/27
```

Access1

6. Review the active clients on the port 1/1/27.

> **NOTE:** In the output below, it may be either the MAC address of PC4 or the Access2 that is authenticated on the port.

```
ICX-Tx-Access1-avp(config)# show aaa authentication port-access interface 1/1/27
client-status

Port Access Client Status Details

Client 88:3a:30:97:b6:00, 883a3097b600
==============================
  Session Details
  ---------------
    Port         : 1/1/27
    Session Time : 50s

  Authentication Details
  ----------------------
    Status          : mac-auth Authenticated
    Auth Precedence : dot1x - Not attempted, mac-auth - Authenticated

  Authorization Details
```

```
  ----------------------
    Role   : RADIUS_242871489
    Status : Applied
```

Q: How many authenticated clients are there on the interface?

_____

A: Only one client is authenticated, even though two clients (the Access2 switch and the client PC4) are connected to the port.

7. Review the mac-address table on Access1 using the **'detail'** option word. Look for the **'Denied'** column. This reveals that the mac-address of the client has been seen on the port but is not allowed to become active.

   NOTE: In the output below, it can be either the PC4 or the Access2 MAC address that is 'denied' on the port 1/1/27. This depends on which MAC address was authenticated on the port, the other MAC address will be denied.

```
ICX-Tx-Access1-avp(config)# show mac-address-table detail
MAC age-time             : 300 seconds
Number of MAC addresses : 11

MAC Address          VLAN Type                   Port    Age  Denied never_ageout
-------------------------------------------------------------------------
88:3a:30:97:b6:00    1    port-access-security   1/1/27  300  true   false
00:50:56:b1:fc:9b    1    port-access-security   1/1/27  300  true   false
...
88:3a:30:97:b6:00    12   port-access-security   1/1/27  300  false  false
...
```

This is because the switch applies a default client limit of one client on each switch port.

Now increase the client limit, to ensure multiple client can be authenticated on the same switch port.

8. On Access1, enter the interface 1/1/27 context and increase the client limit to three. Reset the interface to trigger the Access2 switch to send some packets.

```
ICX-Tx-Access1-avp(config)# interface 1/1/27
ICX-Tx-Access1-avp(config-if)# aaa authentication port-access client-limit 3
ICX-Tx-Access1-avp(config-if)# shutdown
ICX-Tx-Access1-avp(config-if)# no shutdown
```

9. On PC4, disable and enable the 'Lab NIC' again.

10. On Access1, there should now be two clients authenticated on the same port.

```
ICX-Tx-Access1-avp(config-if)# show aaa authentication port-access interface
1/1/27 client-status

Port Access Client Status Details

Client 88:3a:30:97:b6:00, 883a3097b600
============================
  Session Details
  ---------------
    Port         : 1/1/27
    Session Time : 255s

  Authentication Details
  ----------------------
    Status          : mac-auth Authenticated
    Auth Precedence : dot1x - Not attempted, mac-auth - Authenticated

  Authorization Details
  ---------------------
    Role   : RADIUS_242871489
    Status : Applied


Client 00:50:56:b1:fc:9b, 005056b1fc9b
============================
  Session Details
  ---------------
    Port         : 1/1/27
    Session Time : 19s

  Authentication Details
  ----------------------
    Status          : mac-auth Authenticated
    Auth Precedence : dot1x - Not attempted, mac-auth - Authenticated

  Authorization Details
  ---------------------
    Role   : RADIUS_242871489
    Status : Applied
```

NOTE: In case Access2 is not listed as a client, bounce the VLAN 1 Layer3 interface to trigger some DHCP traffic

```
ICX-Tx-Access2(config)# interface vlan 1
ICX-Tx-Access2(config-if-vlan)# shutdown
ICX-Tx-Access2(config-if-vlan)# no shutdown
ICX-Tx-Access2(config-if-vlan)# exit
```

11. Review the RADIUS attributes that were pushed to for the client PC mac auth by checking the radius roles.

```
ICX-Tx-Access1-avp(config-if)# show port-access role radius

Role Information:

Name   : RADIUS_1881452276
Type   : radius
------------------------------------------
    Reauthentication Period             :
    Authentication Mode                 :
    Session Timeout                     :
    Client Inactivity Timeout           :
    Description                         :
    Gateway Zone                        :
    UBT Gateway Role                    :
    Access VLAN                         : 11
    Native VLAN                         :
    Allowed Trunk VLANs                 :
    MTU                                 :
    QOS Trust Mode                      :
    PoE Priority                        :
    Captive Portal Profile              :
    Policy                              :

Name   : RADIUS_242871489
Type   : radius
------------------------------------------
    Reauthentication Period             :
    Authentication Mode                 :
    Session Timeout                     :
    Client Inactivity Timeout           :
    Description                         :
    Gateway Zone                        :
    UBT Gateway Role                    :
    Access VLAN                         : 12
    Native VLAN                         :
    Allowed Trunk VLANs                 :
    MTU                                 :
    QOS Trust Mode                      :
    PoE Priority                        :
    Captive Portal Profile              :
    Policy                              :
```

12. Review the updated MAC address table on port 1/1/27. Notice how both MAC addresses of the Access2 and the PC4 are now active, but each has been assigned to its own VLAN, based on the instructions of the RADIUS server.

```
ICX-Tx-Access1-avp(config-if)# show mac-address-table port 1/1/27
MAC age-time            : 300 seconds
Number of MAC addresses : 2
```

```
MAC Address          VLAN      Type                          Port
------------------------------------------------------------------
88:3a:30:97:b6:00    11        port-access-security          1/1/27
00:50:56:b1:fc:9b    12        port-access-security          1/1/27
```

This demonstrates that multiple untagged devices can be concurrently active on the same physical port, each with their own network access security settings applied to them.

Optional steps: Review the Authentication and Attributes on the ClearPass System

13. **Optional step**: On the ClearPass access tracker, use Access Tracker to verify that two authentications were processed.

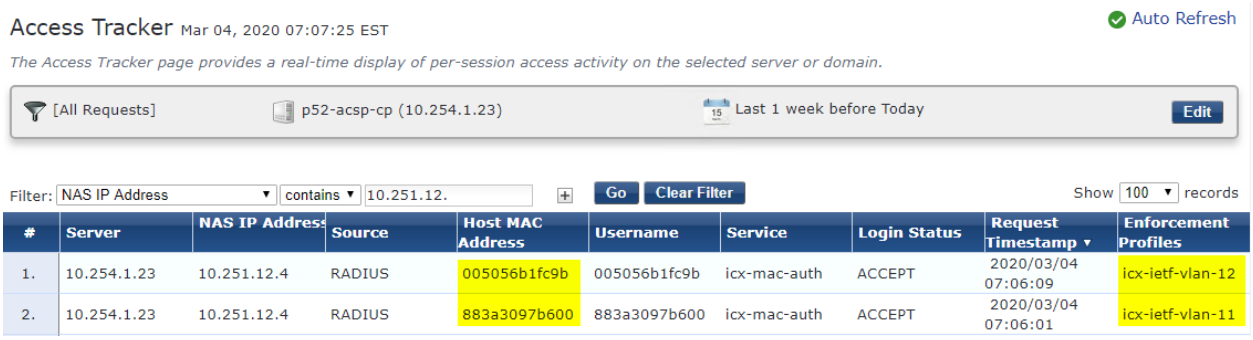

14. **Optional step**: Open the MAC-auth entry for MAC **005056** (the PC4 VM), click on the **'Input'** to review the attributes that are used for MAC authentication.

Example Input request for mac auth:

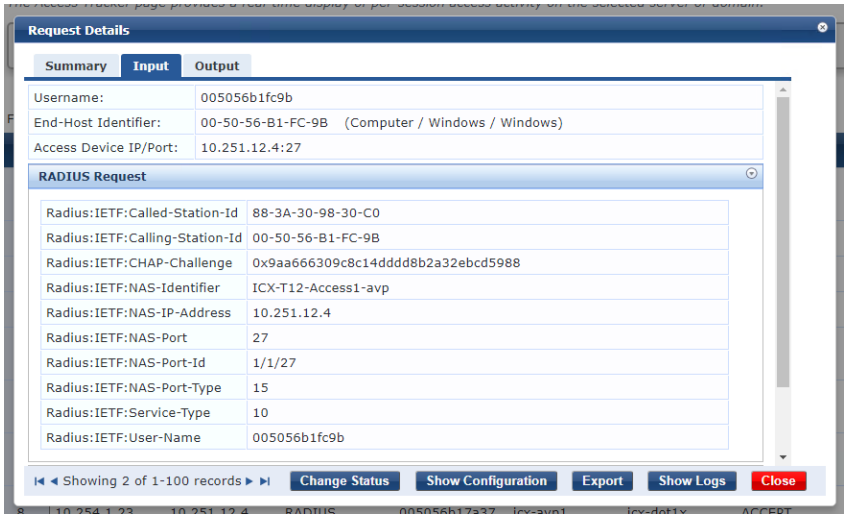

15. **Optional step**: Check the **'Output'** tab, click on 'RADIUS response' to see the VLAN that was returned to the switch. **Close** the 'Request details'.

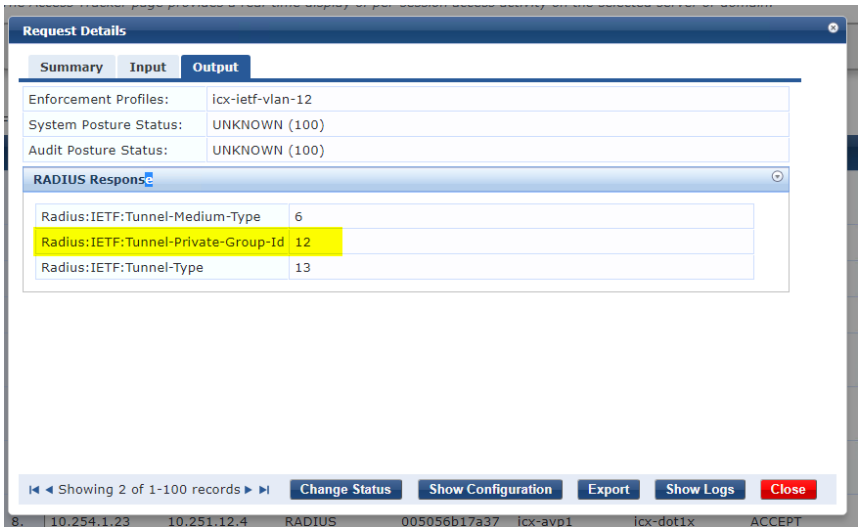

16. **Optional step**: Open the second MAC authentication entry (for the Access2 switch) and review the **'Output'** settings.

Example for the Switch (assigned to VLAN 11):

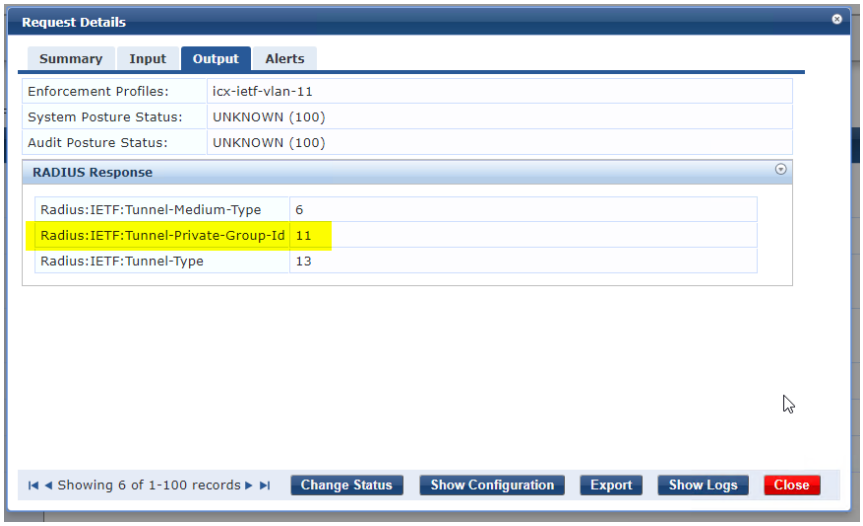

End of the optional ClearPass settings review steps

Task 3: Aruba User Role Based Access

Diagram

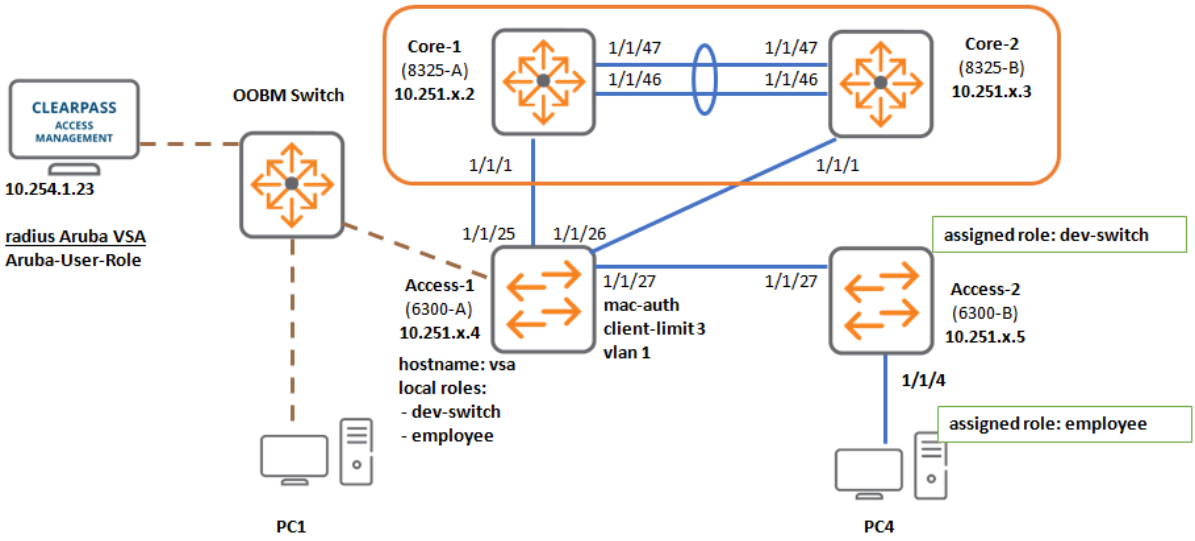

Objectives

Similar to the 802.1X authentication, MAC authentication also supports the use of user roles that are defined on the switch. User roles on the switch have the advantage of grouping all the relevant access settings (VLAN, ACL, QOS, POE priority) into a single logical object. This simplifies the RADIUS server configuration, since the RADIUS server only needs to assign the role to the switch, while the switch knows the configuration of that role.

In this lab, the Access1 switch will be assigned the role **'dev-switch'**, while the PC connected to the Access2 switch will be assigned the role **'employee'**.

In the previous Lab (Configuring 802.1X) the 'employee' role was already defined, so these steps are not included here. This task will define the new user role for the Access2 switch and show that the role can also be used to control both tagged and untagged access for the user role.

Steps

Access1

Define new 'dev-switch' role on Access1.

1. Open a terminal connection to Access1, enter the configuration mode.

Port Client Limit Check

2. If you did not complete the previous optional task, increase the client-limit on the port 1/1/27, otherwise move to the next step.

```
ICX-Tx-Access1-avp(config)# interface 1/1/27
```

```
ICX-Tx-Access1-avp(config-if)# aaa authentication port-access client-limit 3
ICX-Tx-Access1-avp(config-if)# exit
```

NOTE: You can verify the current interface configuration using **'show running interface 1/1/27**' and check if the client-limit was applied.

3. Define a new port-access user-role named **'dev-switch**'. Provide the role a VLAN trunk with native VLAN 11, and allowed vlans 11,12,13.

```
ICX-Tx-Access1-avp(config)# port-access role dev-switch
ICX-Tx-Access1-avp(config-pa-role)# vlan trunk native 11
ICX-Tx-Access1-avp(config-pa-role)# vlan trunk allowed 11,12,13
ICX-Tx-Access1-avp(config-pa-role)# exit
ICX-Tx-Access1-avp(config)#
```

Change the hostname to get the Aruba VSA Aruba-User-Role ClearPass rule set.

4. On Access1, change the hostname, replace the **'avp'** the hostname with **'vsa'**. Change “Tx” to “T” followed by your table number; for example, Table 1 would be “T1” and Table 12 would be T12.

```
ICX-Tx-Access1-avp(config)# hostname ICX-Tx-Access1-vsa
ICX-Tx-Access1-vsa(config)#
```

NOTE: The hostname is used as the NAS-identifier. ClearPass has been pre-configured to look for the value '**vsa**' in the NAS-identifier to return the **Aruba-User-Role** VSA attributes that are needed for this lab.

Test Role Assignment

5. On the Access1 switch, disable and enable the port 1/1/27.

```
ICX-Tx-Access1-vsa(config)# interface 1/1/27
ICX-Tx-Access1-vsa(config-if)# shutdown
ICX-Tx-Access1-vsa(config-if)# no shutdown
ICX-Tx-Access1-vsa(config-if)# exit
```

You may need to generate some traffic to trigger the MAC-auth on the 'clients': That is, the Access2 switch and the PC connected to it.

Access2

6. On the Access2, disable and enable the VLAN 1 IP interface, this will send out a DHCP request.

```
ICX-Tx-Access2(config)# interface vlan 1
ICX-Tx-Access2(config-if-vlan)# shutdown
ICX-Tx-Access2(config-if-vlan)# no shutdown
ICX-Tx-Access2(config-if-vlan)# exit
```

7. On the PC4 connected to Access2, disable and enable the 'Lab NIC'.

> **NOTE:** This is important, since the PC4 may be assigned to a different VLAN, so it needs to renew its IP address.

Verify the Results

Access1

8. On the Access1 switch, review the authenticated clients and the role that was assigned to the clients.

```
ICX-Tx-Access1-vsa(config)# show aaa authentication port-access interface 1/1/27
client-status

Port Access Client Status Details

Client 00:50:56:b1:fc:9b, 005056b1fc9b
============================
  Session Details
  ---------------
    Port         : 1/1/27
    Session Time : 4s

  Authentication Details
  ----------------------
    Status          : mac-auth Authenticated
    Auth Precedence : dot1x - Not attempted, mac-auth - Authenticated

  Authorization Details
  ---------------------
    Role   : employee
    Status : Applied


Client 88:3a:30:97:b6:00, 883a3097b600
============================
  Session Details
  ---------------
    Port         : 1/1/27
    Session Time : 179s

  Authentication Details
  ----------------------
    Status          : mac-auth Authenticated
    Auth Precedence : dot1x - Not attempted, mac-auth - Authenticated

  Authorization Details
  ---------------------
    Role   : dev-switch
    Status : Applied
```

9. Review the VLAN port status of the port 1/1/27.

```
ICX-Tx-Access1-vsa(config)# show vlan port 1/1/27

--------------------------------------------------------------------------------
VLAN  Name                                Mode            Mapping
--------------------------------------------------------------------------------
1     DEFAULT_VLAN_1                      native-untagged port
11    VLAN11                              trunk           port-access,mbv
12    VLAN12                              trunk           port-access
13    VLAN13                              trunk           port-access
ICX-Tx-Access1-vsa(config)#
```

Q: What is the listed method for the tagged VLAN, e.g. VLAN 12, assignment mapping?

_____

A: The VLAN 12 is assigned by the 'port-access' module, while a manual VLAN trunk would be listed as mapped by the 'port' configuration, as can be seen for VLAN 1.

10. **Optional step:** Using the PC1 Management PC, access ClearPass, open Access Tracker, and review the output of both authentication events. This demonstrates that the Aruba VSA User-Role was used for the assignment.

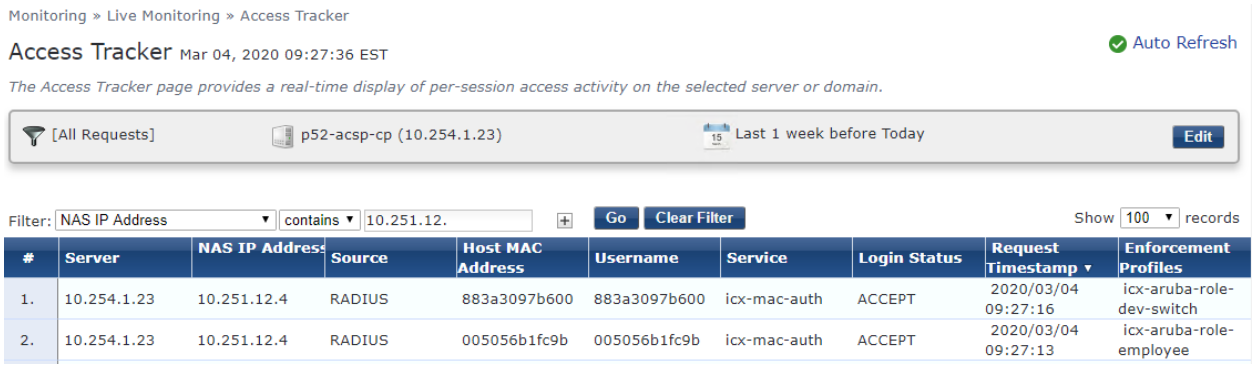

11. **Optional step:** Review each entry, check that the VSA Aruba-User-Role was now assigned to Access1 instead of the IETF VLAN assignment.

For the PC4:

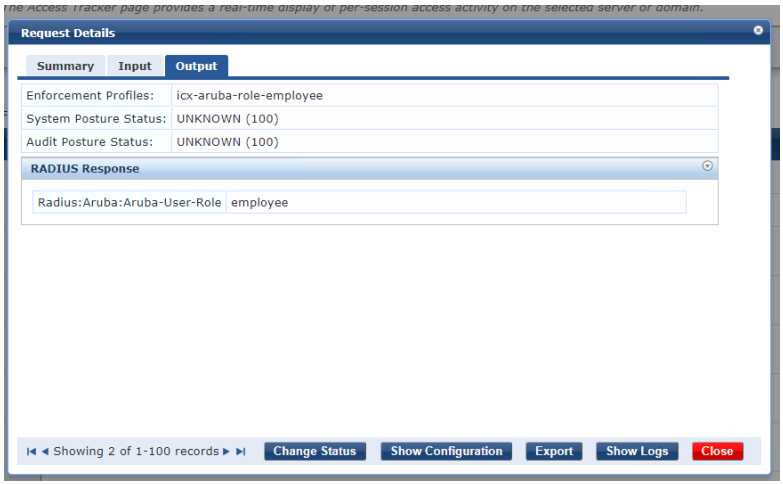

For the Switch Access2:

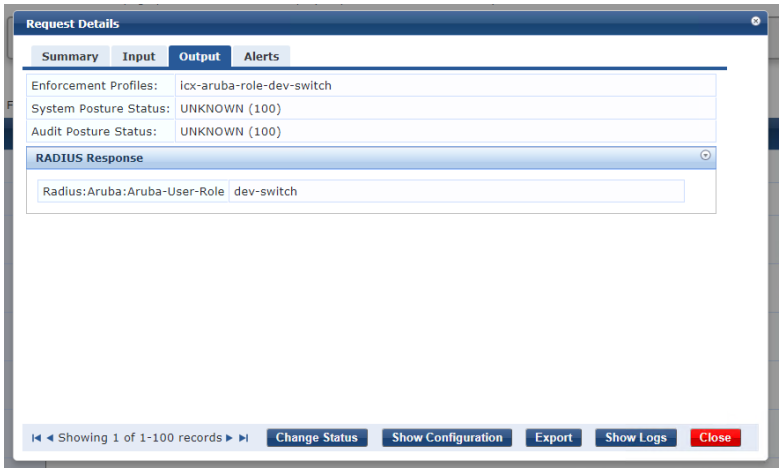

Optional Task 4: Client-Mode versus Device-Mode Port Authentication

This task is **optional** and can be done if time permits. Check with your instructor.

Diagram

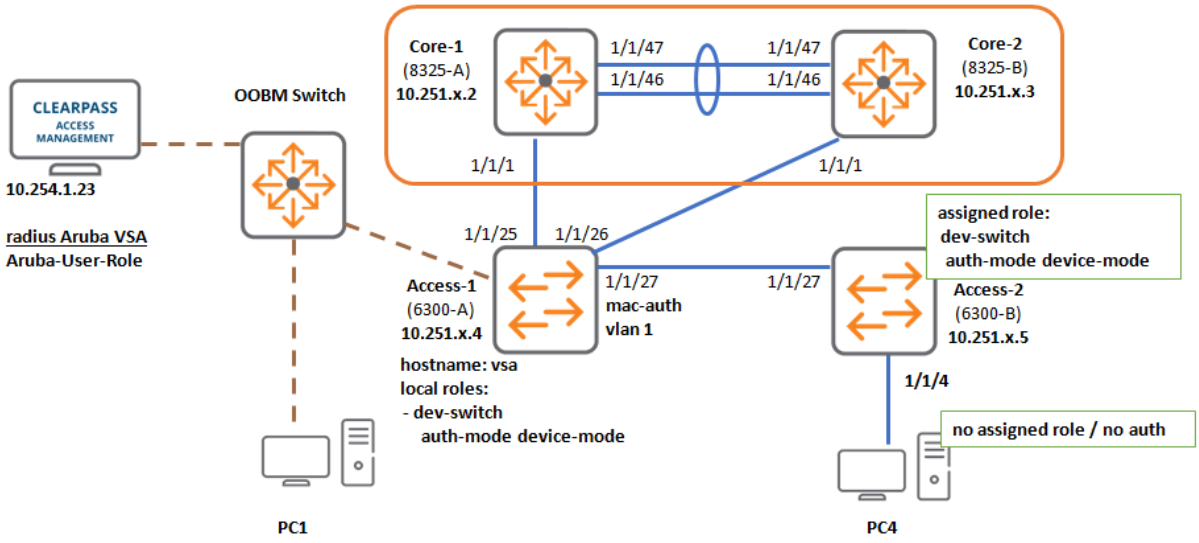

Objectives

In the previous task, it was demonstrated that multiple devices can be connected to the same switch port, and each device can be assigned its own network access policy by using Aruba user roles. However, sometimes it may not be desired to authenticate each device on the port, for example, if the connected device itself is responsible to perform authentication.

This would apply, for example, to Aruba Instant APs that are connected to a switch port. With an Instant AP, the wireless traffic will be forwarded by the AP as either tagged or untagged traffic to the switch port. This is also referred to as local breakout at the AP or bridging at the AP level. The result is that the switch port will actually see all the MAC addresses that are connected to the IAP device.

The IAP itself is responsible to handle the authentication, so it would perform 802.1X authentication with the wireless clients. But then the traffic is forwarded as regular traffic on the switch port, so the switch would also attempt to perform authentication of this client. Since the 802.1X traffic of the client is terminated at the IAP, the switch would attempt to perform MAC authentication for the client MAC address.

This is unnecessary and confusing, since ClearPass would see the same MAC address as 802.1X authenticated on the IAP, and MAC-authenticated on the switch port.

For this scenario, the switch can be set to 'port-based' authentication; i.e., device mode. When a specific device (such as an IAP or an access switch in a meeting room) is connected to the switch port, the switch will open the port for **all** traffic, so it will not perform MAC authentication for the other MAC addresses seen off of the port anymore.

In this task, you will change the user-role of the Access2 switch 'dev-switch' to become a role that operates as **'Device-mode**'. When the switch completes its authentication, any device connected to the switch will automatically be allowed access to the network as well. In this lab, the Access2 is not actually performing any authentication itself for PC4, but it simply demonstrates the concept of the **'Device-mode**' authentication.

Steps

Access2

1. Open Access2, disable the port to the PC. This is done to ensure that the Access2 switch's MAC address will be guaranteed to be the first seen MAC address on the port.

```
ICX-Tx-Access2(config)# interface 1/1/4
ICX-Tx-Access2(config-if)# shutdown
```

Access1: Change the Role

2. On Access1, modify the user-role for 'dev-switch'.

```
ICX-Tx-Access1-vsa(config)# port-access role dev-switch
ICX-Tx-Access1-vsa(config-pa-role)# auth-mode device-mode
ICX-Tx-Access1-vsa(config-pa-role)# exit
```

3. Reset the port 1/1/27.

```
ICX-Tx-Access1-vsa(config)# interface 1/1/27
ICX-Tx-Access1-vsa(config-if)# shutdown
ICX-Tx-Access1-vsa(config-if)# no shutdown
ICX-Tx-Access1-vsa(config-if)# exit
```

4. Verify the connected switch has authenticated.

```
ICX-Tx-Access1-vsa(config)# show aaa authentication port-access interface 1/1/27
client-status

Port Access Client Status Details

Client 88:3a:30:97:b6:00, 883a3097b600
============================
  Session Details
  ---------------
    Port         : 1/1/27
    Session Time : 131s

  Authentication Details
  ----------------------
    Status          : mac-auth Authenticated
```

```
    Auth Precedence : dot1x - Not attempted, mac-auth - Authenticated

  Authorization Details
  ---------------------
    Role   : dev-switch
    Status : Applied
```

Access2

On Access2, enable the PC port to verify it passes network access without authentication.

5. On Access2, enable the port 1/1/4 of the PC again.

```
ICX-Tx-Access2(config-if)# no shutdown
```

6. On the PC4, disable and enable the 'Lab Nic' interface.

Access1

7. On the Access1, verify the MAC address table shows 2 MAC addresses on the interface 1/1/27, and the 'Denied' status is **'false'**. This means these MAC addresses have access to the network.

```
ICX-Tx-Access1-vsa(config)# show mac-address-table detail
MAC age-time             : 300 seconds
Number of MAC addresses : 10

MAC Address          VLAN Type                  Port    Age  Denied never_ageout
-------------------------------------------------------------------------------
20:4c:03:5f:98:02    1    dynamic               lag256  300  false  false
88:3a:30:97:b6:00    11   port-access-security  1/1/27  300  false  false
00:50:56:b1:fc:9b    11   port-access-security  1/1/27  300  false  false
02:02:00:00:12:00    11   dynamic               lag256  300  false  false
90:20:c2:bc:17:00    11   dynamic               lag256  300  false  false

<...output omitted...>
```

8. Navigate to ClearPass's Access Tracker from PC1.

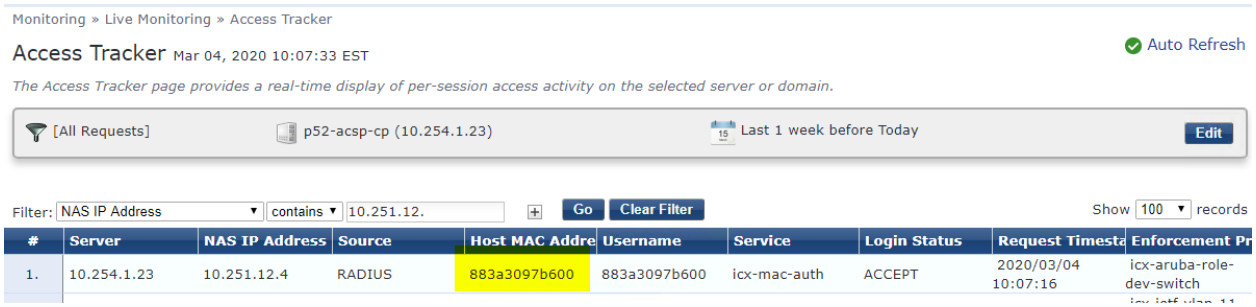

Q: How many authentication events are there for the last test?

A: Only one authentication was done by the switch. The role that was assigned contains the **'auth-mode: device-mode**': device' option, so no other authentications are needed on this port.

This concludes the device-mode feature lab demonstration.

Revert Configuration

Access1

9. On Access1, change the 'dev-switch' role back to client based and disable the port.

```
ICX-Tx-Access1-vsa(config)# port-access role dev-switch
ICX-Tx-Access1-vsa(config-pa-role)# auth-mode client-mode
ICX-Tx-Access1-vsa(config-pa-role)# exit
```

```
ICX-Tx-Access1-vsa(config)# interface 1/1/27
ICX-Tx-Access1-vsa(config-if)# shutdown
ICX-Tx-Access1-vsa(config-if)# exit
ICX-Tx-Access1-vsa(config)#
```

Access2

10. On Access2, disable the port 1/1/4 to PC4.

```
ICX-Tx-Access2(config-if)# shutdown
ICX-Tx-Access2(config-if)# exit
```

Task 5: Authentication Priority Order with Combined MAC-Auth and 802.1X

Diagram

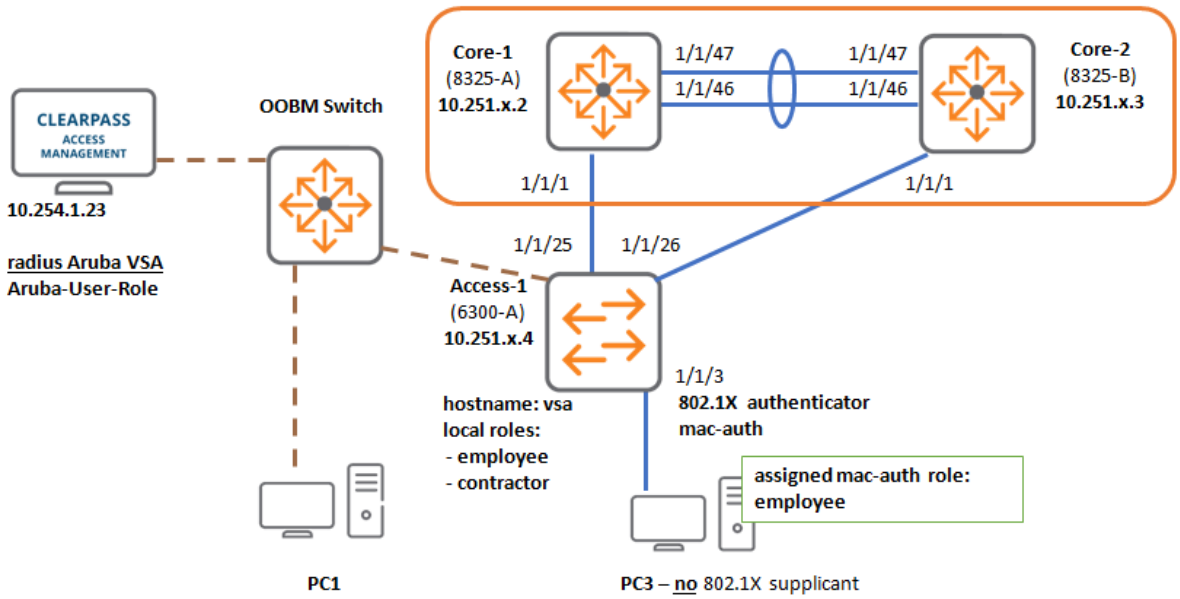

Objectives

This task will show how a port that combines 802.1X and MAC authentication will handle the combined authentication of a device.

For each MAC address that connects to the switch, only one authentication method can be active to control the network access. When a MAC address comes online, the switch will either apply the role that was received for the MAC authentication *or* the role that was received for the 802.1X authentication, but never both at the same time.

Since 802.1X is a more secure authentication method than MAC-Auth, the switch will, by default, give the 802.1X role precedence over the MAC-auth role.

In a real-world deployment, the administrator should realize that a PC with an 802.1X configuration will only start the 802.1X process when it has booted. During the boot process, the PC may already send some frames that may trigger MAC-authentication on the switch port. A few seconds later, the 802.1X process would also succeed and the 802.1X role would overrule the MAC-auth role.

In this task, this behavior will be observed by enabling MAC-auth on port 1/1/3 of Access1. The PC that performs 802.1X will first be authenticated using MAC-auth. Then the 802.1X supplicant will be enabled, so the 802.1X role will overrule the MAC-auth role.

Steps

Disable the 802.1X supplicant on PC3 (connected to port 1/1/3 on Access1).

1. Open a session to the PC3.
2. Get the properties for the '**Lab NIC' > Authentication**' and **uncheck** the 'Enable IEEE 802.1X authentication' checkbox. Confirm the message that saved credentials will be deleted, if prompted.

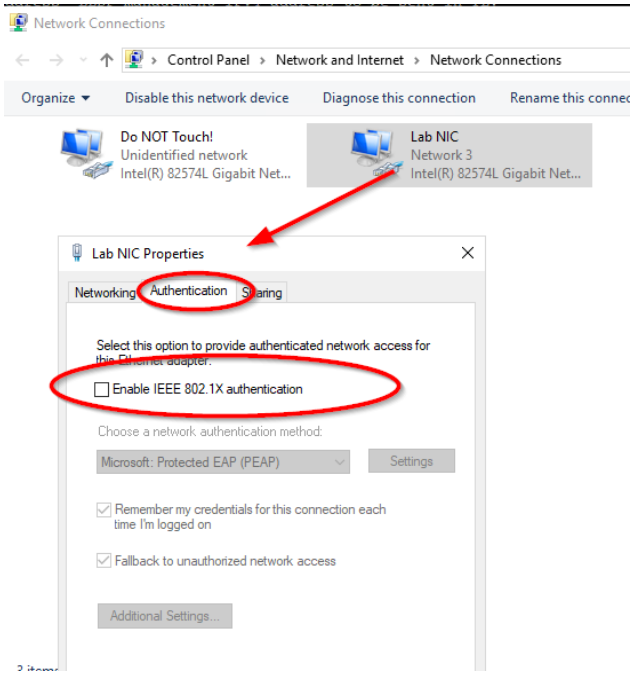

Enable MAC-auth on the Access1

Access1

3. On Access1, configure port 1/1/3 for MAC-auth.

```
ICX-Tx-Access1-vsa(config)# interface 1/1/3
ICX-Tx-Access1-vsa(config-if)# aaa authentication port-access mac-auth enable
```

4. Reset the port.

```
ICX-Tx-Access1-vsa(config-if)# shutdown
ICX-Tx-Access1-vsa(config-if)# no shutdown
ICX-Tx-Access1-vsa(config-if)# exit
```

5. On the PC3, disable and enable the 'Lab NIC'.
6. On the Access1, follow the client authentication status.

```
ICX-Tx-Access1-vsa(config)# show aaa authentication port-access interface 1/1/3
client-status

Port Access Client Status Details
```

```
Client 00:50:56:b1:7a:37
============================
  Session Details
  ---------------
    Port         : 1/1/3
    Session Time : 15s

  Authentication Details
  ----------------------
    Status          : Authenticating
    Auth Precedence : dot1x - Authenticating, mac-auth - Not attempted

  Authorization Details
  ----------------------
    Role   :
    Status :
```

NOTE: With the default timers, it may take several minutes before the 802.1X times out and performs the MAC-auth. During testing, about 3 minutes were observed.

7. Eventually, the switch will consider that 802.1X has failed and it will perform MAC-auth.

NOTE: Below output is for reference only, you do not need to wait for this timeout in the lab. You can move to the next step.

```
ICX-Tx-Access1-vsa(config)# show aaa authentication port-access interface 1/1/3
client-status

Port Access Client Status Details

Client 00:50:56:b1:7a:37, 005056b17a37
============================
  Session Details
  ---------------
    Port         : 1/1/3
    Session Time : 208s

  Authentication Details
  ----------------------
    Status          : mac-auth Authenticated
    Auth Precedence : dot1x - Unauthenticated, mac-auth - Authenticated

  Authorization Details
  ----------------------
    Role   : employee
    Status : Applied
```

Adjust the Timers to Achieve Faster Mac-authentication

8. Configure updated timers for the 802.1X authenticator on the port 1/1/3.

```
ICX-Tx-Access1-vsa(config)# interface 1/1/3
ICX-Tx-Access1-vsa(config-if)# aaa authentication port-access dot1x authenticator
ICX-Tx-Access1-vsa(config-if-dot1x-auth)# eapol-timeout 30
ICX-Tx-Access1-vsa(config-if-dot1x-auth)# max-eapol-requests 1
ICX-Tx-Access1-vsa(config-if-dot1x-auth)# max-retries 1
ICX-Tx-Access1-vsa(config-if-dot1x-auth)# exit
ICX-Tx-Access1-vsa(config-if)#
```

NOTE: These are the validated timers when the labs were developed, adjust these timers based on the requirements of the deployment.

9. Review the resulting port configuration.

```
ICX-Tx-Access1-vsa(config-if)# show running-config current-context
interface 1/1/3
    no shutdown
    no routing
    vlan access 1
    aaa authentication port-access dot1x authenticator
        eapol-timeout 30
        max-eapol-requests 1
        max-retries 1
        enable
    aaa authentication port-access mac-auth
        enable
    exit
```

10. Disable and enable the port.

```
ICX-Tx-Access1-vsa(config-if)# shutdown
ICX-Tx-Access1-vsa(config-if)# no shutdown
```

11. Verify the authentication status. It may help to disable/enable the 'Lab Nic' on the client PC to generate some client traffic.
12. On Access1, repeat the ' show aaa authentication port-access interface 1/1/3 client-status' command to follow the state transitions.

NOTE: If you use the 'repeat' command on the switch, the repeat loop can be stopped using the **<CTRL>-c** key combination.

Example outputs:

Before any traffic was received on the port:

```
ICX-Tx-Access1-vsa(config-if)# show aaa authentication port-access interface
1/1/3 client-status

Port Access Client Status Details
ICX-Tx-Access1-vsa(config-if)# do repeat delay 5
```

Once the first frame from the client was received:

```
ICX-Tx-Access1-vsa(config-if)# show aaa authentication port-access interface
1/1/3 client-status

Port Access Client Status Details

Client 00:50:56:b1:7a:37
============================
  Session Details
  ---------------
    Port         : 1/1/3
    Session Time : 1s

  Authentication Details
  ----------------------
    Status          : Authenticating
    Auth Precedence : dot1x - Initialized, mac-auth - Not attempted
```

While the 802.1X is re-trying:

```
ICX-Tx-Access1-vsa(config-if)# show aaa authentication port-access interface
1/1/3 client-status

Port Access Client Status Details

Client 00:50:56:b1:7a:37
============================
  Session Details
  ---------------
    Port         : 1/1/3
    Session Time : 3s

  Authentication Details
  ----------------------
    Status          : Authenticating
    Auth Precedence : dot1x - Authenticating, mac-auth - Not attempted
```

Once the 802.1X authentication is considered 'failed' , it will take about 60 seconds with the configured timers for MAC-auth to be performed.

```
ICX-Tx-Access1-vsa(config-if)# show aaa authentication port-access interface
1/1/3 client-status

Port Access Client Status Details

Client 00:50:56:b1:7a:37, 005056b17a37
============================
  Session Details
```

```
---------------
  Port         : 1/1/3
  Session Time : 63s

Authentication Details
----------------------
  Status          : mac-auth Authenticated
  Auth Precedence : dot1x - Unauthenticated, mac-auth - Authenticated

Authorization Details
---------------------
  Role   : employee
  Status : Applied
```

13. If you used the 'repeat' command, use <CTRL>-C to stop the output.

This ensures that any device that is not 802.1X capable, such as IOT devices, will still be able to access the network using MAC-auth.

Notice that the current role of the device is 'employee': this was assigned by the MAC-auth authentication process.

Verify 802.1X Authentication Precedence over MAC-auth

In the next steps, the 802.1X default precedence over MAC-auth will be demonstrated. This means that in case if both MAC-auth and 802.1X would succeed, the 802.1X authentication method will control the final access for the client device.

Diagram

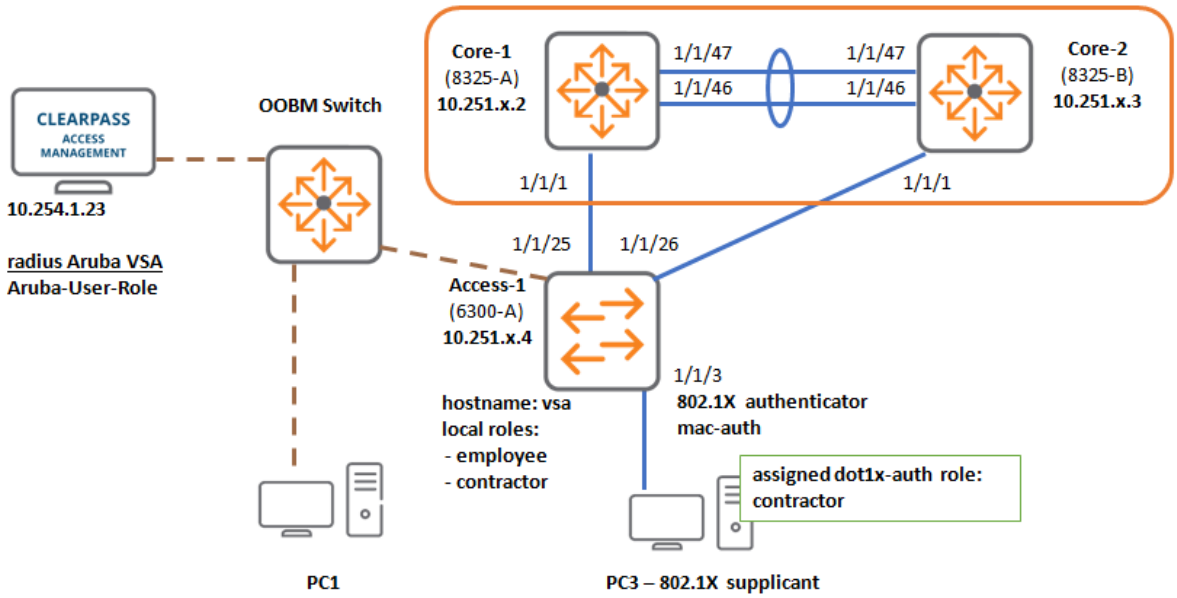

14. On PC3, enable 802.1X again, make sure to enter the credentials again.
    - Username: **icx-contractor**
    - Password: **aruba123**

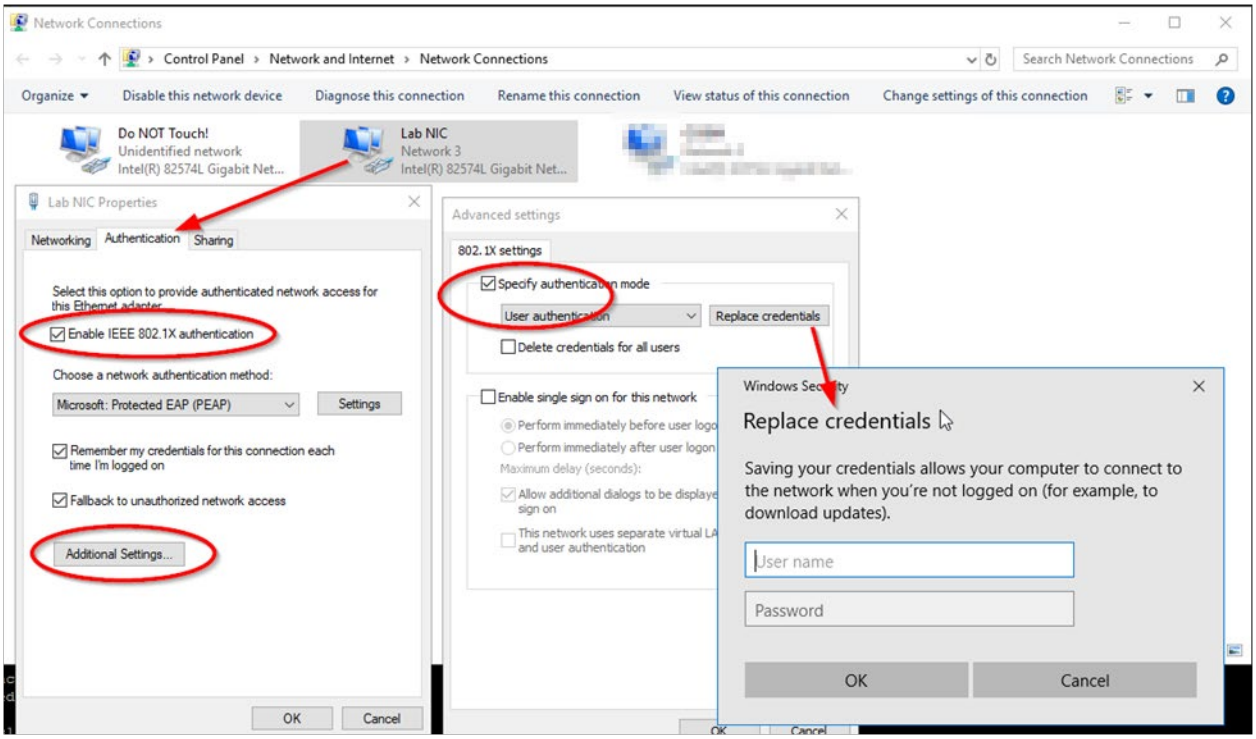

15. On the Access1, check the authenticated client status.

```
ICX-Tx-Access1-vsa(config-if)# show aaa authentication port-access interface
1/1/3 client-status

Port Access Client Status Details

Client 00:50:56:b1:7a:37, icx-contractor
============================
  Session Details
  ---------------
    Port         : 1/1/3
    Session Time : 295s

  Authentication Details
  ----------------------
    Status          : dot1x Authenticated
    Auth Precedence : dot1x - Authenticated, mac-auth - Not attempted

  Authorization Details
  ----------------------
    Role   : contractor
    Status : Applied
```

```
ICX-Tx-Access1-vsa(config-if)#
```

Q: What is the role that was applied to the client?

A: The role is 'contractor'. This is based on the 802.1X authentication since 802.1X authentication has precedence over a successful MAC-authentication.

Optional step: Verify Combination of Failed 802.1X and Successful MAC-auth

In the next steps (optional), you will see that in case a device fails 802.1X authentication, it can still succeed MAC-auth and get controlled network access based on MAC-auth.

Diagram

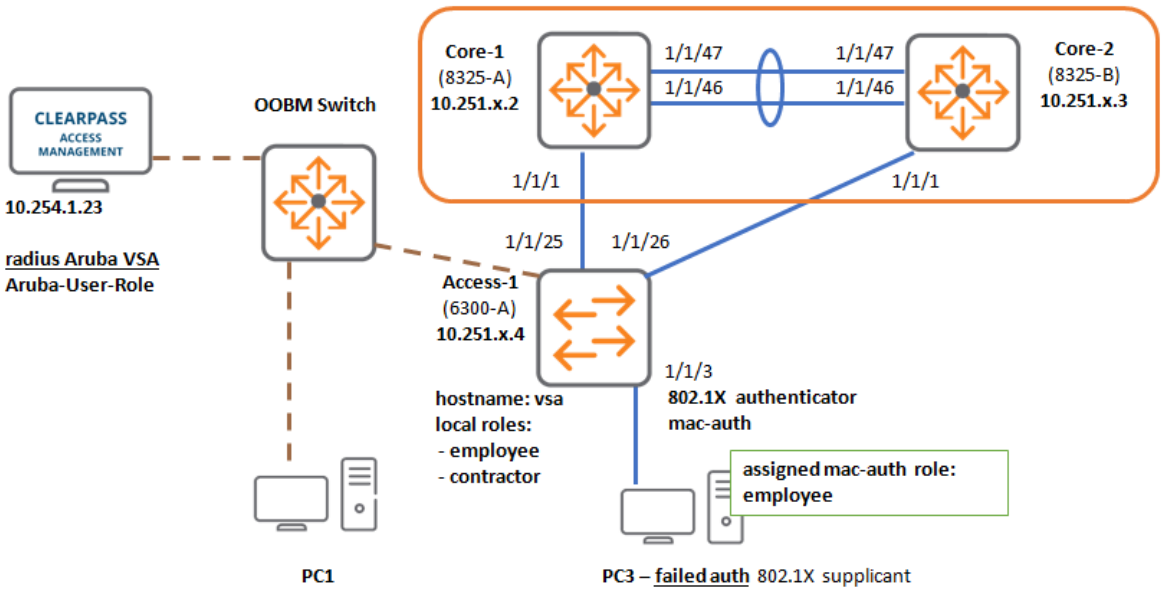

- On the PC3, update the 802.1X credentials on the 'Lab Nic' to incorrect credentials:
  - username **icx-contractor**
  - password **bad**
- On Access1, verify the result.

Example output of the 802.1X session with the failed login:

```
ICX-Tx-Access1-vsa(config-if)# show aaa authentication port-access interface
1/1/3 client-status

Port Access Client Status Details

Client 00:50:56:b1:7a:37, 005056b17a37
============================
  Session Details
  ---------------
    Port         : 1/1/3
    Session Time : 452s

  Authentication Details
  ----------------------
    Status          : mac-auth Authenticated
    Auth Precedence : dot1x - Held, mac-auth - Authenticated

  Authorization Details
  ---------------------
    Role   : employee
    Status : Applied
```

And a few moments later, once the 802.1X authentication has failed:

```
ICX-Tx-Access1-vsa(config-if)# show aaa authentication port-access interface
1/1/3 client-status

Port Access Client Status Details

Client 00:50:56:b1:7a:37, 005056b17a37
============================
  Session Details
  ---------------
    Port         : 1/1/3
    Session Time : 35s

  Authentication Details
  ----------------------
    Status          : mac-auth Authenticated
    Auth Precedence : dot1x - Unauthenticated, mac-auth
 - Authenticated

  Authorization Details
  ---------------------
    Role   : employee
    Status : Applied

ICX-Tx-Access1-vsa(config-if)# exit
```

End of the optional step.

Optional Task 6: Device Profiles with LLDP

This task is **optional** and can be done if time permits. Check with your instructor.

> **IMPORTANT:** Make sure to complete Task 7 at the end of this lab activity.

Diagram

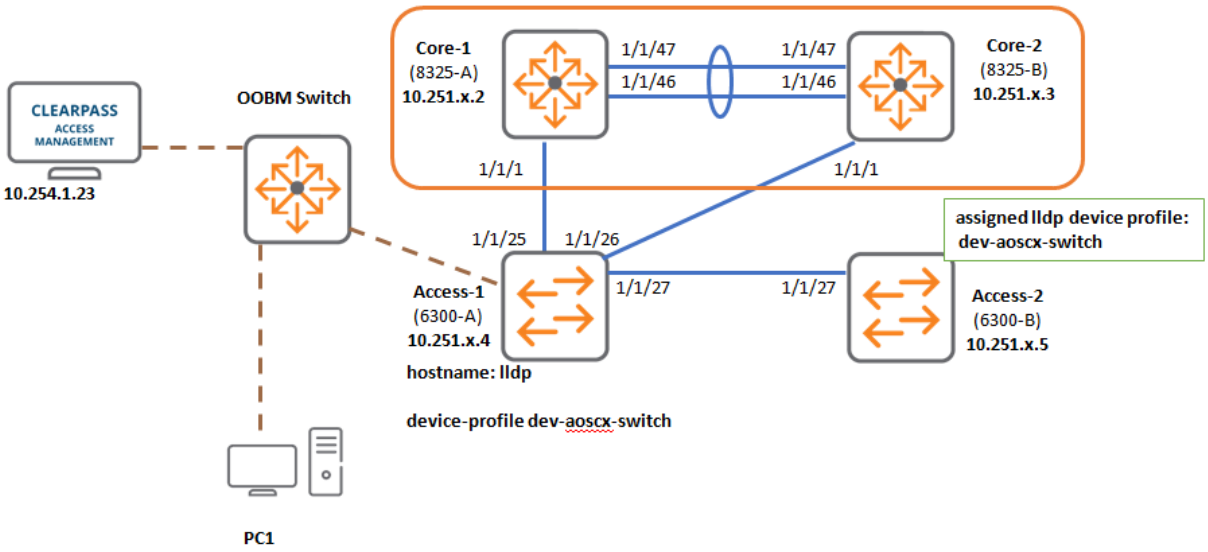

Objectives

In this task, the AOS-CX device profiles will be demonstrated. This is not a security feature, but a convenience feature, similar to MAC-auth.

In this case, instead of using the client MAC address for authentication, like MAC-auth, it is possible to configure a switch port based on the incoming LLDP or CDP messages. Since many network devices support LLDP or CDP, this means that the switch port of those devices can be easily provisioned based on the received LLDP communication from these network devices.

LLDP is more difficult to spoof compared to a client MAC address, but it should still be understood that there is no actual authentication happening; so with a more advanced toolset, the LLDP messages can still be spoofed. This is why the Aruba device profiles feature is considered a 'convenience' feature.

Some examples for real world deployment:

- Aruba AP is connected to any switch port. The switch assigns a device profile with the correct VLAN untagged and tagged states.
- Any vendor AP that supports LLDP or CDP is connected to any switch port. The switch assigns a device profile with the correct VLAN untagged and tagged states.

- Any vendor media device, such as a conference room controller, that supports LLDP or CDP, can now automatically be assigned to the correct VLAN.

All of these examples are handled by the switch itself; they do not require a RADIUS server or other authentication infrastructure. While it is possible to have the same results with 802.1X or MAC-auth, not every customer may be willing to setup the RADIUS infrastructure. In that case, device profiles can assist to simplify the life of the network administrator.

In this task, the Access2 switch will be used as the 'test LLDP' device that is connected to the Access1. Based on the Access2 LLDP communication, the switch port on Access1 will be automatically configured with the correct settings based on the device profile that will be prepared.

Steps

These are the steps used in this task:

- Collect information of the peer LLDP device
- Prepare the device profile
- Associate the device profile to LLDP settings
- Prepare the test > hostname
- Verify the configuration

Collect Information of the Peer LLDP Device

Access1

1. Open a terminal session on Access1 and enter the configuration mode.

Change the hostname to get the Aruba VSA Aruba-User-Role ClearPass rule set.

2. Change the hostname, replace the **'vsa'** the hostname with **'lldp'**. Change “Tx” to “T” followed by your table number; for example, Table 1 would be “T1” and Table 12 would be T12.

```
ICX-Tx-Access1-vsa(config)# hostname ICX-Tx-Access1-lldp
ICX-Tx-Access1-lldp(config)#
```

> **NOTE:** The hostname is used as the NAS-identifier. ClearPass has been pre-configured to look for the value **'lldp'** in the NAS-identifier, in this case it will return an **access-reject** for the MAC-authentication. This is required for this lab, since the local device profile will be used to assign the user-role instead of the RADIUS server.

3. Enable port 1/1/27, this is the port connected to Access2.

```
ICX-Tx-Access1-lldp(config)# interface 1/1/27
ICX-Tx-Access1-lldp(config-if)# no shutdown
ICX-Tx-Access1-lldp(config-if)# exit
```

4. Review the current LLDP neighbors.

```
ICX-Tx-Access1-lldp(config)# show lldp neighbor-info

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

Total Neighbor Entries          : 2
Total Neighbor Entries Deleted  : 5
Total Neighbor Entries Dropped  : 0
Total Neighbor Entries Aged-Out : 3

LOCAL-PORT  CHASSIS-ID         PORT-ID       PORT-DESC     TTL      SYS-NAME
--------------------------------------------------------------------------------
1/1/25      90:20:c2:bc:17:00  1/1/1         1/1/1         120      ICX-T12-C...
1/1/26      90:20:c2:bc:97:00  1/1/1         1/1/1         120      ICX-T12-C...
ICX-Tx-Access1-lldp(config)#
```

Q: Is Access2 switch listed as an LLDP neighbor?

_____

A: No, on ports with port-access authentication enabled, LLDP packets are disabled by default. LLDP would start to operate once a device has been successfully authenticated.

5. Enable LLDP packets on the port-access protected interface. Review the port-access options, next enable LLDP BPDUs.

```
ICX-Tx-Access1-lldp(config)# interface 1/1/27
ICX-Tx-Access1-lldp(config-if)# aaa authentication port-access allow-?
  allow-cdp-bpdu   Allow or block cdp bpdu.
  allow-lldp-bpdu  Allow or block lldp bpdu.

ICX-Tx-Access1-lldp(config-if)# aaa authentication port-access allow-lldp-bpdu
ICX-Tx-Access1-lldp(config-if)# exit
```

6. Wait up to 30 seconds, then review the LLDP neighbors. On port 1/1/27 there should be an LLDP neighbor now.

```
ICX-Tx-Access1-lldp(config)# show lldp neighbor-info

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

Total Neighbor Entries          : 3
Total Neighbor Entries Deleted  : 5
```

```
Total Neighbor Entries Dropped  : 0
Total Neighbor Entries Aged-Out : 3

LOCAL-PORT  CHASSIS-ID         PORT-ID       PORT-DESC     TTL      SYS-NAME
--------------------------------------------------------------------------------
1/1/25      90:20:c2:bc:17:00  1/1/1         1/1/1         120      ACSP-P52-...
1/1/26      90:20:c2:bc:97:00  1/1/1         1/1/1         120      ACSP-P52-...
1/1/27      88:3a:30:97:b6:00  1/1/27        1/1/27        120      ACSP-P52-...
```

7. Review the details for the neighbor on interface 1/1/27. The device profile can be applied to the **chassis-name** and the **chassis-description** information.

It is also possible to assign the device profile based on a vendor specific LLDP TLV, but that feature is less frequently implemented. The Aruba APs, AOS-Switch and AOS-CX switches do report a specific Vendor OUI, so it is possible to apply the vendor OUI filter for those devices.

In this lab, the assignment will be done based on the LLDP system description. The Access2 switch is a specific switch model, and the AOS-CX switches report their model code (JXXXX) in the LLDP system description.

Take note of the Jxxxx code in the system description in this lab guide: it is JL668A or JL660A (Pod 11).

> **NOTE:** The lab assumes that the Core switches are using a different model number than the Access switches. If your lab would be using the same model number, it is possible to use the system name (access) instead of the model number.

```
ICX-Tx-Access1-lldp(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 Chassis-Name          : ICX-Tx-Access2
Neighbor Chassis-Description   : Aruba JL668A  FL.10.04.0030
Neighbor Chassis-ID            : 88:3a:30:97:b6:00
Neighbor Management-Address    : 10.251.x.5
Chassis Capabilities Available : Bridge, Router
Chassis Capabilities Enabled   : Bridge, Router
Neighbor Port-ID               : 1/1/27
Neighbor Port-Desc             : 1/1/27
Neighbor Port VLAN ID          : 1
TTL                            : 120

Neighbor Mac-Phy details
Neighbor Auto-neg Supported    : true
Neighbor Auto-Neg Enabled      : false
Neighbor Auto-Neg Advertised   :
```

```
Neighbor MAU type               : 10 GIGBASEER

ICX-Tx-Access1-lldp(config-if)#
```

8. Define a new role that will be used in this task. The device profile feature also uses the same port-access roles that are used by the 802.1X and MAC-auth modules. For the purpose of this lab a dedicated role with be defined.

NOTE: For this lab demo device-mode authentication is enabled. This means that the port will be '**open**' for any mac-address after this device profile has been applied. This would apply to Aruba Instant APs, for example, since the AP will be performing the wireless user authentication.

```
ICX-Tx-Access1-lldp(config)# port-access role demo-lldp-switch
ICX-Tx-Access1-lldp(config-pa-role)# auth-mode device-mode
ICX-Tx-Access1-lldp(config-pa-role)# vlan trunk native 11
ICX-Tx-Access1-lldp(config-pa-role)# vlan trunk allowed 11,12
ICX-Tx-Access1-lldp(config-pa-role)# exit
ICX-Tx-Access1-lldp(config)#
```

Define an LLDP Group to Match the Peer LLDP Device

With an LLDP group, the administrator can assign LLDP devices based on the reported System Description, System Name, or a Vendor OUI TLV in the LLDP frame.

In this lab, the assignment will be done based on the LLDP system description. The Access2 switch is a specific switch model, and the AOS-CX switches report their model code (JXXXX) in the LLDP system description.

9. Define a new LLDP group profile. This LLDP group will match on the LLDP system description of the neighbor LLDP switch.

```
ICX-Tx-Access1-lldp(config)# port-access lldp-group demo-switch
```

10. Review the selection criteria of the match command.

```
ICX-Tx-Access1-lldp(config-lldp-group)# match ?
  sys-desc    Configure LLDP system description.
  sysname     Configure LLDP system name.
  vendor-oui  Configure LLDP vendor OUI.
  <cr>
```

11. Define a new match criteria based on the LLDP neighbor system description.

NOTE: All students should verify the model number of Access1 and Access2 – For example, in **Pod11** they are model **JL660A**.

```
ICX-Tx-Access1-lldp(config-lldp-group)# match sys-desc JL668A
```

12. Review the current configuration of the LLDP group and exit the group context.

```
ICX-Tx-Access1-lldp(config-lldp-group)# show running-config current-context
port-access lldp-group demo-switch
     seq 10 match sys-desc JL668A
ICX-Tx-Access1-lldp(config-lldp-group)# exit
```

Define a Device Profile to Bind the LLDP Group to a User-role

13. Define a new port-access device profile.

```
ICX-Tx-Access1-lldp(config)# port-access device-profile dev-aoscx-switch
```

14. Bind the LLDP group and the user-role to this device profile and enable it.

```
ICX-Tx-Access1-lldp(config-device-profile)# associate role demo-lldp-switch
ICX-Tx-Access1-lldp(config-device-profile)# associate lldp-group demo-switch
ICX-Tx-Access1-lldp(config-device-profile)# enable
ICX-Tx-Access1-lldp(config-device-profile)# exit
```

15. Review the current device profile configuration.

```
ICX-Tx-Access1-lldp(config)# show port-access device-profile

    Profile Name            : dev-aoscx-switch
    LLDP Groups             : demo-switch
    CDP Groups              :
    Role                    : demo-lldp-switch
    State                   : Enabled
ICX-Tx-Access1-lldp(config)#
```

16. Reset the port connected to Access2.

```
ICX-Tx-Access1-lldp(config)# int 1/1/27
ICX-Tx-Access1-lldp(config-if)# shutdown
ICX-Tx-Access1-lldp(config-if)# no shutdown
ICX-Tx-Access1-lldp(config-if)# exit
```

17. Review the device-profile status for all interfaces.

```
ICX-Tx-Access1-lldp(config)# show port-access device-profile interface all
  Port 1/1/27, Neighbor-Mac  88:3a:30:97:b6:00
    Profile Name:           : dev-aoscx-switch
    LLDP Group:             : demo-switch
    CDP Group:              :
    Role:                   : demo-lldp-switch
    State:                  : applied
    Failure Reason:         :
```

18. Review that no other successful authentication device is online on this port.

```
ICX-Tx-Access1-lldp(config)# show aaa authentication port-access interface 1/1/27
client-status

Port Access Client Status Details
ICX-Tx-Access1-lldp(config)#
```

19. Review that the port VLAN status is now active based on the VLAN settings of the device profile user-role.

```
ICX-Tx-Access1-lldp(config)# show vlan port 1/1/27

--------------------------------------------------------------------------------
VLAN  Name                            Mode            Mapping
--------------------------------------------------------------------------------
11    VLAN11                          native-untagged port-access
12    VLAN12                          trunk           port-access

Overridden VLAN list: 1
```

This method can be used to auto-provision the port and VLAN settings for many LLDP device types.

20. Disable the device profile feature.

```
ICX-Tx-Access1-lldp(config)# port-access device-profile dev-aoscx-switch
ICX-Tx-Access1-lldp(config-device-profile)# disable
ICX-Tx-Access1-lldp(config-device-profile)# exit
```

21. Verify that the VLAN configuration of the port 1/1/27 has been reverted.

```
ICX-Tx-Access1-lldp(config)# show vlan port 1/1/27

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

22. Remove the 'allow-lldp-bpdu' feature on the port 1/1/27.

```
ICX-Tx-Access1(config)# interface 1/1/27
ICX-Tx-Access1(config-if)# no aaa authentication port-access allow-lldp-bpdu
ICX-Tx-Access1(config-if)# exit
```

Task 7: Save checkpoint configuration

This task is required, upcoming lab activities are based on this configuration.

Steps

Access1

1. On Access1, change the hostname back to the default lab hostname. Change “Tx” to “T” followed by your table number; for example, Table 1 would be “T1” and Table 12 would be T12.

```
hostname ICX-Tx-Access1
```

2. Save checkpoint **'icx-lab11-mac**'.

```
ICX-Tx-Access1(config)# end
ICX-Tx-Access1# copy run checkpoint icx-lab11-mac
Configuration changes will take time to process, please be patient.
```

Access2

3. Save checkpoint **'icx-lab11-mac**'.

```
ICX-Tx-Access2(config)# end
ICX-Tx-Access2# copy run checkpoint icx-lab11-mac
Configuration changes will take time to process, please be patient.
```

You’ve completed Lab 11!

Lab 12.1: Integration with Aruba CPPM- ClearPass Downloadable User Roles

Lab Diagram

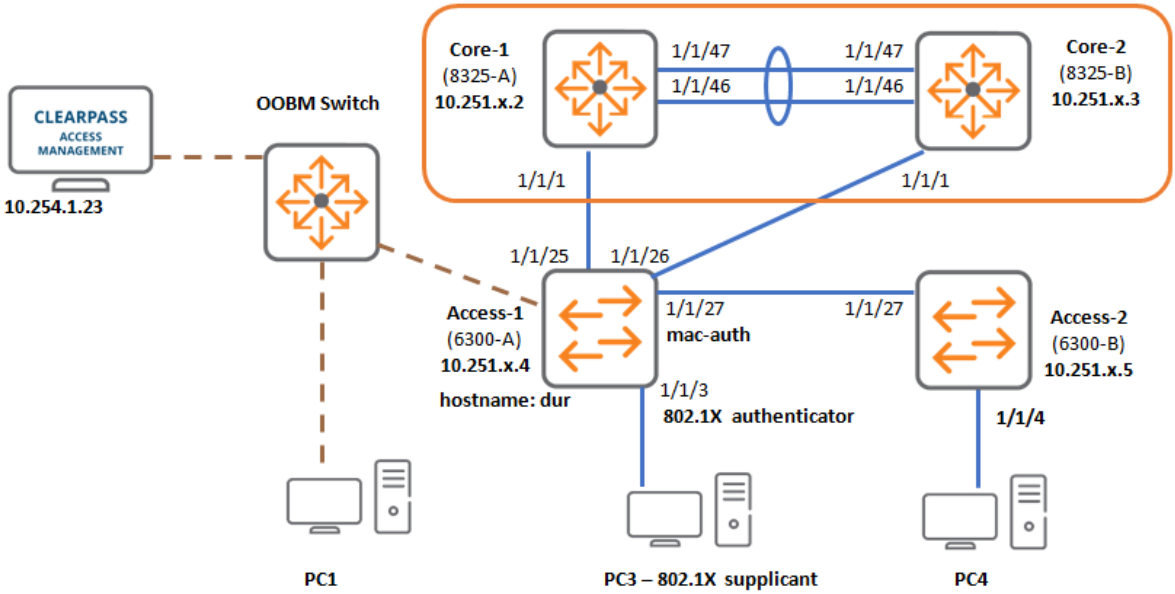

Overview

This lab requires the completion of the 802.1X (Lab10) and MAC authentication (Lab11) lab activities.

This lab activity will demonstrate how ClearPass downloadable user roles can be used to simplify the deployment of the authentication user roles to the switches. The roles are only defined on the ClearPass host and they will be downloaded on the fly, as needed, by the access switches.

This means that the administrator still needs to understand the concept of roles. but it is not necessary anymore to check that all the switches have the latest configuration/policy version of all these user roles in their configuration, since this will be done automatically for the administrator. This will also reduce the operational configuration file of the access switches, since these roles are not saved in the configuration file.

The switch will download the definition from ClearPass as part of the authentication process using the REST API; therefore, credentials must be configured on the switch to access the ClearPass host to download the role configuration.

Objectives

- Define secure HTTPS access between the switch and ClearPass by importing the Trust Anchor certificate
- Define the credentials to access the ClearPass REST API
- Verify the operation of the downloadable user roles

Task 1: CPPM REST API Communication

Diagram

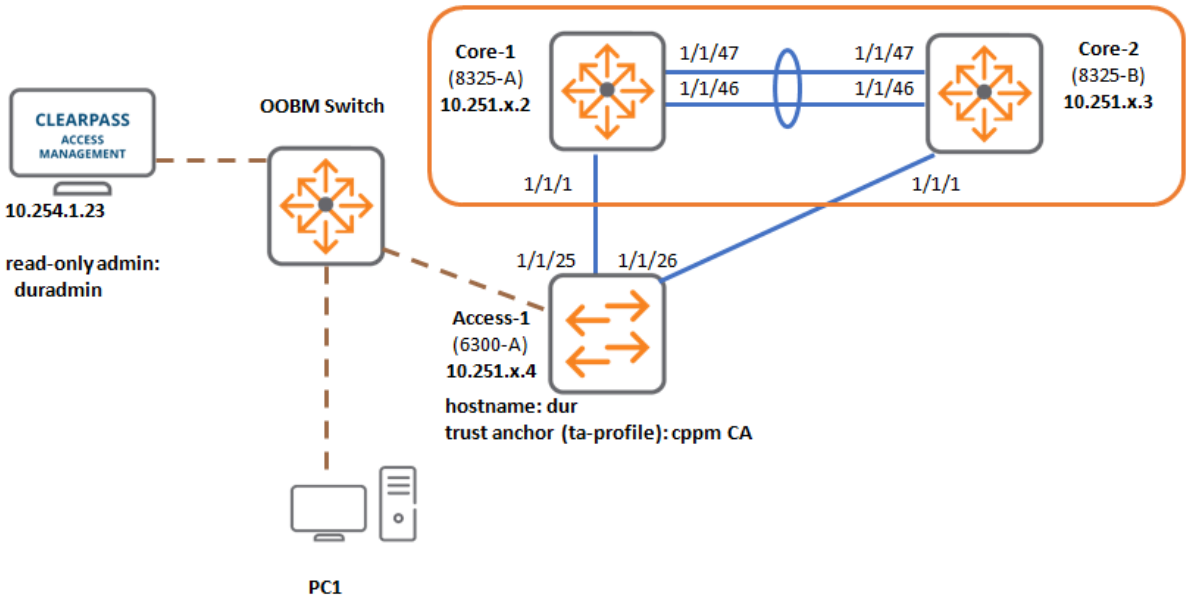

Objectives

Review the requirements on the switch to validate the CPPM certificate: check NTP, and the name resolution for the certificate subject name

Verify CPPM has been configured with a read-only admin user account.

On the Access switch, install the CPPM CA certificate as a 'trust anchor'. This will enable the switch to validate the CPPM server certificate. This procedure can be done using the CLI or using NetEdit.

On the Access switch, configure the (read-only) admin account that the switch should use to access the CPPM RESTAPI when downloading user-roles.

Steps

Verify Requirements for Certificate Validation

The AOS-CX switch should be able to verify the validity of the ClearPass server certificate. This requires a valid time to be set on the switch and the switch should have a valid certificate chain (Trust Anchor) to validate the certificate on the ClearPass server.

In the next steps, these requirements will be configured and verified.

Review Time is Synchronized Correctly

1. Open a terminal session on Access1, enter the configuration mode
2. Verify the time.

```
ICX-Tx-Access1(config)# show clock
Mon Jan 27 18:14:35 EDT 2020
System is configured for timezone : US/Eastern
```

3. Verify the NTP status.

```
ICX-Tx-Access1(config)# show ntp status
NTP Status Information

NTP                          : Enabled
NTP Authentication           : Disabled
NTP Server Connections       : Using the mgmt VRF

System time                  : Mon Jan 27 18:14:31 EDT 2020
NTP uptime                   : 6 hours, 31 minutes, 39 seconds
```

Trust Anchor (TA) Installation

This installation can be done using the CLI or using NetEdit. The lab guide provides instructions for both options. You should only use one of these options, you are free to select which option you prefer.

After the instructions of Option1 you will find Option 2: CLI based TA Installation.

Option1: NetEdit-Based TA Installation

4. On PC1, open a browser to NetEdit (https://10.251.x.200), login using admin/aruba123.

> **TIP (Optional):** At the end of the NetEdit deployment, you will verify the installation of the TA certificate using the command 'show crypto pki ta-profile'.
>
> You can add this command to the NetEdit 'Change Validation' command list to NetEdit complete this validation command. If you want this, apply>
>
> NetEdit > Settings > Validation > Change Validation > 'Command Scripts' > 'System Information'. Add a line with 'show crypto pki ta-profile' , click 'OK' to save.

5. Navigate to **Devices**, select **'ICX-Tx-Access1'**.

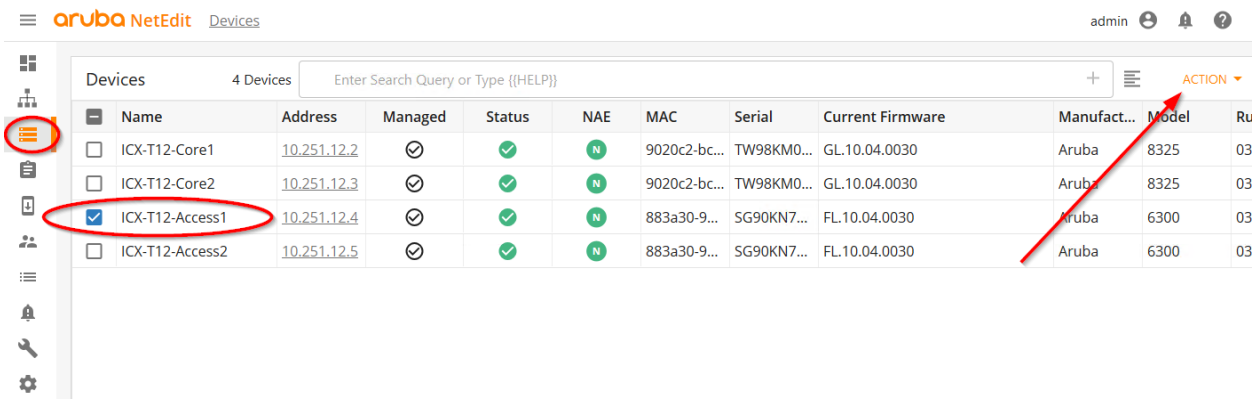

6. Use **'Action'** menu > **'Edit Config'**.
7. For the plan name, enter **'import ta'**, click **'Create'** to continue.
8. Place the cursor at the start of the first line, press <ENTER> to add a new line.

9. Define a new Trust Anchor with name **'cppm'** profile by entering, press <ENTER> after this line.

```
crypto pki ta-profile cppm
```

10. NetEdit will move the line to the correct location in the configuration.

11. The cursor will be on the next line, press **<TAB>** to enter the ta-profile context, and then enter **'ta-certificate'**, press **<ENTER>** to submit the command.

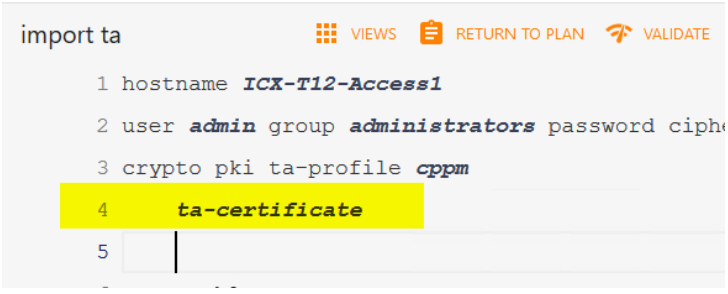

12. Now move the mouse over the '**ta-certificate**' word, the mouse icon will change, Next **right-click** to access the import screen. This screen allows the CA Trust Certificate to be pasted and then imported by the device.

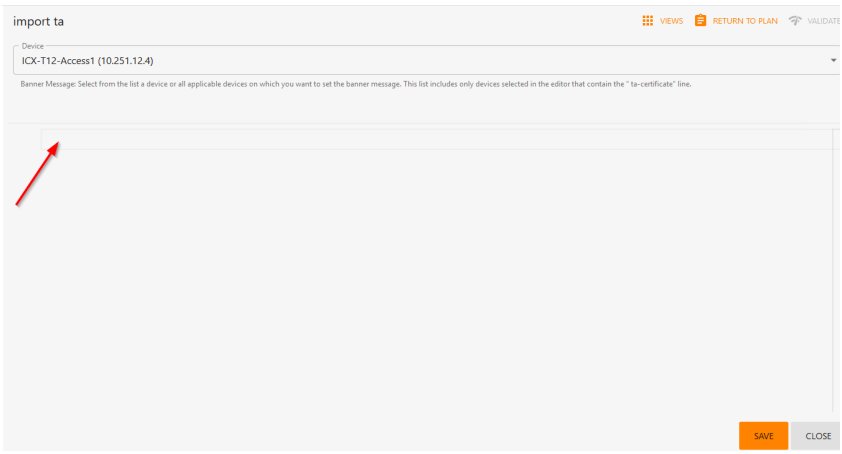

13. On the PC1, open the **ICX-Files** folder on the desktop and open the file:

lab 12.1 - clearpass - root - ca.txt

14. If you are unable to paste the content to PC1, CTRL-ALT-SHIFT will open a clipboard. Paste the content into the clipboard, CTRL-ALT-SHIFT to close the clipboard and then paste the contents into PC1
15. **'Select all'** the contents (CTRL-A) and use **'copy'** (CTRL-C). Example screenshot.

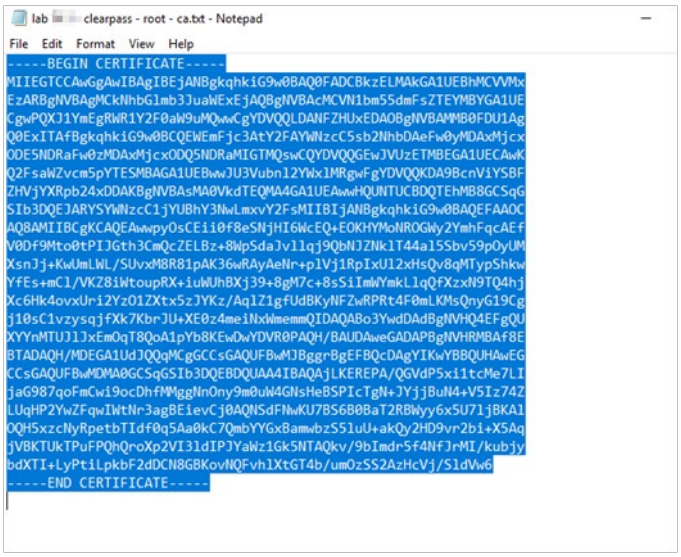

16. Switch back to the NetEdit Browser, paste the certificate text in the message field.

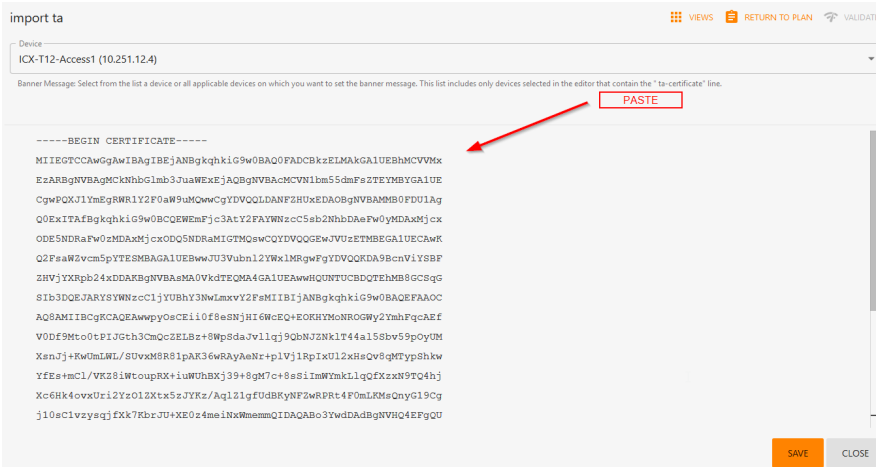

17. Click **'Save'** to return to the Plan Editor.
18. Click **'Return to Plan'**.
19. Click **'Deploy'** and confirm with **'Deploy'.**
20. Once the deployment has completed, use **'Commit'** and confirm with **'Commit'**.

Validation of import on Access1

21. Open a terminal connection to Access1, enter the configuration mode.
22. Review the installed ta profiles.

```
ICX-Tx-Access1(config)# show crypto pki ta-profile

TA Profile Name                  TA Certificate       Revocation Check
-------------------------------- -------------------- ----------------
cppm                             Installed, valid     disabled
```

TIP:

If you had the validation command defined in NetEdit based on the previous tip the plan validation would show this result:

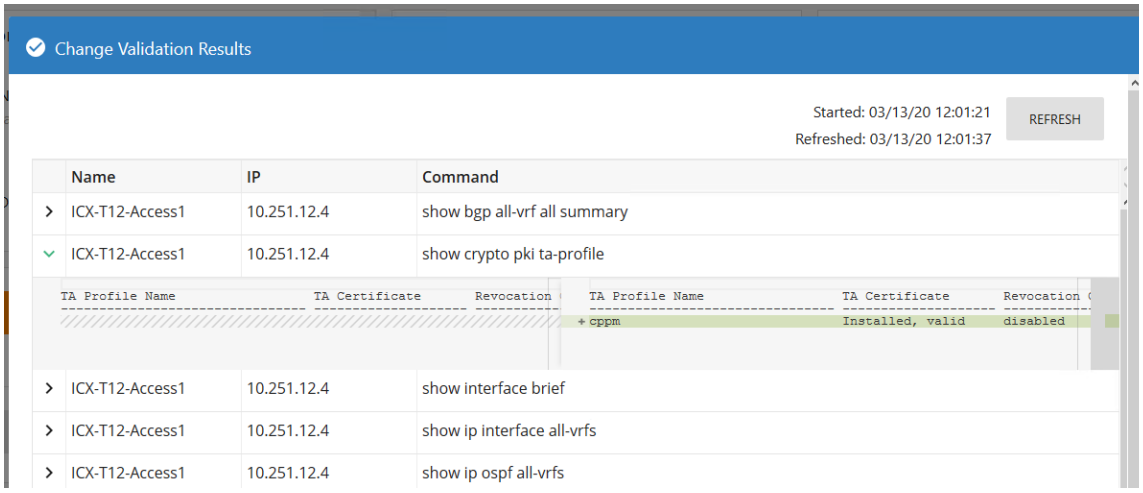

End of Option1: NetEdit based TA Installation

Option2: CLI based TA Installation

IMPORTANT: If you used the NetEdit method to import the TA certificate, you should skip this option.

23. Review the default Trust Anchor configuration. There should be no root CAs installed by default.

```
ICX-Tx-Access1(config)# show crypto pki ta-profile

TA Profile Name                  TA Certificate       Revocation Check
-------------------------------- -------------------- ----------------
```

24. Define a new TA profile 'cppm'.

```
ICX-Tx-Access1(config)# crypto pki ta-profile cppm
```

25. Run the 'import' command.

```
ICX-Tx-Access1(config-ta-cppm)# ta-certificate import
Paste the certificate in PEM format below, then hit enter and ctrl-D:
```

26. On PC1, open the **ICX-Files** folder on the desktop).
27. Open the cert file: "**lab 12.1 - clearpass - root - ca.txt**", copy the contents and paste the content in the CLI.

NOTE:

The file can manually be downloaded from ClearPass on this 'well-known' URL:

http://<cppm-ip>/.well-known/aruba/clearpass/https-root.pem

For the lab this would be:

http://10.254.1.23/.well-known/aruba/clearpass/https-root.pem

28. If you are unable to paste the content to Access1, CTRL-ALT-SHIFT will open a clipboard. Paste the content into the clipboard, CTRL-ALT-SHIFT to close the clipboard and then paste the contents into Access1.

29. Font size has been reduced in the example output:

```
ICX-T12-Access1(config-ta-cert)# -----BEGIN CERTIFICATE-----
MIIEGTCCAwGgAwIBAgIBEjANBgkqhkiG9w0BAQ0FADCBkzELMAkGA1UEBhMCVVMx
ICX-T12-Access1(config-ta-cert)# MIIEGTCCAwGgAwIBAgIBEjANBgkqhkiG9w0BAQ0FADCBkzELMAkGA1UEBhMCVVMx
ICX-T12-Access1(config-ta-cert)# EzARBgNVBAgMCkNhbGlmb3JuaWExEjAQBgNVBAcMCVN1bm55dmFsZTEYMBYGA1UE
ICX-T12-Access1(config-ta-cert)# CgwPQXJ1YmEgRWR1Y2F0aW9uMQwwCgYDVQQLDANFZHUxEDAOBgNVBAMMB0FDU1Ag
ICX-T12-Access1(config-ta-cert)# Q0ExITAfBgkqhkiG9w0BCQEWEmFjc3AtY2FAYWNzcC5sb2NhbDAeFw0yMDAxMjcx
ICX-T12-Access1(config-ta-cert)# ODE5NDRaFw0zMDAxMjcxODQ5NDRaMIGTMQswCQYDVQQGEwJVUzETMBEGA1UECAwK
```

```
ICX-T12-Access1(config-ta-cert)# Q2FsaWZvcm5pYTESMBAGA1UEBwwJU3Vubnl2YWxlMRgwFgYDVQQKDA9BcnViYSBF
ICX-T12-Access1(config-ta-cert)# ZHVjYXRpb24xDDAKBgNVBAsMA0VkdTEQMA4GA1UEAwwHQUNTUCBDQTEhMB8GCSqG
ICX-T12-Access1(config-ta-cert)# SIb3DQEJARYSYWNzcC1jYUBhY3NwLmxvY2FsMIIBIjANBgkqhkiG9w0BAQEFAAOC
ICX-T12-Access1(config-ta-cert)# AQ8AMIIBCgKCAQEAwwpyOsCEii0f8eSNjHI6WcEQ+EOKHYMoNROGWy2YmhFqcAEf
ICX-T12-Access1(config-ta-cert)# V0Df9Mto0tPIJGth3CmQcZELBz+8WpSdaJvllqj9QbNJZNklT44al5Sbv59pOyUM
ICX-T12-Access1(config-ta-cert)# XsnJj+KwUmLWL/SUvxM8R81pAK36wRAyAeNr+plVj1RpIxUl2xHsQv8qMTypShkw
ICX-T12-Access1(config-ta-cert)# YfEs+mCl/VKZ8iWtoupRX+iuWUhBXj39+8gM7c+8sSiImWYmkLlqQfXzxN9TQ4hj
ICX-T12-Access1(config-ta-cert)# Xc6Hk4ovxUri2YzO1ZXtx5zJYKz/AqlZ1gfUdBKyNFZwRPRt4F0mLKMsQnyG19Cg
ICX-T12-Access1(config-ta-cert)# j10sC1vzysqjfXk7KbrJU+XE0z4meiNxWmemmQIDAQABo3YwdDAdBgNVHQ4EFgQU
ICX-T12-Access1(config-ta-cert)# XYYnMTUJlJxEmOqT8QoA1pYb8KEwDwYDVR0PAQH/BAUDAweGADAPBgNVHRMBAf8E
ICX-T12-Access1(config-ta-cert)# BTADAQH/MDEGA1UdJQQqMCgGCCsGAQUFBwMJBggrBgEFBQcDAgYIKwYBBQUHAwEG
ICX-T12-Access1(config-ta-cert)# CCsGAQUFBwMDMA0GCSqGSIb3DQEBDQUAA4IBAQAjLKEREPA/QGVdP5xi1tcMe7LI
ICX-T12-Access1(config-ta-cert)# jaG987qoFmCwi9ocDhfMMggNnOny9m0uW4GNsHeBSPIcTgN+JYjjBuN4+V5Iz74Z
ICX-T12-Access1(config-ta-cert)# LUqHP2YwZFqwIWtNr3agBEievCj0AQNSdFNwKU7BS6B0BaT2RBWyy6x5U7ljBKAl
ICX-T12-Access1(config-ta-cert)# OQH5xzcNyRpetbTIdf0q5Aa0kC7QmbYYGxBamwbzS5luU+akQy2HD9vr2bi+X5Aq
ICX-T12-Access1(config-ta-cert)# jVBKTUkTPuFPQhQroXp2VI3ldIPJYaWz1Gk5NTAQkv/9bImdr5f4NfJrMI/kubjy
ICX-T12-Access1(config-ta-cert)# bdXTI+LyPtiLpkbF2dDCN8GBKovNQFvhlXtGT4b/umOzSS2AzHcVj/SldVw6
ICX-T12-Access1(config-ta-cert)# -----END CERTIFICATE-----
ICX-T12-Access1(config-ta-cert)#
```

30. When the paste is complete, enter the '**<CTRL>D**' key combination. This will instruct the switch to parse the pasted cert and show the subject name. Now confirm that this cert ok.

```
ICX-Tx-Access1(config-ta-cert)#
The certificate you are importing has the following attributes:
Subject: C = US, ST = California, L = Sunnyvale, O = Aruba Education, OU = Edu,
CN = ACSP CA, emailAddress = acsp-ca@acsp.local
Issuer:  C = US, ST = California, L = Sunnyvale, O = Aruba Education, OU = Edu,
CN = ACSP CA, emailAddress = acsp-ca@acsp.local
Serial Number: 0x12
TA certificate import is allowed only once for a TA profile
Do you want to accept this certificate (y/n)? y
ICX-Tx-Access1(config-ta-cppm)#
```

31. Exit the context and verify that the TA profile is installed.

> **IMPORTANT:** The actual import will be done when the TA-import context is closed, so make sure to use **'exit'** before checking the CA cert import status.

```
ICX-Tx-Access1(config-ta-cppm)# show crypto pki ta-profile

TA Profile Name                  TA Certificate       Revocation Check
-------------------------------- -------------------- ----------------
```

```
ICX-Tx-Access1(config-ta-cppm)# exit
```

```
ICX-Tx-Access1(config)# show crypto pki ta-profile
```

```
TA Profile Name                  TA Certificate       Revocation Check
-------------------------------- -------------------- ----------------
cppm                             Installed, valid     disabled

ICX-Tx-Access1(config)#
```

End of Option2: CLI based TA Installation

CPPM REST API Read-Only User Account

A read-only administrator has been defined on ClearPass. In the next steps, you will verify that this administrator has been defined.

32. Using the PC1 management station, open a connection to ClearPass on https://10.254.1.23/tips

username **icx-adminX** (replace X with your table number)

Examples: icx-admin1 for table1, icx-admin12 for table12

password **aruba123**

33.. Navigate to **Administration > Users and Privileges > Admin Users**.

34. Verify that an admin account has been created named **'duradmin'** (**downloadable user role** admin), with read-only privileges.

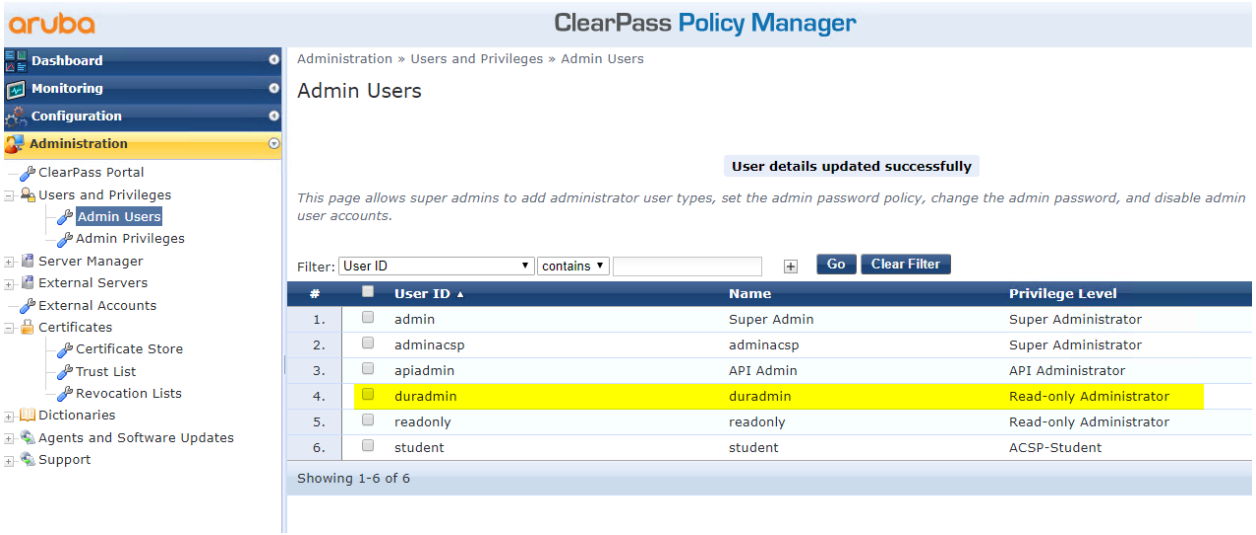

NOTE: In this lab environment, a Read-only Administrator is used. ClearPass also provides a more restrictive admin privilege **Aruba User Role Download** that can be assigned to the download account. This role can only access ClearPass using the RESTAPI and it can only read the enforcement profiles.

Switch: Define the ClearPass Credentials to Access the REST API

35. On the Access1 switch, define the RESTAPI credentials on the ClearPass RADIUS server object with the name ' **cppm.arubatraining.com** '.

NOTE: Pay attention to define the correct **hostname** and VRF **mgmt.**

```
ICX-Tx-Access1(config)# radius-server host cppm.arubatraining.com vrf mgmt
clearpass-username duradmin clearpass-password plaintext aruba123
```

36. Verify the running configuration for RADIUS server lines, there should be only one RADIUS server.

```
ICX-Tx-Access1(config)# show run | include radius-server
radius-server tracking user-name icx-radius-track password ciphertext
AQBapUmeqwuSjUoetq4KWXbTnUyBILPjxzok4qzRZeSXsBIzCQAAACMxHv6lEnY7jA==

radius-server host cppm.arubatraining.com key ciphertext
AQBapVWcNJavUClNBQenFaJwwRrR+nWcJUvsQlHUbuaiOvlDCAAAAMCnYwT2Ful+ tracking enable
clearpass-username duradmin clearpass-password ciphertext
AQBapUmeqwuSjUoetq4KWXbTnUyBILPjxzok4qzRZeSXsBIzCQAAACMxHv6lEnY7jA== vrf mgmt

ICX-Tx-Access1(config)#
```

37. You can also review the RADIUS list in a different format.

```
ICX-Tx-Access1(config)# show radius-server detail
******* Global RADIUS Configuration *******

Shared-Secret: None
Timeout: 5
Auth-Type: pap
Retries: 1
Tracking Time Interval (seconds): 300
Tracking Retries: 1
Tracking User-name: icx-radius-track
Tracking Password:
AQBapUmeqwuSjUoetq4KWXbTnUyBILPjxzok4qzRZeSXsBIzCQAAACMxHv6lEnY7jA==
Number of Servers: 1

****** RADIUS Server Information ******
Server-Name               : cppm.arubatraining.com
Auth-Port                 : 1812
Accounting-Port           : 1813
VRF                       : mgmt
Shared-Secret             :
AQBapVWcNJavUClNBQenFaJwwRrR+nWcJUvsQlHUbuaiOvlDCAAAAMCnYwT2Ful+
Timeout (default)         : 5
Retries (default)         : 1
Auth-Type (default)       : pap
Server-Group              : cppm
Group-Priority            : 1
Tracking                  : enabled
Reachability-Status       : reachable
ClearPass-Username        : duradmin
ClearPass-Password        :
AQBapUmeqwuSjUoetq4KWXbTnUyBILPjxzok4qzRZeSXsBIzCQAAACMxHv6lEnY7jA==
```

Task 2: CPPM User Role Definitions

Diagram

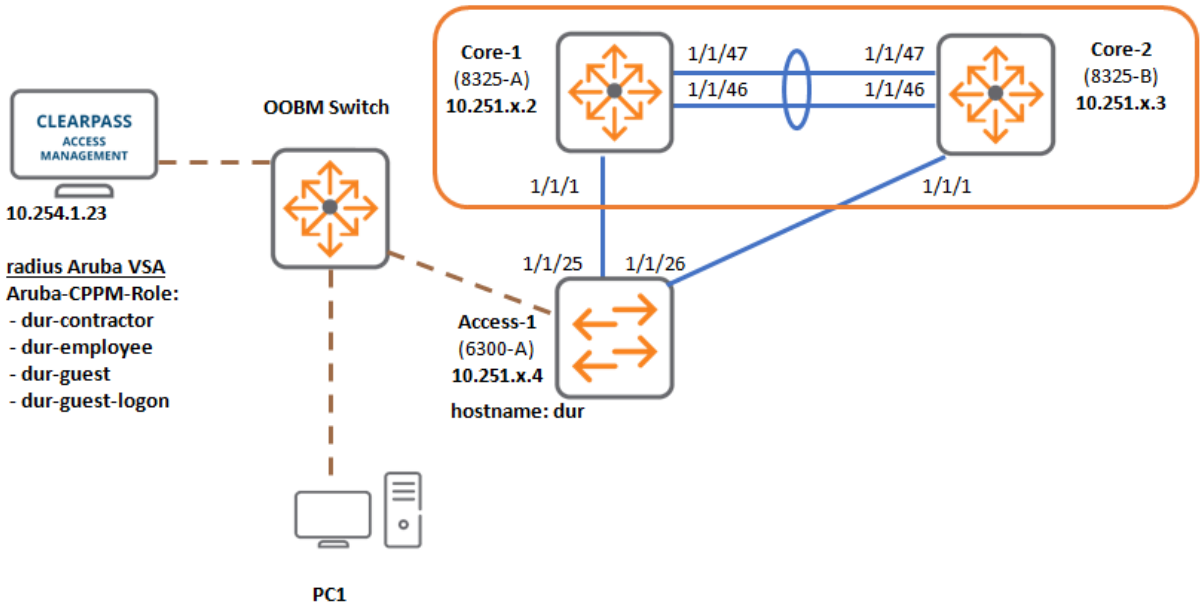

Objectives

There are no configuration steps in this task. You will review the user role configuration that has been prepared on the ClearPass Policy Manager.

Downloadable user roles have been pre-defined for:

- employee
- contractor

Steps

Review the ClearPass Enforcement Profile Configuration

1. Use PC1 to open a session to the ClearPass server.
2. Navigate to **Configuration > Enforcement > Profiles**.
3. Enter **'dur'** in the 'Name' filter and select **'Go'**.

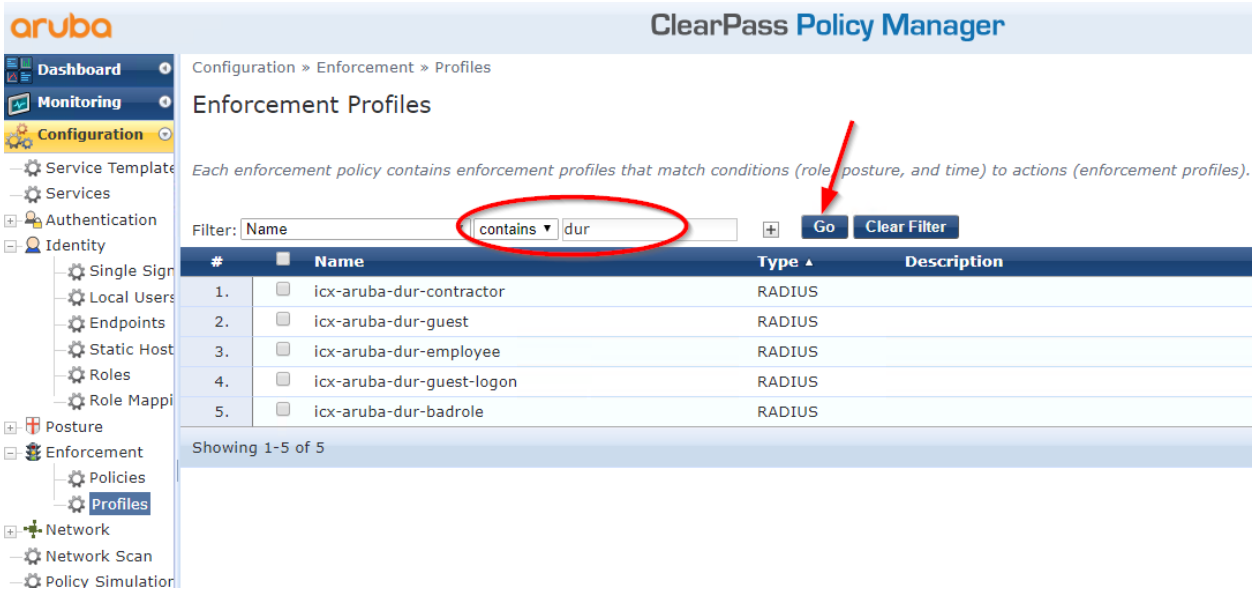

4. Open the **'icx-aruba-dur-employee'**. The RADIUS VSA attribute used here is *Aruba-CPPM-Role*. The data contains the CLI commands that define the complete user role.

> **NOTE:** In the enforcement profile, the user-role name is not important; in the example 'dur' is used. When the profile is downloaded by the switch, the switch will put the provided configuration under a role name that will be based on the Enforcement Policy name, that is "**icx-aruba-dur-employee**" in this example.

Configuration » Enforcement » Profiles » Edit Enforcement Profile - icx-aruba-dur-employee

Enforcement Profiles - icx-aruba-dur-employee

Summary | Profile | Attributes

Profile:

Name:	icx-aruba-dur-employee
Description:	
Type:	RADIUS
Action:	Accept
Device Group List:	-
Product:	Mobility Access Switch

Attributes:

	Type	Name		Value
1.	Radius:Aruba	Aruba-CPPM-Role	=	class ip any match any any any port-access policy employee class ip any port-access role dur associate policy employee vlan access 11

5. Click **Cancel** to close the Employee profile.
6. Click the **icx-aruba-dur-contractor** profile to see the configuration.

Enforcement Profiles - icx-aruba-dur-contractor

Summary | Profile | Attributes

Profile:

Name:	icx-aruba-dur-contractor
Description:	
Type:	RADIUS
Action:	Accept
Device Group List:	-
Product:	Mobility Access Switch

Attributes:

	Type	Name		Value
1.	Radius:Aruba	Aruba-CPPM-Role	=	class ip any 10 match any any any class ip servers 10 match any any 10.0.1.2/255.0.255.255 port-access policy contractor 10 class ip servers action drop 20 class ip any port-access role dur associate policy contractor vlan access 11

IMPORTANT: The 'port-access role' name is again **'dur'**, but as stated before, this name will not be used in the operational status in the switch, since every downloadable role will have its unique name based on the ClearPass enforcement profile name. For this example, it would be "**icx-aruba-dur-contractor**".

7. Click **Cancel** to close the Contractor profile.

Task 3: Testing 802.1X DUR with Employee and Contractor

Diagram

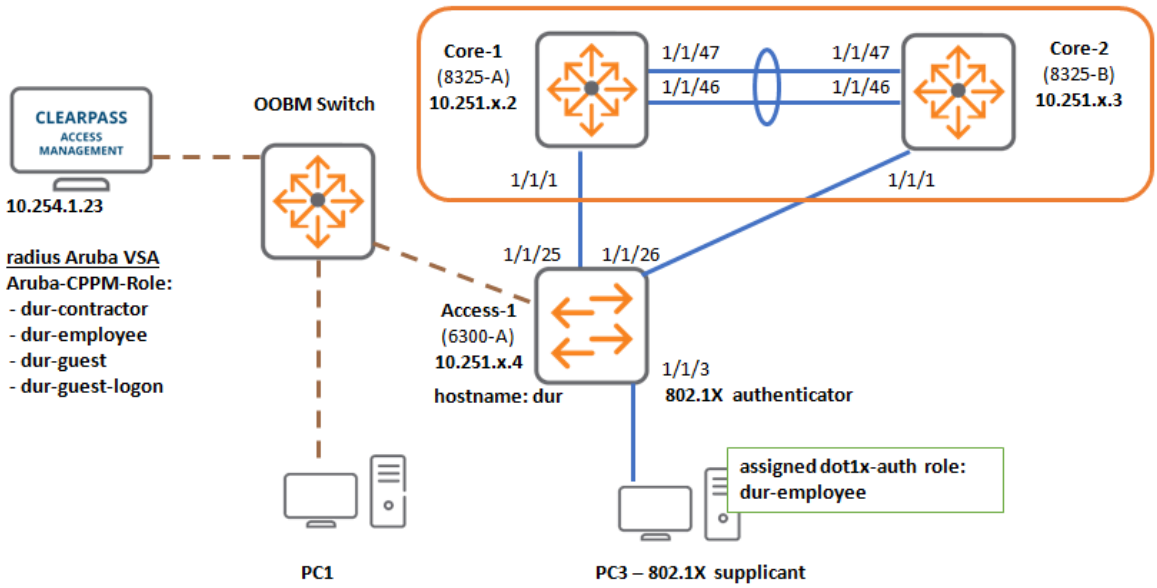

Objectives

In this task, the downloadable user role will be tested using the 802.1X client PC (PC3). Once the user has authenticated, ClearPass will return the name of the downloadable user role and the switch will dynamically collect the details of the role using a REST API connection over HTTPS to ClearPass.

In the lab setup, ClearPass must be made aware that it needs to return the Aruba-CPPM-role, so the hostname on the switch must be updated to include '**dur**' (this is based on the Services that were predefined on the Lab CPPM server).

Steps

Testing Employee Access

Access1

1. Open a terminal session to Access1 and enter the configuration mode.
2. Change the hostname to include the string **'dur'**. This will ensure the correct ClearPass policy rules are used for the Downloadable User Roles. Also, make sure you change “Tx”, where “x” is your table number; for example, T1 for Table 1 or T12 for Table 12.

```
ICX-Tx-Access1(config)# hostname ICX-Tx-Access1-dur
ICX-Tx-Access1-dur(config)#
```

NOTE: The hostname is used as the NAS-identifier in the access-request to ClearPass. ClearPass has been pre-configured to look for the value '**dur**' in the NAS-

identifier to return the downloadable user roles that are needed for this lab.

3. Review the current configuration of port 1/1/3. This port should have 802.1X enabled (mac authentication can also be enabled, but will not be tested on this port).

```
ICX-Tx-Access1-dur(config)# show run int 1/1/3
interface 1/1/3
    no shutdown
    no routing
    vlan access 1
    aaa authentication port-access dot1x authenticator
        eapol-timeout 30
        max-eapol-requests 1
        max-retries 1
        enable
    aaa authentication port-access mac-auth
        enable
    exit
ICX-Tx-Access1-dur(config)#
```

4. On the PC3, connected to Access1 port 1/1/3, enable 802.1X on the 'Lab NIC' and configure the employee credentials (refer to the 802.1X lab for instructions on how to do this):
   - Username: **icx-employee**
   - Password: **aruba123**

Troubleshooting steps in case the authentication fails:

- Verify the TA cert import
- Verify the hostname
- Verify the user credentials
- Verify the switch log (show logging -r -n 20) for any errors

5. Review the authenticated users on port 1/1/3

```
ICX-Tx-Access1-dur(config)# show aaa authentication port-access interface 1/1/3
client-status

Port Access Client Status Details

Client 00:50:56:b1:7a:37, icx-employee
============================
  Session Details
  ---------------
    Port         : 1/1/3
    Session Time : 31273s
```

```
  Authentication Details
  ----------------------
    Status          : dot1x Authenticated
    Auth Precedence : dot1x - Authenticated, mac-auth - Not attempted

  Authorization Details
  ----------------------
    Role   : icx_aruba_dur_employee-3044-7
    Status : Applied

ICX-Tx-Access1-dur(config)#
```

6. Next review the current roles on the switch. Notice how the 'local', 'radius' and 'clearpass' (DUR) roles all co-exist on the same switch.

```
ICX-Tx-Access1-dur(config)# show port-access role
```

7. Next review just the ClearPass downloaded roles.

```
ICX-Tx-Access1-dur(config)# show port-access role clearpass

Role Information:

Name  : icx_aruba_dur_employee-3044-2
Type  : clearpass
Status: Completed
-----------------------------------------------
    Reauthentication Period             :
    Authentication Mode                 :
    Session Timeout                     :
    Client Inactivity Timeout           :
    Description                         :
    Gateway Zone                        :
    UBT Gateway Role                    :
    Access VLAN                         : 11
    Native VLAN                         :
    Allowed Trunk VLANs                 :
    MTU                                 :
    QOS Trust Mode                      :
    PoE Priority                        :
    Captive Portal Profile              :
    Policy                              : employee_icx_aruba_dur_employee-3044-7
```

Q: What is the VLAN assigned to this role?

_____

A: VLAN access 11

8. Next review the port-access policies. Between the locally defined policies, there should also be 1 downloaded policy.

```
ICX-Tx-Access1-dur(config)# show port-access policy

Access Policy Details:
======================
<...>
Policy Name   : employee_icx_aruba_dur_employee-3044-7
Policy Type   : Downloaded
Policy Status : Applied

SEQUENCE    CLASS                         TYPE ACTION
----------- ----------------------------- ---- ---------------------------------
10          any_icx_aruba_dur_employ... ipv4 permit
<...>
```

Q: What is the name of the traffic class in this policy?

_____

A: The IP class name begins with ' any_icx_aruba_dur_employee'.

9. Next review the IP classes in the system.

```
ICX-Tx-Access1-dur(config)# show class ip
Type       Name
           Additional Class Parameters
  Sequence Comment
           Action                          L3 Protocol
           Source IP Address               Source L4 Port(s)
           Destination IP Address          Destination L4 Port(s)
           Additional Entry Parameters
-------------------------------------------------------------------------------
```

In the output, there should be a downloaded IP Class.

```
-------------------------------------------------------------------------------
IPv4       any_icx_aruba_dur_employee-3044-7
        10
           match                           any
           any
           any
-------------------------------------------------------------------------------
```

This demonstrates that the downloaded user roles and their member objects are added to the operational configuration of the switch.

10. Review the running configuration, look for any user roles, policies or classes with 'dur' in the name.

```
ICX-Tx-Access1-dur(config)# show run | include dur
```

Q: Are there any *user-roles* in the configuration?

A: No, the ClearPass Downloadable User Roles are always downloaded on the fly and are not stored in the switch's configuration. This demonstrates that the downloaded user roles and their member objects are added to the operational configuration of the switch, but they are not added to the running or startup configuration of the switch.

Verify the Contractor DUR

Diagram

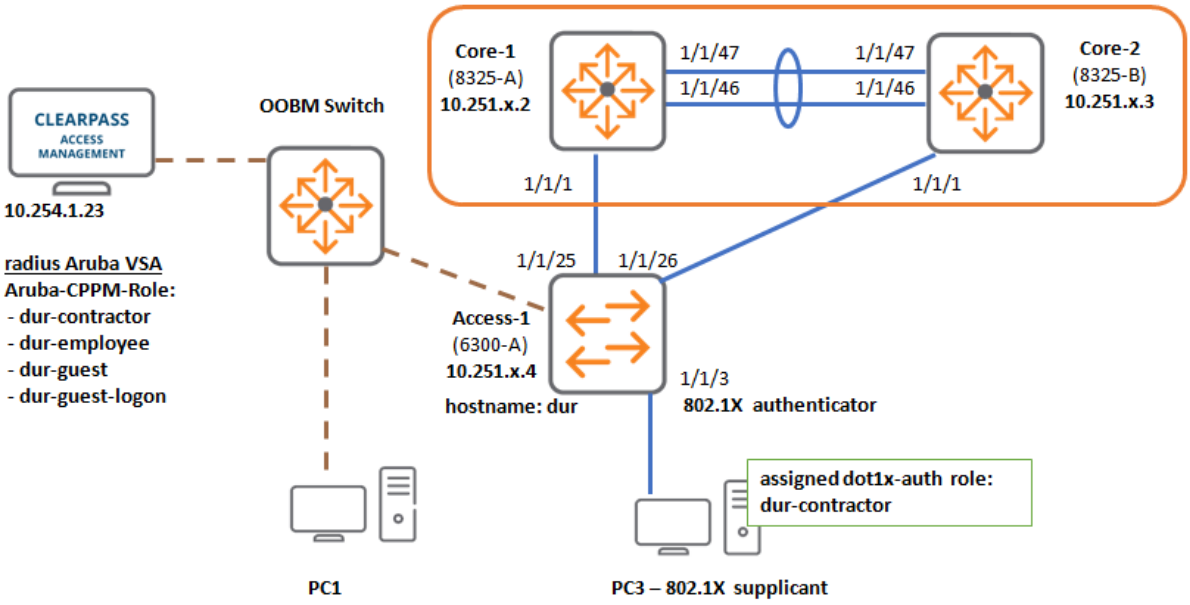

11. On PC3, change the 802.1X credentials to the contractor account:
    - Username: **icx-contractor**
    - Password: **aruba123**

12. On the switch, review the authenticated clients on port 1/1/3.

```
ICX-Tx-Access1-dur(config)# show aaa authentication port-access interface 1/1/3
client-status

Port Access Client Status Details

Client 00:50:56:b1:7a:37, icx-contractor
============================
  Session Details
  ---------------
    Port         : 1/1/3
    Session Time : 32275s

  Authentication Details
  ----------------------
    Status          : dot1x Authenticated
    Auth Precedence : dot1x - Authenticated, mac-auth - Not attempted

  Authorization Details
  ----------------------
    Role   : icx_aruba_dur_contractor-3042-8
    Status : Applied
```

13. Review the downloaded roles.

```
ICX-Tx-Access1-dur(config)# show port-access role clearpass

Role Information:

Name  : icx_aruba_dur_contractor-3042-8
Type  : clearpass
Status: Completed
-----------------------------------------------
    Reauthentication Period             :
    Authentication Mode                 :
    Session Timeout                     :
    Client Inactivity Timeout           :
    Description                         :
    Gateway Zone                        :
    UBT Gateway Role                    :
    Access VLAN                         : 11
    Native VLAN                         :
    Allowed Trunk VLANs                 :
    MTU                                 :
    QOS Trust Mode                      :
    PoE Priority                        :
    Captive Portal Profile              :
    Policy                              : contractor_icx_aruba_dur_contractor-
3042-8
```

Q: Is the employee DUR still available in the role list?

A: No, when the last online user with a downloadable user role goes offline, the role is no longer maintained in the switch memory.

14. Review the policy list.

```
ICX-Tx-Access1-dur(config)# show port-access policy
```

The output should contain:

```
Policy Name   : contractor_icx_aruba_dur_contractor-3042-6
Policy Type   : Downloaded
Policy Status : Applied

SEQUENCE    CLASS                          TYPE ACTION
----------- ------------------------------ ---- ---------------------------------
10          servers_icx_aruba_dur_co... ipv4 drop
20          any_icx_aruba_dur_contra... ipv4 permit
```

This demonstrates how the complete user-role definition can be downloaded on the fly from the ClearPass Policy Manager server.

This concludes the ClearPass downloadable user roles lab.

Optional Task 4: ClearPass DUR Configuration and Troubleshooting

This task is **optional** and can be done if time permits. Check with your instructor.

Introduction

When defining the user-role on ClearPass, the administrator should note that:

- The downloadable user role fully 'self-contained'. In the role it can have its own classes, policies and captive profile objects.
- The downloadable user role cannot use any classes, policies or captive portal profiles that are locally defined in the switch. This means that any referenced object in a DUR must be defined within the ClearPass enforcement profile.
- ClearPass will automatically generate a name for the downloadable user role, this is based on:
  - the name of the "Enforcement Profile"
  - the ClearPass internal enforcement object ID (3xxx)
  - the version number of the "Enforcement Profile". Every time the enforcement profile is saved on the ClearPass server, the version number will increment with 1.
- In the object names within the downloadable user role, any <space> or hyphen (-) is automatically replaced with an underscore (_).
- It is not allowed to have 'comment' lines in the downloadable user role.
- For classes and policies, it is not required to include line numbers, in that case, the line order will be used.
- It is important to order the objects correctly in the ClearPass Enforcement Profile:
  - Classes should be defined before the policy.
  - Policy and Captive Portal profile should be defined before the actual user-role.
- The downloadable user role can refer to a 'ubt' zone. This is covered in the Mobility Controller lab activity.
- The downloadable user role cannot be used to create a ubt zone. That must be done in the switch configuration.

Config Version Information

Every time the enforcement profile is saved in ClearPass, the internal enforcement profile version number is increased. At the next authentication, ClearPass will return the new version number in the DUR name to the switch. This ensures both previous and current version of the same enforcement profile can be online at the same time on the switch.

Migration to the new version occurs based on the re-authentication configuration, allowing for a smooth transition of any ClearPass changes to the actual authenticated devices. When all the devices that were online on the previous version of the DUR have re-authenticated, the old version of the DUR will automatically disappear from the switch.

Steps

Detailed operation of Downloadable User Roles

1. Using PC1, open a session to ClearPass, navigate to Configuration > Enforcement > Profiles. Open the **'icx-aruba-dur-contractor'** profile.

Enforcement Profiles - icx-aruba-dur-contractor

Summary | Profile | Attributes

Profile:

Name:	icx-aruba-dur-contractor
Description:	
Type:	RADIUS
Action:	Accept
Device Group List:	-
Product:	Mobility Access Switch

Attributes:

	Type	Name		Value
1.	Radius:Aruba	Aruba-CPPM-Role	=	class ip any 10 match any any any class ip servers 10 match any any 10.0.1.2/255.0.255.255 port-access policy contractor 10 class ip servers action drop 20 class ip any port-access role dur associate policy contractor vlan access 11

Q: What is the name of the enforcement profile?

_____

A: In this lab, the enforcement profile name is ' **icx-aruba-dur-contractor'**

Q: Why is this name relevant to the switch?

_____

A: The name of the downloadable user role on the switch will be based on the enforcement profile name. Using a meaningful enforcement profile name helps when troubleshooting on the switch.

2. Navigate to Monitoring > Live Monitoring > Access Tracker, open the latest authentication for user '**icx-contractor**'. Check the 'output' tab.

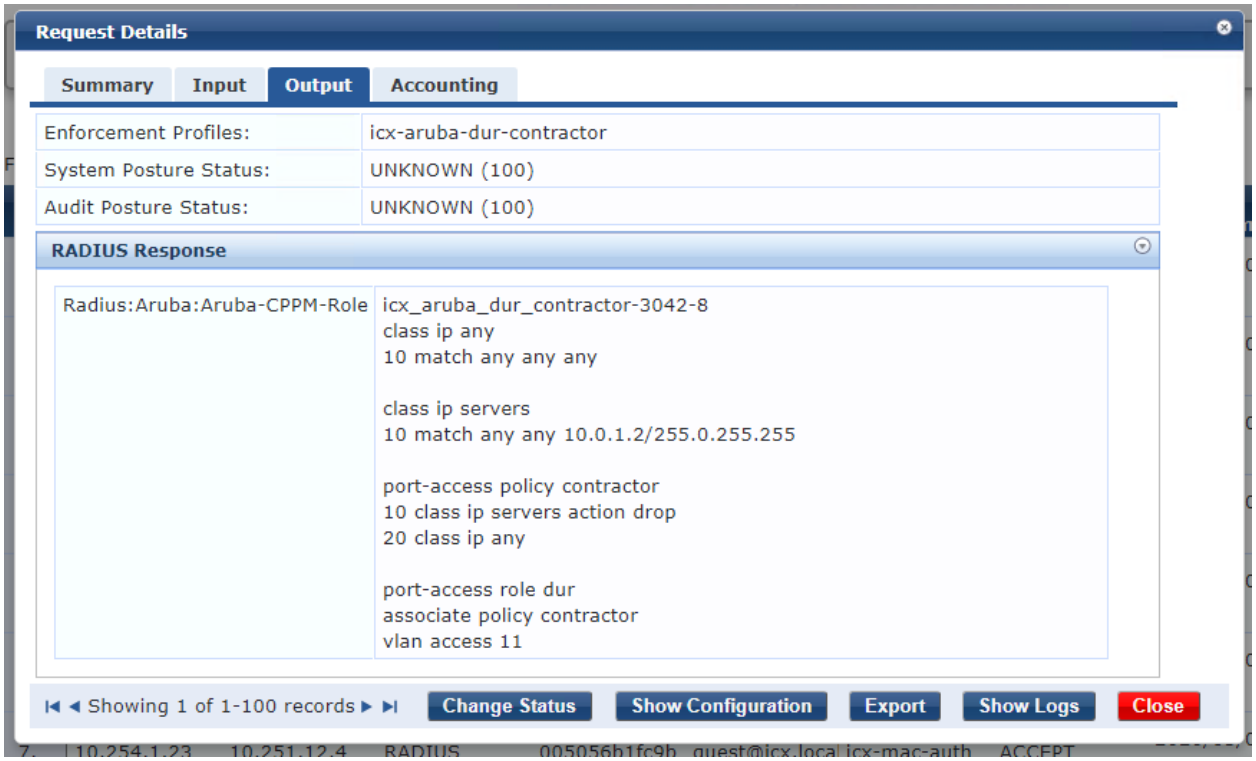

This seems to suggest that the contents of the DUR are sent to the switch as part of the RADIUS Access-Accept.

3. On the Access1, review the name of the role on port 1/1/3.

```
ICX-Tx-Access1-dur(config)# show aaa authentication port-access interface 1/1/3
client-status

Port Access Client Status Details

Client 00:50:56:b1:7a:37, icx-contractor
============================
  Session Details
  ---------------
    Port         : 1/1/3
    Session Time : 71574s

  Authentication Details
  ----------------------
```

```
    Status          : dot1x Authenticated
    Auth Precedence : dot1x - Authenticated, mac-auth - Not attempted

  Authorization Details
  ---------------------
    Role   : icx_aruba_dur_contractor-3042-8
    Status : Applied
```

The output demonstrates that the name of the ClearPass enforcement profile is automatically part of the name of the DUR.

Q: What other elements exist in the DUR name?

A: In this example, the numbers 3042 and 4.

3042: Every enforcement profile in ClearPass has an internal object number, the range starts at 3000. This is the unique identifier of the enforcement profile, while the name makes it easy to recognize it.

8: This is the version number. Every time an enforcement profile is saved, the version number is incremented with 1.

The complete DUR consists of:

- Enforcement Profile name
- Enforcement Profile internal object number
- Enforcement Profile version number

When a ClearPass receives an authentication request that has an Aruba-CPPM-Role in the result, it will not send the detailed commands in the RADIUS request, but only a reference name to the current enforcement profile.

Example RADIUS trace of the Access-Accept sent by ClearPass:

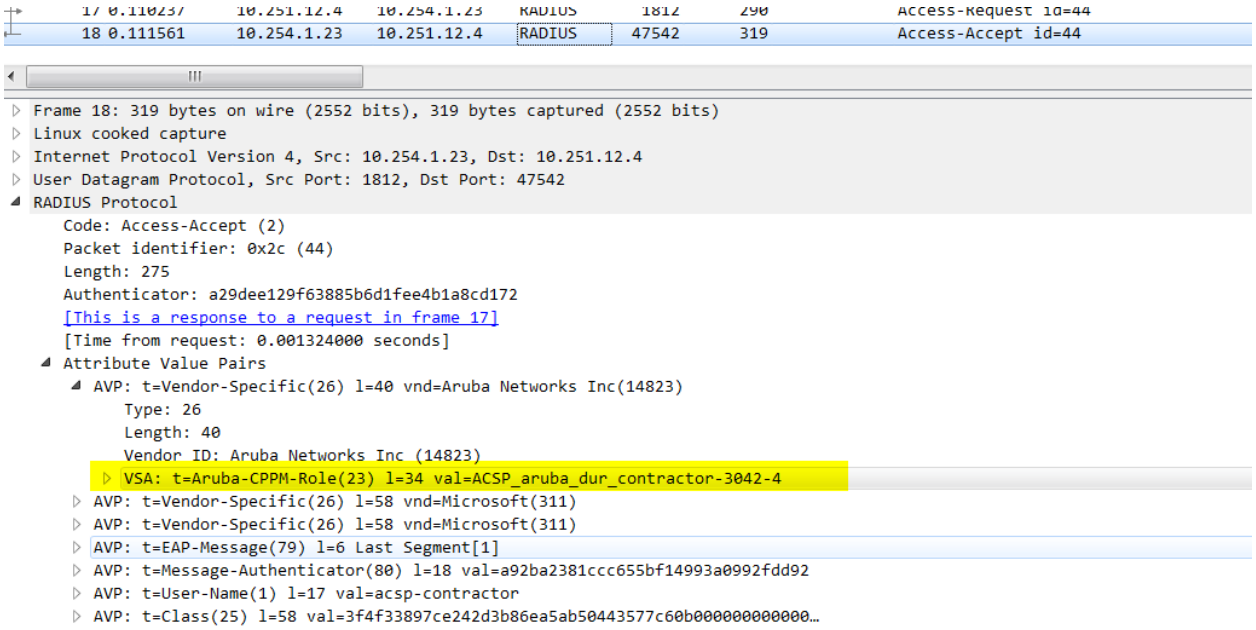

Next, the switch will access the REST API of ClearPass over HTTPS (this required the TA profile installation) and download this enforcement profile contents.

The contents are then installed as a 'clearpass role' in the port-access module.

Any names of referenced objects (classes, policies or captive portal names) are automatically appended with this generated enforcement-profile DUR name.

This ensures no conflicts can occur between roles or even locally defined roles.

4. This is an example of the cli configuration of the Role Contractor:

```
class ip any
 10 match any any any

class ip servers
 10 match any any 10.255.0.200

port-access policy contractor
 10 class ip servers action drop
 20 class ip any

port-access role dur
 associate policy contractor
 vlan access 11
```

NOTE: ClearPass version 6.9 includes a UI to configure the role. The ClearPass UI will then generate these commands automatically.

5. On the switch, there is also a local role 'contractor', that is also referring to a class 'servers'. Thanks to the appended names, there is no conflict. Review the ip classes on the Access1. Notice that there are two server classes:

servers locally define on the switch during the 802.1X Lab Activity.

servers_icx_aruba_... part of the downloadable user role

```
ICX-Tx-Access1-dur(config)# show class ip

<...output omitted...>

------------------------------------------------------------------------------
IPv4       servers
        10
           match                           any
           any
           10.0.1.2/255.0.255.255
------------------------------------------------------------------------------
IPv4       servers_icx_aruba_dur_contractor-3042-8
        10
           match                           any
           any
           10.0.1.2/255.0.255.255
```

This demonstrates that objects in DUR roles will not interfere with each other or with locally defined roles or objects.

Example of Invalid DUR configuration on ClearPass

In the next steps, a DUR that contains an error will be returned to the switch, to demonstrate to effect in the switch.

A **'badrole'** DUR has been prepared on the ClearPass system. It will be returned when the 'icx-badrole' user account connects to the network using 802.1X on the PC3 (connected to Access1).

Diagram

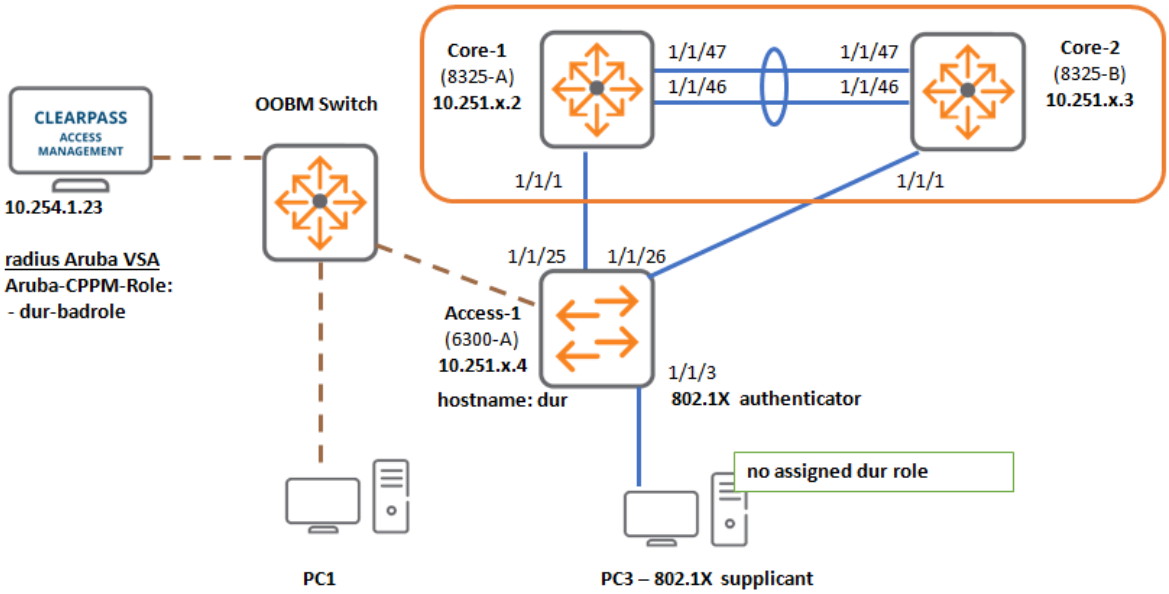

6. On the PC3, connected to Access1, change the 802.1X credentials to:
   - Username: **icx-badrole**
   - Password: **aruba123**

7. After the authentication completes, check the authentication status on the Access1 switch port 1/1/3.

```
ICX-Tx-Access1-dur(config)# show aaa authentication port-access interface 1/1/3
client-status

Port Access Client Status Details

Client 00:50:56:b1:7a:37, icx-badrole
============================
  Session Details
  ---------------
    Port         : 1/1/3
    Session Time : 73386s

  Authentication Details
  ----------------------
    Status          : dot1x Authenticated
    Auth Precedence : dot1x - Authenticated, mac-auth - Not attempted

  Authorization Details
  ----------------------
```

```
    Role   :
    Status : Not Ready
```

Q: What is the status of the role?

A: The status is reported as 'Not Ready', this indicates some error with the role processing.

8. Use PC1 to open ClearPass Access Tracker, check the 'icx-badrole' authentication entry, open the **'output'** tab to see the details of the role.

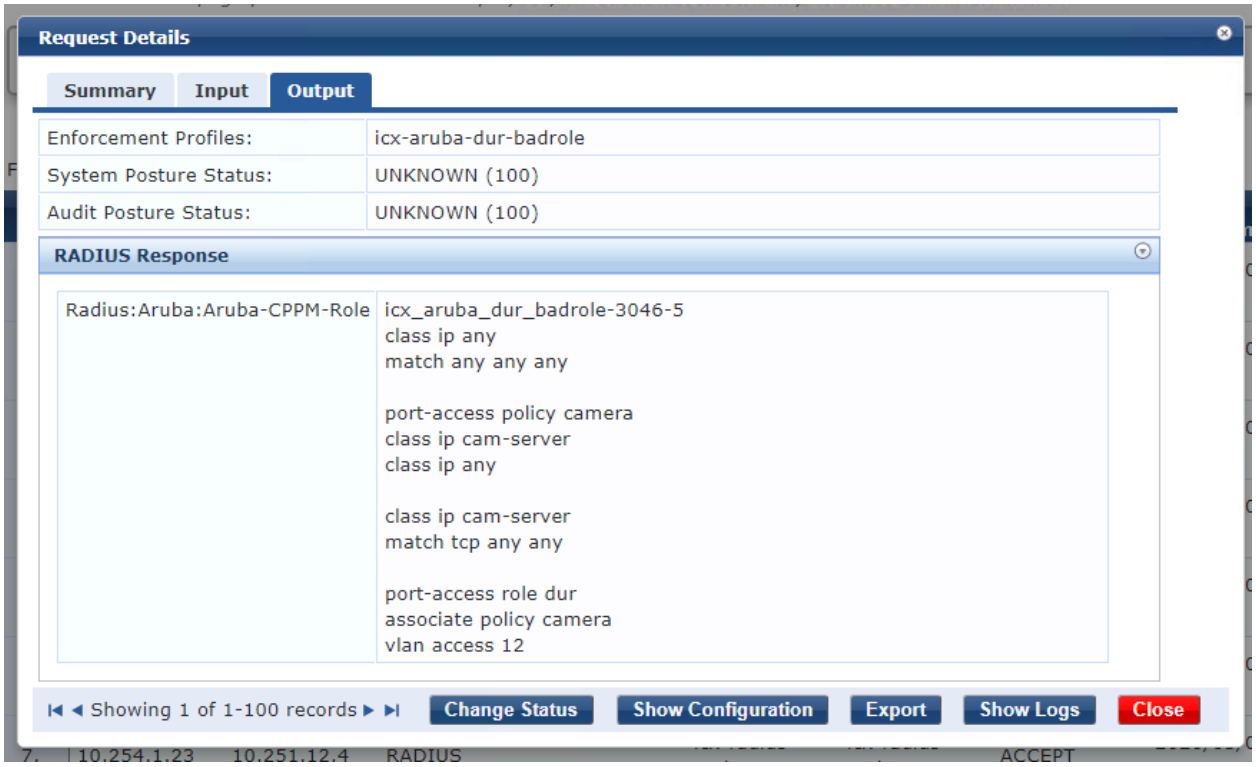

This is the text example of the role:

```
icx_aruba_dur_badrole-3046-5
class ip any
match any any any

port-access policy camera
class ip cam-server
class ip any

class ip cam-server
```

```
match tcp any any

port-access role dur
associate policy camera
vlan access 12
```

Q: What is wrong with this example role?

_____

A: The order of the objects is incorrect. The policy references the class 'cam-server' before it is defined.

9. On the Access1 switch, check the log file. There should be LOG_ERR line that states that the class does not exist.

```
ICX-Tx-Access1-dur(config)# show logging -r -n 10
-------------------------------------------------
Event logs from current boot
-------------------------------------------------
2020-01-28T08:44:06.851047+00:00 ICX-Tx-Access1-dur port-accessd[2956]:
Event|9301|LOG_ERR|MSTR|1|Failed to apply ClearPass role - icx_aruba_dur_badrole-
3046-5 - Error at Line 4 - class ip cam-ser
2020-01-28T08:44:06.850753+00:00 ICX-Tx-Access1-dur port-accessd[2956]:
Event|9301|LOG_ERR|MSTR|1|Failed to apply ClearPass role - icx_aruba_dur_badrole-
3046-5 - Class does not exi
2020-01-28T08:44:06.514464+00:00 ICX-Tx-Access1-dur port-accessd[2956]:
Event|Unknown Event Name CERT_CHAIN_VERIFIED
```

Lab Cleanup

10. Revert the Access1 switch hostname to the base hostname.

```
ICX-Tx-Access1-dur(config)# hostname ICX-Tx-Access1
ICX-Tx-Access1(config)#
```

You've completed Lab 12.1!

Lab 12.2: Integration with Aruba MC- User-Based Tunneling with the MC Firewall

Lab Diagram

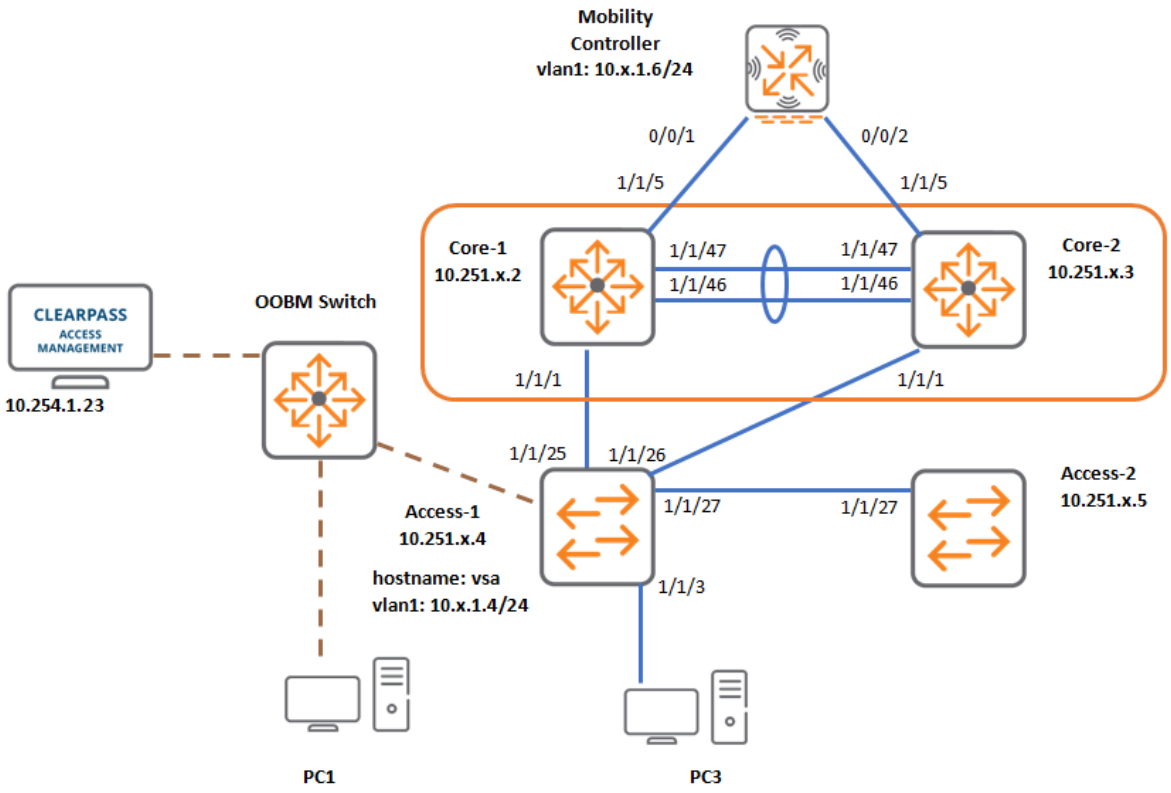

Overview

If you just completed Lab 12.1, the switch configurations are valid and you can start with this lab activity. This lab activity requires at least the completion of Lab11 (MAC Auth).

In this lab activity, the AOS-CX switch will be integrated with the Aruba Mobility Controller. The controller licensing and firewall user-roles will be reviewed.

On the switch, the connection to the controller will be defined using a gateway zone.

Next, the switch user-roles will be updated. The user-role will now point to the MC as the 'gateway zone' and a firewall user-role instruction will be included. Once the user has authenticated, the user and session information will be reviewed on the MC.

Objectives

- Define a connection to the MC for the User-Based Tunneling
- Define user-roles that use the UBT feature
- Verify the state and operation of the tunnels on the Switch
- Verify the state and operation of the tunnels on the Mobility Controller

Task 1: Prepare the Lab Devices

Introduction

This lab requires the completion of the 'Integration with CPPM' lab activity.

In this task, the MC licenses and firewall visibility are verified and adjusted as needed.

Steps

Prepare the MC Licensing and Firewall

MC

1. Open a terminal connection to the MC:
   - Username: **admin**
   - Password: **aruba123**

2. Review that Access Point (AP) and Policy Enforcement Firewall (PEF) licenses have been installed.

```
(ICX-Tx-MC) [mynode] #show license

License Table
-------------
Key                                                 Installed
Expires(Grace period expiry)  Flags  Service Type
---                                                 ---------            ----------
------------------  -----  ------------
uoX+qhM0-rJrBIco4-C8I6V8+1-68KaOfcy-9p081sXP-ork  2019-11-18 15:56:42  Never
E     Access Points: 16
lSQ0Oeuv-hmGIDCLc-tdoNhcLN-vjB0LzVG-nuGwp4d7-G94  2019-11-18 15:56:44  Never
E     Next Generation Policy Enforcement Firewall Module: 16

License Entries: 2

Flags: A - auto-generated; E - enabled; S - Subscription; R - reboot required to
activate; D - Not enabled on license client

Note: Time under 'Installed' for Subscription licenses is the license generation
time.
(ICX-Tx-MC) [mynode] #
```

3. Verify if the PEF license feature has been enabled. This command must be executed at the /mm context level. In the lab setup, it is expected that this feature has been enabled.

```
(ICX-Tx-MC) [mynode] #cd /mm
(ICX-Tx-MC) [mm] #show license-pool-profile-root

License root(/) pool profile
----------------------------
```

```
Parameter             Value
---------             -----
enable PEFNG feature  Enabled
enable RFP feature    Disabled
enable ACR feature    Disabled
enable WebCC feature  Disabled
(ICX-Tx-MC) [mm] #
```

4. **Optional step** in case PEF license feature was not enabled.

```
(ICX-Tx-MC) [mm] #configure terminal
Enter Configuration commands, one per line. End with CNTL/Z

(ICX-Tx-MC) [mm] (config) #license-pool-profile-root
(ICX-Tx-MC) [mm] (License root(/) pool profile) #pefng-licenses-enable
Please ensure to add licenses before enabling feature bit.

(ICX-Tx-MC) ^[mm] (License root(/) pool profile) #write mem

Saving Configuration...

Configuration Saved.
(ICX-Tx-MC) [mm] (License root(/) pool profile) # exit
(ICX-Tx-MC) [mm] (config) # end
```

5. Enable the Deep Packet Inspection (DPI) and the application visibility.

```
(ICX-Tx-MC) [mm] #configure terminal
(ICX-Tx-MC) [mm] (config) #firewall
(ICX-Tx-MC) ^[mm] (config-submode)#dpi
Warning: Application visibility/control is enabled, this change would take effect
after reloading the controller(s) in "/mm"


(ICX-Tx-MC) ^ [mm] (config-submode)# exit
(ICX-Tx-MC) ^ [mm] (config) # firewall-visibility
(ICX-Tx-MC) ^ [mm] (config) # end
(ICX-Tx-MC) ^ [mm] #write memory

Saving Configuration...
```

NOTE: Your Mobility Controller may have been enabled for deep packet inspection already, in this case the controller will log the following message when entering the **dpi** command:

Application visibility/control is already enabled in one or more controllers in "/mm". Changes would take effect after reloading controller(s) in "/mm"

You can move on to the next step.

6. Reload the MC to activate the DPI engine.

```
(ICX-Tx-MC) [mm] #reload
Do you really want to restart the system(y/n): y
```

NOTE: It takes about 7-8 minutes for the controller in the lab to reboot.

Task 2: Aruba MC Integration

Diagram

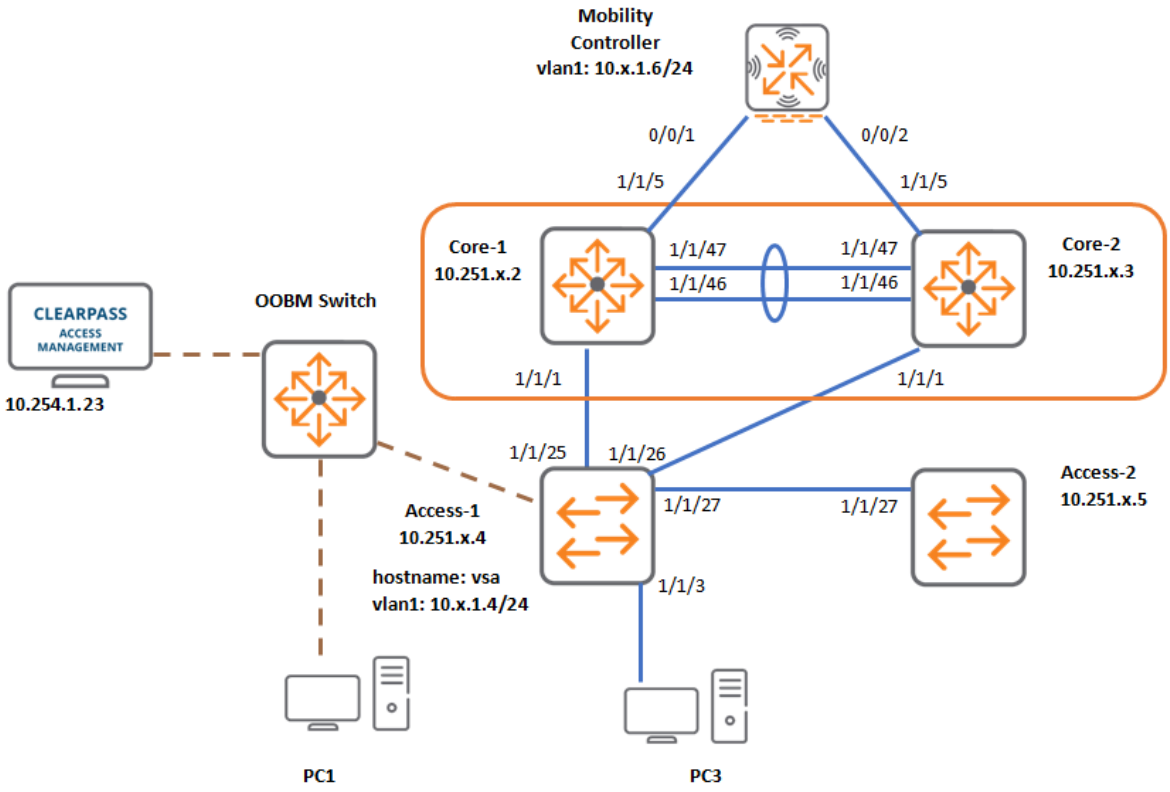

Objectives

In this task, the MC will be configured with the transit VLAN for the user tunnel. In an earlier version of the user based tunneling, the final VLAN that the wired user would be assigned to on the MC had to exist on the switch as well.

To simplify the VLAN administration, the current version only needs a single VLAN to transport wired users via the GRE tunnel to the MC. Once the wired user traffic arrives at the MC, the MC will map the traffic to the final VLAN, so the MC will do a VLAN re-write of the transit VLAN to the final user VLAN.

The switch will be configured with the IP address of the MC as the tunneled-node server. In the switch configuration, this is defined as a 'user-based-tunnel' (ubt) zone.

Steps

Configure tunnel VLAN on MC

1. Open a terminal connection to the MC, enter the configuration mode.
2. At the /mm level, define the VLAN 4000, this is used as the transport VLAN between the switch and the controller.

NOTE: This is the transport VLAN inside the IP GRE tunnel between the switch and the MC, so this VLAN does not need to exist on the intermediate network between the switch and the MC.

```
(ICX-Tx-MC) [mynode] (config) #cd /mm
(ICX-Tx-MC) [mm] (config) #vlan 4000
(ICX-Tx-MC) ^[mm] (config-submode)#write memory

Saving Configuration...

Configuration Saved.
(ICX-Tx-MC) [mm] (config-submode)#end
(ICX-Tx-MC) [mm] #
```

NOTE: It is not strictly required to define the transport VLAN on the Mobility Controller, if all authenticated devices are assigned controller firewall user-roles that include a VLAN instruction (This is the method used in this lab activity).
However, if a device is assigned to a firewall user-role without VLAN instruction, then the device would still be in the assigned transport VLAN.
In this case, the transport VLAN should be defined on the MC and it should be allowed on the Mobility Controller uplink VLAN trunk port(s) to the upstream switch.

Access1 UBT zone configuration

Define the gateway zone and set the MC IP.

Access1

3. On Access1, define an in-band IP on VLAN 1.

```
ICX-Tx-Access1(config)# interface vlan1
ICX-Tx-Access1(config-if-vlan)# no ip dhcp
ICX-Tx-Access1(config-if-vlan)# ip address 10.x.1.4/24
ICX-Tx-Access1(config-if-vlan)# exit
```

4. Verify connectivity to the MC with a ping to 10.x.1.6 (MC IP).

```
ICX-Tx-Access1(config)# do ping 10.x.1.6
PING 10.x.1.6 (10.x.1.6) 100(128) bytes of data.
108 bytes from 10.x.1.6: icmp_seq=1 ttl=64 time=25.8 ms
108 bytes from 10.x.1.6: icmp_seq=2 ttl=64 time=0.357 ms
108 bytes from 10.x.1.6: icmp_seq=3 ttl=64 time=0.362 ms
108 bytes from 10.x.1.6: icmp_seq=4 ttl=64 time=0.333 ms
108 bytes from 10.x.1.6: icmp_seq=5 ttl=64 time=0.350 ms

--- 10.x.1.6 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4052ms
rtt min/avg/max/mdev = 0.333/5.442/25.808/10.183 ms
```

```
ICX-Tx-Access1(config)#
```

5. Set the IP source interface for UBT packets to this VLAN 1 interface.

```
ICX-Tx-Access1(config)# ip source-interface ubt interface vlan1
```

6. Define the UBT zone for the MC. In the current release, only one UBT zone is supported. Make sure to enable the zone.
   - Zone name: **mc**
   - Controller IP: **10.x.1.6**

```
ICX-Tx-Access1(config)# ubt zone mc vrf default
ICX-Tx-Access1(config-ubt-mc)# primary-controller ip 10.x.1.6
ICX-Tx-Access1(config-ubt-mc)# enable
ICX-Tx-Access1(config-ubt-mc)# exit
```

7. Verify the UBT zone configuration.

```
ICX-Tx-Access1(config)# show ubt

Zone Name              : mc
Primary Controller     : 10.x.1.6
Backup Controller      : ---/---
SAC HeartBeat Interval : 1
UAC KeepAlive Interval : 60
VLAN Identifier        : ---/---
VRF Name               : default
Admin State            : Enabled
PAPI Security Key      : Disabled
```

Q: What information seems to be incomplete in the given UBT zone configuration?

_____

A: The switch has not been configured to use a transport VLAN yet for the user tunneled traffic.

8. Define VLAN 4000 as the UBT transport vlan.

```
ICX-Tx-Access1(config)# vlan 4000
ICX-Tx-Access1(config-vlan-4000)# exit
ICX-Tx-Access1(config)# ubt-client-vlan 4000
ICX-Tx-Access1(config)#
```

9. Review the UBT configuration.

```
ICX-Tx-Access1(config)# show ubt

Zone Name               : mc
Primary Controller      : 10.x.1.6
Backup Controller       : ---/---
SAC HeartBeat Interval : 1
UAC KeepAlive Interval : 60
VLAN Identifier         : 4000
VRF Name                : default
Admin State             : Enabled
PAPI Security Key       : Disabled
```

Verify State on MC

10. On Access1, verify that the switch has registered with the MC.

```
ICX-Tx-Access1(config)# show ubt state

Local Master Server (LMS) State:

LMS Type      IP Address    State
-------------------------------------------------
Primary     : 10.x.1.6  ready_for_bootstrap


Switch Anchor Controller (SAC) State:

              IP Address      MAC Address          State
-----------------------------------------------------------------
Active      : 10.x.1.6        20:4c:03:5f:98:02    Registered
```

MC

11. On the MC, verify that the switch is listed as a tunneled-node.

```
(ICX-Tx-MC) [mynode] #show tunneled-node-mgr tunneled-nodes

Tunneled Node Table Entries
---------------------------

Flags: A - Active  Switch Anchor Controller(A-SAC),
       S - Standby Switch Anchor Controller(S-SAC),
       U - Active  User   Anchor Controller(A-UAC),
       X - Standby User   Anchor Controller(S-UAC),
       C - Convert BC & MC into Unicast,


Name                     Tunneled Node Mac  IP Address  Age(d:h:m)  Key   Tunnel
Index  SAC IP Address  S-SAC IP Address  A-Users  S-Users  Flags
----                     -----------------  ----------  ----------  ---   ------
------  --------------  ----------------  -------  -------  -----
ICX-Tx-Access1  88:3a:30:98:30:c0  10.x.1.4   00:00:23    deed  tunnel 9
10.x.1.6       0.0.0.0           0        0        AC
(ICX-Tx-MC) [mynode] #
```

12. On the MC, verify the license usage. The output shows that the switch consumes a license on the controller like an AP.

NOTE: Only the first columns are displayed on this example output.

:

```
(ICX-Tx-MC) [mynode] # show license-usage

License Clients License Usage for pool /
----------------------------------------
Hostname         IP Address  Mac addr          AP    PEF   RF Protect   ACR   ---
-----         ----------  --------          ----  ----  -----------  ----  ----
ICX-Tx-MC  10.x.1.6   20:4c:03:5f:98:02       1     1     0            0     0
TOTAL                                         1     1     0            0     0

Total no. of clients: 1
(ICX-Tx-MC) [mynode] #
```

13. Review the current datapath tunnels.

```
(ICX-Tx-MC) [mm] (config) # show datapath tunnel

...

Datapath Tunnel Table Entries
-----------------------------

Flags: E - Ether encap,  I - Wi-Fi encap,  R - Wired tunnel,  F - IP fragment OK
       W - WEP,  K - TKIP,  A - AESCCM,  G - AESGCM,  M - no mcast src filtering
       S - Single encrypt,  U - Untagged,  X - Tunneled node,  1(cert-id) -
802.1X Term-PEAP
       2(cert-id) - 802.1X Term-TLS,  T - Trusted,  L - No looping, d - Drop
Bcast/Unknown Mcast,
       D - Decrypt tunnel,  a - Reduce ARP packets in the air, e - EAPOL only
       C - Prohibit new calls, P - Permanent, m - Convert multicast, B - Bgw peer
uplink tunnel
       n - Convert RAs to unicast(VLAN Pooling/L3 Mobility enabled), s - Split
tunnel
       V - enforce user vlan(open clients only), x - Striping IP, z - Datazone
       H - Standby (HA-Lite), u - Cluster UAC tunnel, b - Active AAC tunnel, t -
Cluster s-AAC tunnel
       c - IP Compression, g - PAN GlobalProtect Tunnel, w - Tunneled Node
Heartbeat
       B - Cluster A-SAC Mcast, G - Cluster S-SAC Mcast, l - Tunneled Node user
tunnel
       f - Static GRE Tunnels, k- keepalive enabled, Y - Convert BC/MC to Unicast

 #          Source       Destination    Prt  Type  MTU   VLAN       Acls
BSSID          Decaps     Encaps
Heartbeats Flags            EncapKBytes  DecapKBytes
```

```
------  -------------  -------------  ---  ----  ----  ---- ------------------
----  ----------------- ---------- ---------- ---------- --------------- --------
----- -----------
9       10.x.1.6       10.x.1.4       47   deed  1500  0    0   0    0    0    0
88:3a:30:98:30:c0       1854           0       1854 TESw
```

Q: Why is there already a tunnel if there are no clients active? Check the '**w**' flag.

A: This is the keepalive check GRE tunnel. Thanks to this tunnel, all the other tunnels will not have to perform their own keepalive checks. This concept is similar to the Aruba AP GRE data tunnels and keep-alive tunnel.

Task 3: User-Role Configuration on the Switch and the MC

Diagram

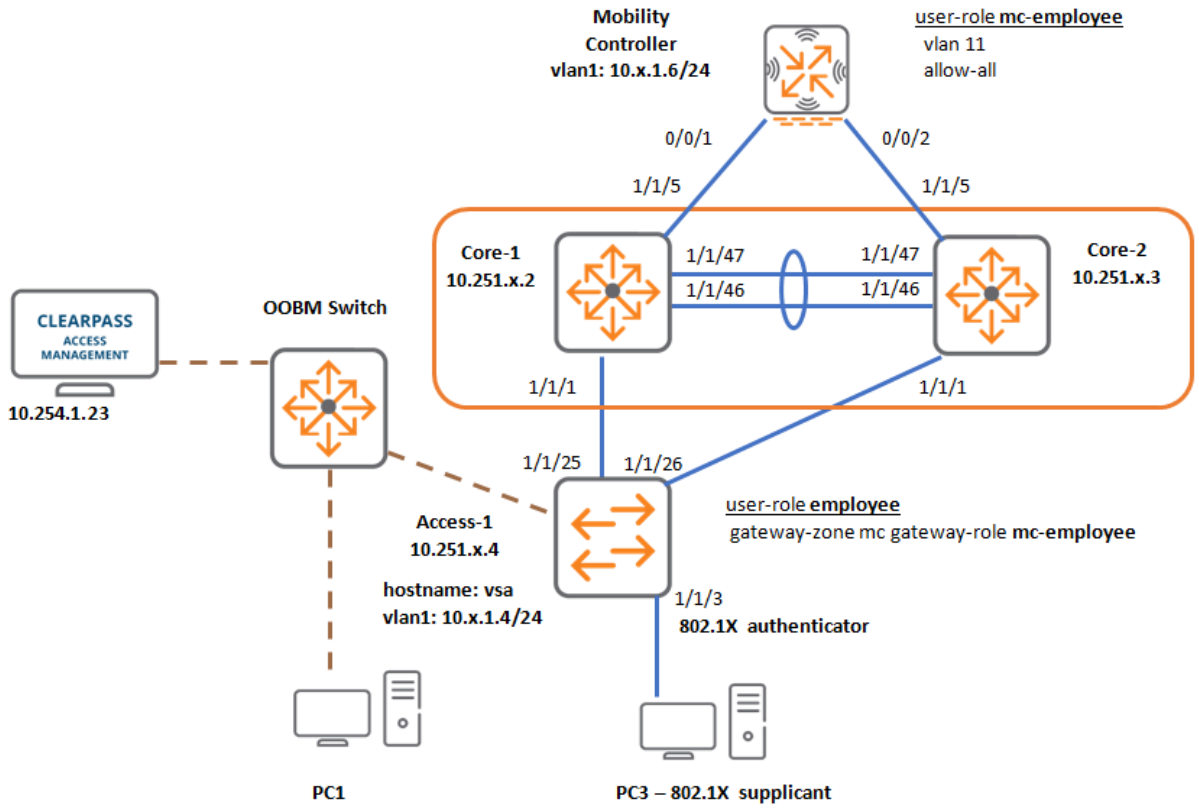

Objectives

In this task, the 'employee 'user-role on the switch will be configured to forward the traffic of the authenticated user to the Aruba Mobility Controller.

First, you will disable the Access1 switch ports 1/1/3 and 1/1/27 so no devices are authenticated.

The hostname will be updated to match the correct ClearPass policies.

On the Aruba Mobility Controller, a new firewall user-role will be defined. In this lab, the role will have 'full access' to the network, but this will still demonstrate that the client traffic visibility is in the firewall.

Steps

Access1: Disable Ports 1/1/3 and 1/1/27, Remove VLAN 11

1. Access the Access1 switch. Change the hostname. Add “**vsa**” at the end. Make sure you change “Tx” to your table informaiton, like T1 for Table 1 and T12 for Table 12.

```
ICX-Tx-Access1(config)# hostname ICX-Tx-Access1-vsa
ICX-Tx-Access1-vsa(config)#
```

> **NOTE:** The hostname is used as the NAS-identifier in the access-request to ClearPass. ClearPass has been pre-configured to look for the value '**vsa**' in the NAS-identifier to return the Aruba VSA Aruba-User-Role attributes that are needed for this lab.

2. Disable the two ports that currently may have active client authentications. These are the ports connected to PC3 (802.1X on 1/1/3) and PC4 (mac-auth via Access2 on 1/1/27).

```
ICX-Tx-Access1-vsa(config)# interface 1/1/3,1/1/27
ICX-Tx-Access1-vsa(config-if-<1/1/3,1/1/27>)# shutdown
ICX-Tx-Access1-vsa(config-if-<1/1/3,1/1/27>)# exit
```

3. Remove VLAN 11. The clients will be transported via the VLAN 4000 to the MC, and the MC will assign them to VLAN 11, so the VLAN 11 does not need to exist on the access switch.

```
ICX-Tx-Access1-vsa(config)# no vlan 11
```

MC: Define a Role for the Employee

4. Open a terminal session on the MC, enter the configuration mode.
5. Define VLAN 11 at the /mm level.

```
(ICX-Tx-MC) [mynode] #cd /mm
(ICX-Tx-MC) [mm] #configure t
Enter Configuration commands, one per line. End with CNTL/Z

(ICX-Tx-MC) [mm] (config) #vlan 11
(ICX-Tx-MC) ^[mm] (config-submode)#exit
```

6. Define a new user-role that will be used for employee. This role has full access to the network. By sending the traffic through the Aruba MC firewall, all traffic can still be inspected and traffic analysis is available.

> **NOTE:** The user roles on the Aruba Mobility Controller are different from the user roles on the Aruba switches, since they support firewall features. They could have the same name as on the switch, but in this lab setup, the name is kept different to make it easier to differentiate the different roles.

```
(ICX-Tx-MC) ^[mm] (config) #user-role mc-employee
(ICX-Tx-MC) ^[mm] (config-submode)#access-list session allowall
(ICX-Tx-MC) ^[mm] (config-submode)#vlan 11
(ICX-Tx-MC) ^[mm] (config-submode)#exit
(ICX-Tx-MC) ^[mm] (config) #write mem
```

Access1 : Update Role Employee to Redirect Traffic to MC Using Secondary Role

7. On the Access1 switch, update the employee role, so traffic is sent to the gateway-zone 'mc' with the secondary role mc-employee.

```
ICX-Tx-Access1-vsa(config)# port-access role employee
ICX-Tx-Access1-vsa(config-pa-role)# gateway-zone zone mc gateway-role mc-employee
ICX-Tx-Access1-vsa(config-pa-role)# exit
```

Task 4: Test Aruba MC Integration

Diagram

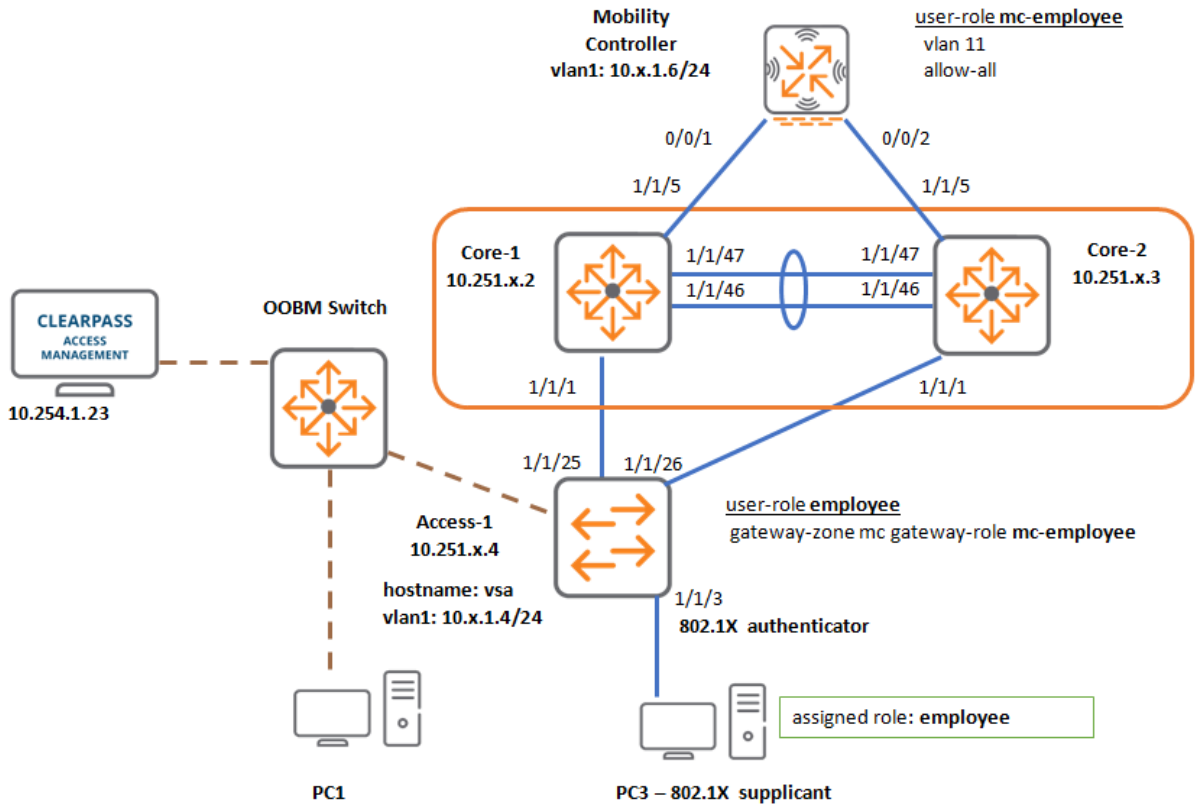

Objectives

In this task, the integration will be tested using the PC3 (connected to Access1 port 1/1/3) with 802.1X authentication.

Steps

Access1

1. On Access1, enable the port 1/1/3.

```
ICX-Tx-Access1-vsa(config)# interface 1/1/3
ICX-Tx-Access1-vsa(config-if)# no shutdown
ICX-Tx-Access1-vsa(config-if)# exit
```

2. On PC3, connected to Access1 port 1/1/3, configure the 802.1X authentication with these credentials:
   - Username: **icx-employee**
   - Password: **aruba123**

Verification the Status on the Switch

3. Verify that the user authentication and role are successfully applied.

```
ICX-Tx-Access1-vsa(config)# show aaa authentication port-access interface 1/1/3
client-status

Port Access Client Status Details

Client 00:50:56:b1:7a:37, icx-employee
==============================
  Session Details
  ---------------
    Port         : 1/1/3
    Session Time : 100s

  Authentication Details
  ----------------------
    Status          : dot1x Authenticated
    Auth Precedence : dot1x - Authenticated, mac-auth - Not attempted

  Authorization Details
  ---------------------
    Role   : employee
    Status : Applied
```

4. Review all the ubt users. This shows the mac and secondary role (mc-employee) that will be send to the MC via the PAPI protocol.

```
ICX-Tx-Access1-vsa(config)# show ubt users all

Displaying All UBT Users for Zone: mc
Downloaded user roles are preceded by *
Port        Mac-Address        Tunnel Status      Secondary-UserRole   Failure
Reason
-------------------------------------------------------------------------------
-------------
1/1/3       00:50:56:b1:7a:37  activated           mc-employee             ---/--
-
```

5. Review the UBT state.

```
ICX-Tx-Access1-vsa(config)# show ubt state

Local Master Server (LMS) State:

LMS Type      IP Address    State
-----------------------------------------------------
Primary     : 10.x.1.6  ready_for_bootstrap
```

```
Switch Anchor Controller (SAC) State:

              IP Address      MAC Address          State
-------------------------------------------------------------------
Active      : 10.x.1.6       20:4c:03:5f:98:02    Registered


User Anchor Controller(UAC): 10.x.1.6

User               Port    State                          Bucket ID  Gre Key
-------------------------------------------------------------------------
00:50:56:b1:7a:37   1/1/3  registered                     252        3
```

NOTE: The bucket map is used in an AOS cluster to distribute clients over the cluster controllers. Each client MAC address is hashed and mapped to a bucket ID based on this hash.

Q: What is the GRE Key for this user?

_____

A: 3, this is based on the switch port ID

Verification Steps on the MC

6. Review the user-tunnel table. This shows the active (GRE) tunnels and the number of users that are active on the tunnel.

```
(ICX-Tx-MC) [mm] (config) #show tunneled-node-mgr user-tunnel-table

Tunnel Info Table Entries
-------------------------

u - Untagged VLAN

Tunnel Id  Tunneled Node  BCMC TO UCast   Key   MTU  Curr Users  VLANs
---------  -------------  -------------  ----  ----  ----------  -----
tunnel 11      10.x.1.4              1     3  1500           1  11
(ICX-Tx-MC) [mm] (config) #
```

Q: What is the Key ID?

_____

A: Key ID is 3, this is the same GRE key ID that was observed on the switch, based on the switch port id.

7. The tunneled-user output shows the active tunneled users.

> **NOTE:** Notice that this output includes the username, client MAC address, the switch (tunneled node mac) and the VLAN mapping for the user. The tunnel # may be different in your setup.

```
(ICX-Tx-MC) [mm] (config) #show tunneled-node-mgr tunneled-users

Tunneled User Table Entries
---------------------------

Flags: U - User Anchor Controller(UAC),
       S - Standby User Anchor Controller(S-UAC),
       T - Tagged VLAN,
       A - Authenticated on Tunneled Node,
       C - Convert BC & MC into Unicast,


User           Tunneled User Mac  Tunneled Node Mac  Vlan    UAC IP Address  Key
Tunnel Index  Flags
----           -----------------  -----------------  ----      --------------  --
-  ------------  -----
icx-employee  00:50:56:b1:7a:37  88:3a:30:98:30:c0  4000(11) 10.x.1.6       3
tunnel 11     UAC
```

Q: What is the meaning of the value **4000(11)** in the VLAN column?

_____

A: The first VLAN, VLAN 4000, is the transport VLAN (inside the GRE tunnel), while the second VLAN is the actual user VLAN on the MC wired uplink port to the core switch.

8. Review the active datapath tunnels.

```
(ICX-Tx-MC) [mm] (config) #show datapath tunnel

...
 #          Source       Destination    Prt  Type  MTU   VLAN       Acls
BSSID          Decaps     Encaps   Heartbeats Flags            EncapKBytes
DecapKBytes
------  --------------  --------------  ---  ----  ----  ---- --------------------
----  ----------------- ---------- ---------- ---------- --------------- --------
----- -----------
9       10.x.1.6       10.x.1.4       47   deed  1500  0    0   0    0    0    0
88:3a:30:98:30:c0       1769          0       1769 TESw
```

```
11      10.x.1.6        10.x.1.4        47   3      1500  0    0    0    0    0    0
88:3a:30:98:30:c0       4309          387           0 EUPRlY
    (ICX-Tx-MC) [mm] (config) #
```

Q: How many tunnels are active now?

_____

A: There are 2 GRE (Protocol 47) tunnels now:

- 1 keep-alive tunnel, type **'deed'**
- 1 user data tunnel, type **'3'** (this was based on the switch port id)

9. Finally, review the regular user table on the MC. This is the user table that would show both wireless and wired (tunneled-node) clients.

```
(ICX-Tx-MC) [mm] (config) #show user
This operation can take a while depending on number of users. Please be patient
....

Users
-----
    IP              MAC             Name           Role          Age(d:h:m)  Auth
VPN link  AP name    Roaming   Essid/Bssid/Phy                      Profile
Forward mode  Type    Host Name  User Type
----------   ------------       ------          ----          ----------  ----
--------  -------    -------   ---------------                     -------
------------  ----    ---------  ---------
10.x.11.41  00:50:56:b1:7a:37  icx-employee  mc-employee  00:00:00    Tunneled-
User-802.1X            10.x.1.4  Tunneled  tunnel 11/88:3a:30:98:30:c0/1/1/3
default-tunneled-user  tunnel        Win 10             TUNNELED USER

User Entries: 1/1
 Curr/Cum Alloc:1/14 Free:0/13 Dyn:1 AllocErr:0 FreeErr:0
```

Application Visibility on MC

Since the wired traffic is now passing through the MC, the network administrator now has visibility in those traffic streams. In the next steps, some traffic will be generated with the wired client PC3 and the resulting analysis will be shown on the MC.

10. On PC3, connected to Access1 port 1/1/3, open a browser and navigate to 10.X.1.1 (the VSX Core switch) . Login using **admin**/**aruba123.**
11. Using the same PC3, open a new browser tab and navigate to 10.X.1.6 (the MC). Login using **admin**/**aruba123**.

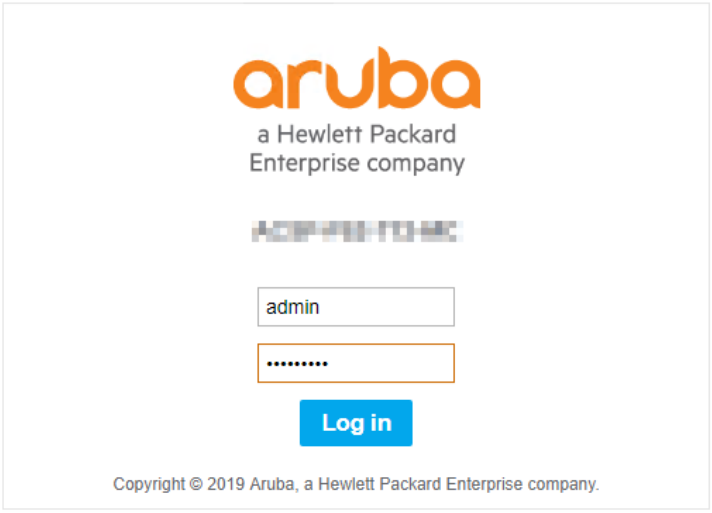

12. Navigate to the **'Dashboard** > **Traffic Analysis'** screen, review the active hosts.

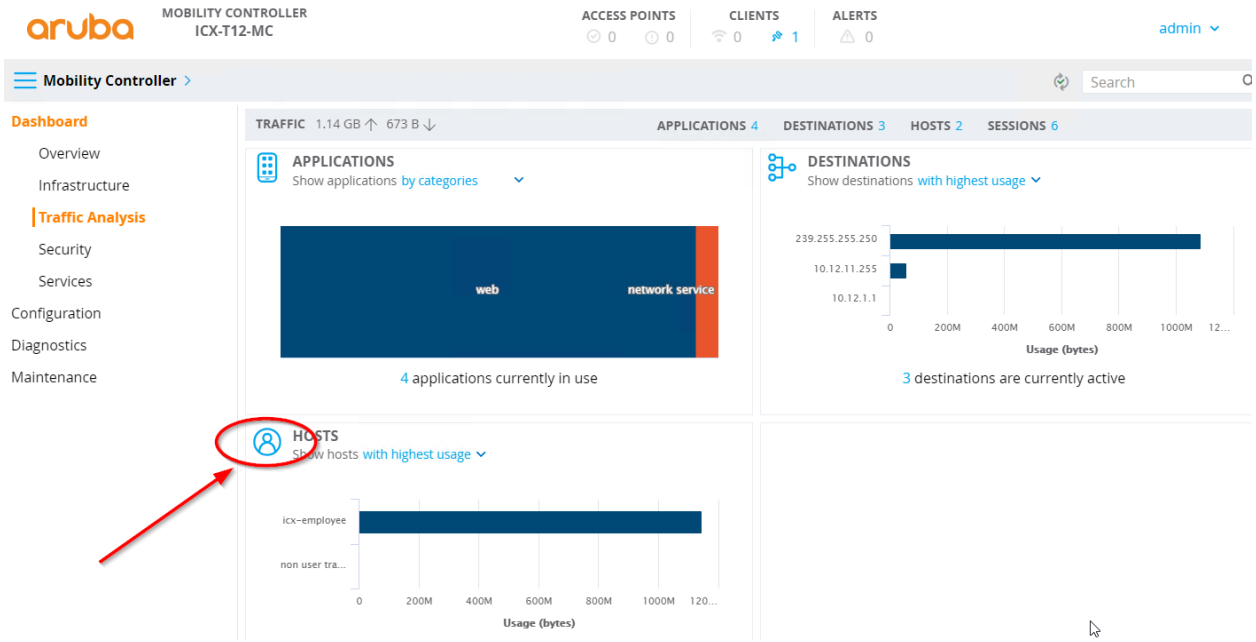

13. Select the **'icx-employee'** host.

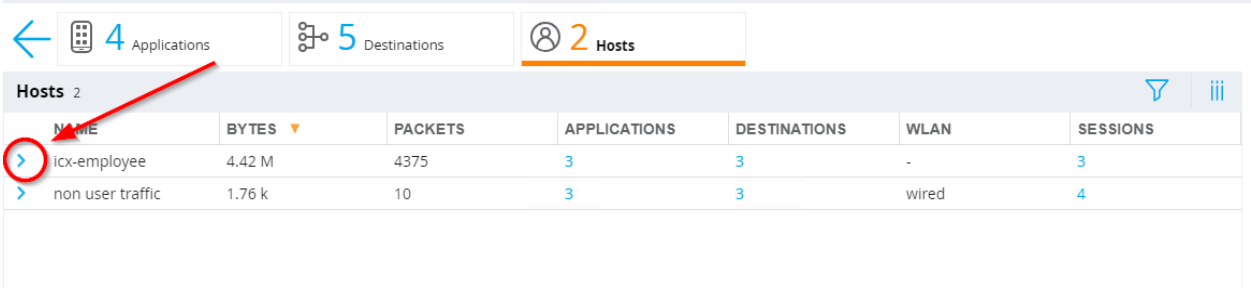

14. Review the details of the traffic sent by this device.

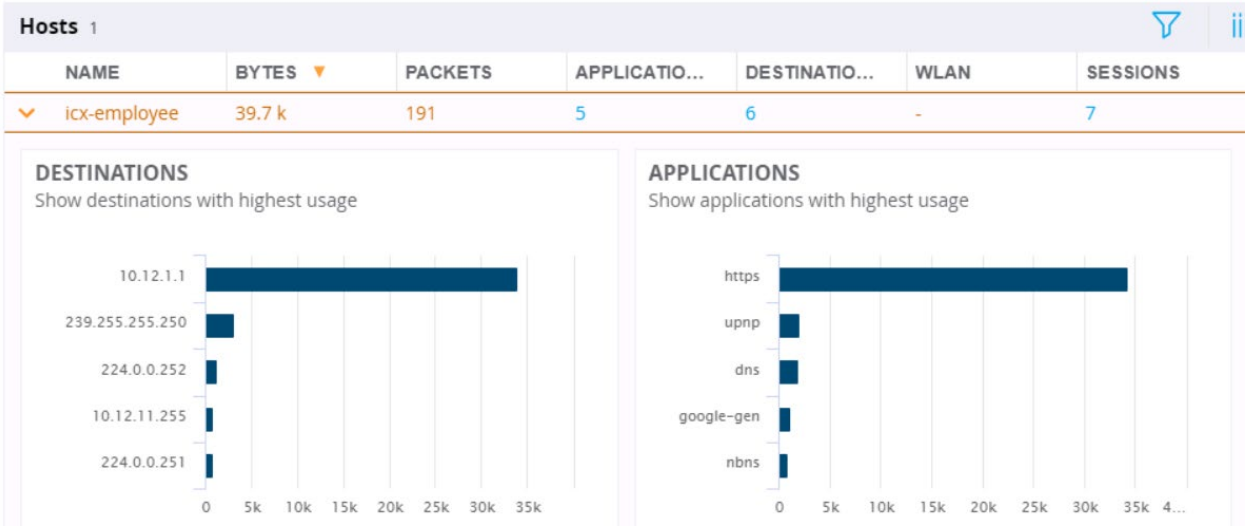

15. In the top bar, click on the 'wired' link under 'clients'. The administrator can clearly see the difference between the wireless and wired clients here.

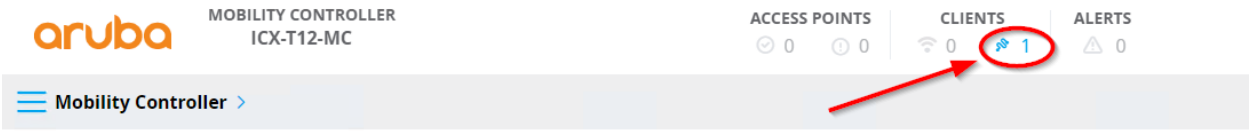

16. On the Aruba MC, more information about the clients can be shown by enabling additional columns.

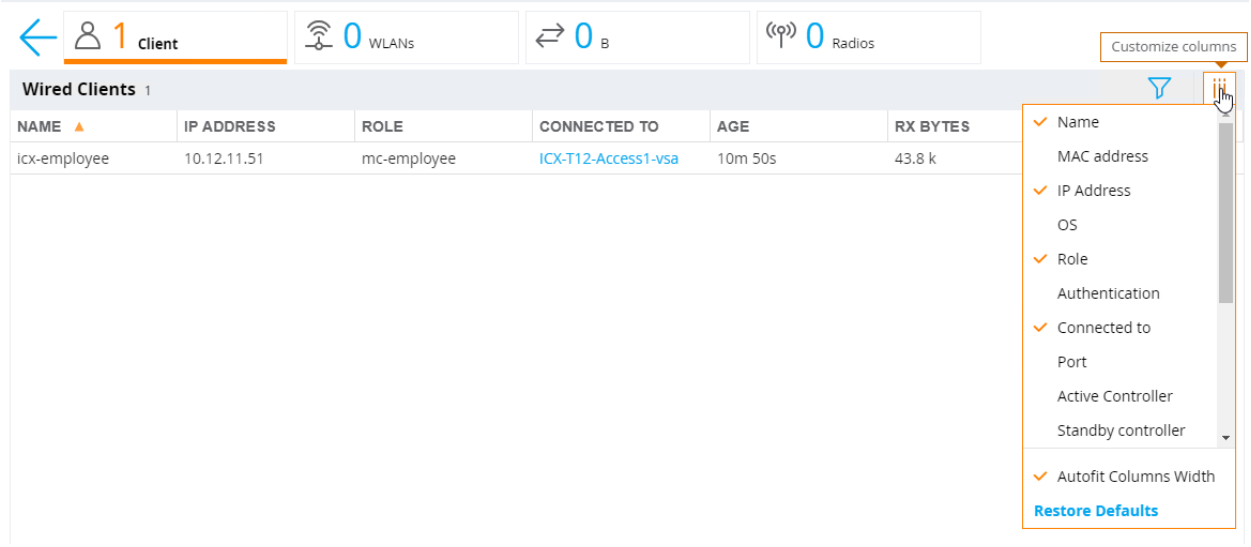

17. This example shows the output with the MAC Address, OS, Authentication and Port enabled:

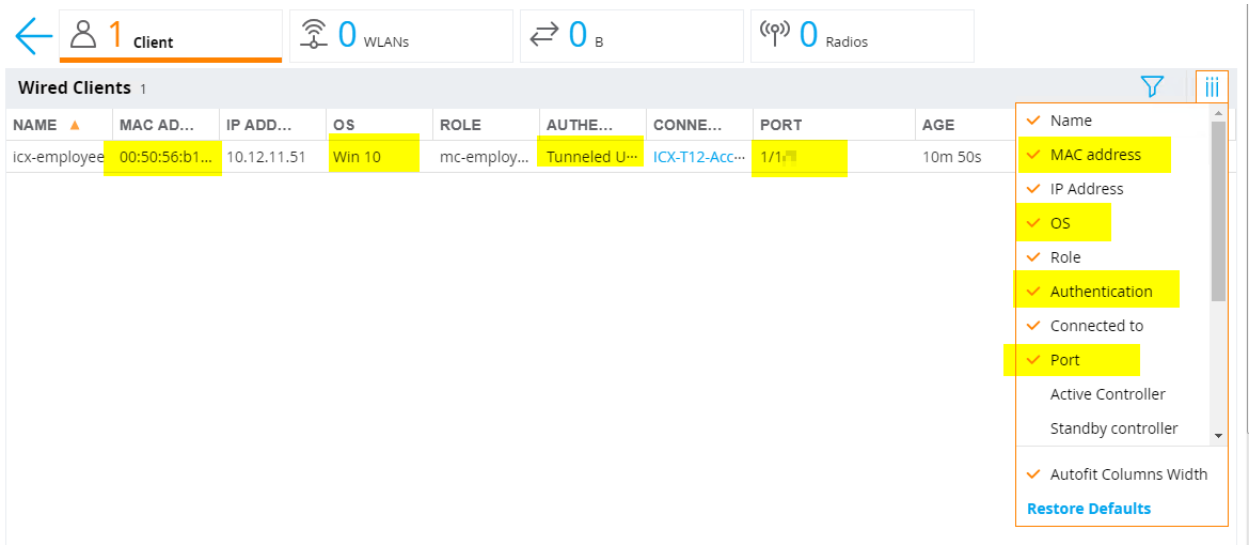

Take note of the IP address of the client device, this will be used in the next step.

Client IP address:

_____

Detailed Firewall Sessions for the Wired Client MC

18. Switch to the terminal of the MC. Review the firewall sessions for the client by using the client IP address as a filter. Replace <client-ip> with the previously noted IP of the client.

In the output, the TCP sessions to the core switch (10.X.1.1 with TCP 443) and the MC (10.x.1.6 with TCP 4343) should be listed.

```
(ICX-T12-MC) [mm] (config) #show user ip <client-ip>
This operation can take a while depending on number of users. Please be patient
....


Datapath Session Table Entries
------------------------------

Flags: F - fast age, S - src NAT, N - dest NAT
       D - deny, R - redirect, Y - no syn
       H - high prio, P - set prio, T - set ToS
       C - client, M - mirror, V - VOIP
       Q - Real-Time Quality analysis
       u - Upstream Real-Time Quality analysis
       I - Deep inspect, U - Locally destined
       E - Media Deep Inspect, G - media signal
       r - Route Nexthop, h - High Value
       A - Application Firewall Inspect
       J - SDWAN Default Probe stats used as fallback
       B - Permanent, O - Openflow
       L - Log, o - Openflow config revision mismatched

Source IP or MAC  Destination IP  Prot SPort DPort Cntr     Prio ToS Age
Destination TAge Packets    Bytes      Flags                   CPU ID
----------------- --------------- ---- ----- ----- -------- ---- --- --- --------
--- ---- ---------- ---------- ---------------         -------
10.x.1.1          10.x.11.51     6    443   9640   0/0     0    0   0   tunnel 10
3c   58         64657      F                          7
10.x.11.51        10.x.1.6       6    9672  4343   1/4101  0    0   0   tunnel 10
38   66         19464      FC                         6
10.x.11.51        239.255.255.250 17   51322 1900   0/0     0    0   1   tunnel 10
18   5          1010       FC                         7
10.x.11.51        10.x.1.6       6    9659  4343   1/4101  0    0   1   tunnel 10
c    2          80         FC                         6
10.x.1.1          10.x.11.51     6    443   9643   0/0     0    0   1   tunnel 10
3c   103        135147                                7
10.x.1.1          10.x.11.51     6    443   9645   0/0     0    0   0   tunnel 10
3c   11         3975       F                          7
10.x.1.6          10.x.11.51     6    4343  9672   0/0     0    0   1   tunnel 10
38   108        106310     F                          6
10.x.11.51        10.x.1.1       6    9645  443    0/0     0    0   0   tunnel 10
3c   10         1969       C                          7
10.x.1.6          10.x.11.51     6    4343  9675   0/0     0    0   1   tunnel 10
37   30         9246       F                          6
<...output omitted...>
```

This demonstrates how the MC can be used to get deep visibility into the traffic of the wired devices.

Optional Task 5: MAC Authentication Example Role for IOT

This task is **optional** and can be done if time permits. Check with your instructor.

Diagram

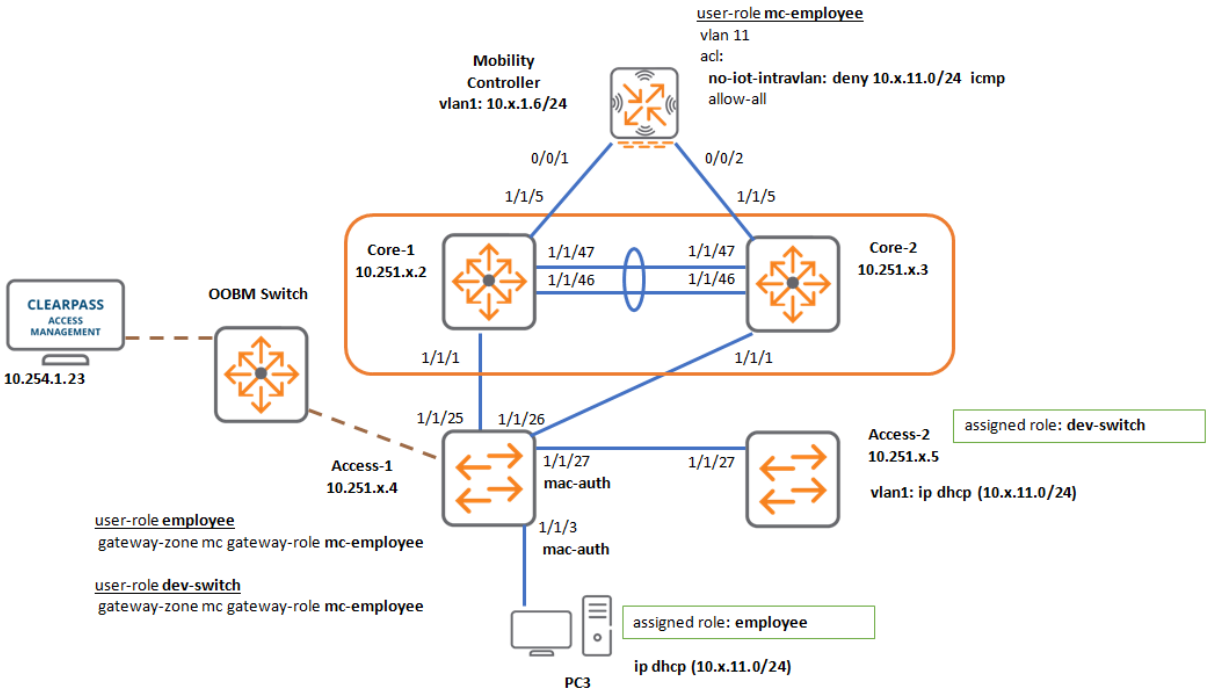

Objectives

In this task, an example IoT use case will be shown. This assumes that a customer wants to assign multiple types of IoT devices to the same VLAN. Since these IoT devices may come from different vendors, the customer is concerned about the security between these IoT devices and would like to apply micro-segmentation. This means that all the traffic, even the intra-VLAN traffic, needs to be firewalled and controlled.

In this lab setup, the PC3 and Access2 will be assigned to the MC in the same VLAN. Initially they will be able to reach each other.

After the Aruba controller firewall user-role is updated, the intra-VLAN traffic will be blocked by the firewall.

To simulate IoT devices, this activity will use MAC-auth.

Steps

PC3: Convert the Lab Setup to MAC-auth

1. On PC3, connected to Access1 port 1/1/3, disable 802.1X on the Lab interface.

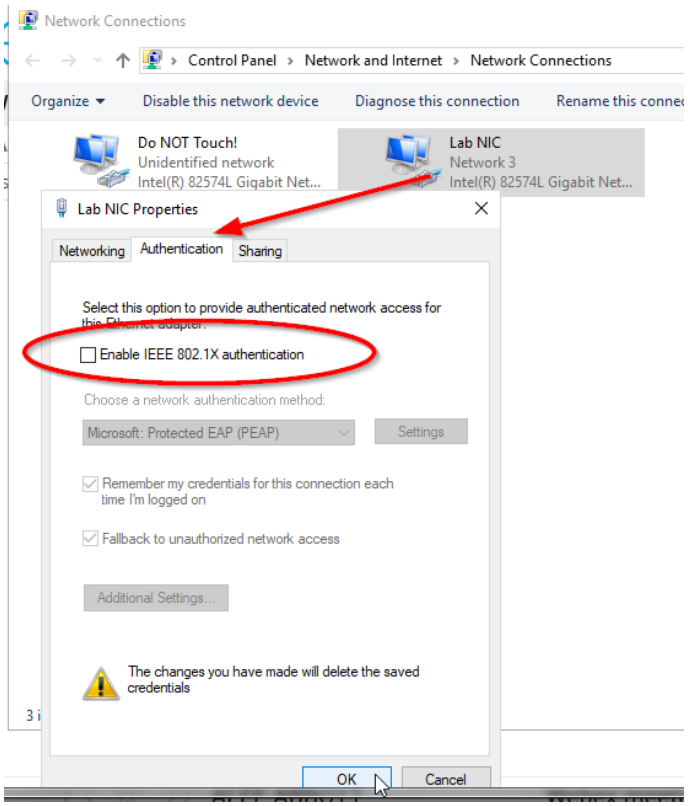

2. On PC3, disable and enable the Lab NIC interface.

Access1

3. On Access1, bounce the switch port 1/1/3 to restart the authentication process. Since the PC does not have 802.1X configured, mac-auth will be performed after about one minute. There is no need to wait for the mac-auth at this moment.

```
ICX-Tx-Access1-vsa(config)# interface 1/1/3
ICX-Tx-Access1-vsa(config-if)# shutdown
ICX-Tx-Access1-vsa(config-if)# no shutdown
ICX-Tx-Access1-vsa(config-if)# exit
```

Prepare connection for Access2

On Access1, the port 1/1/27 connects to Access2. You will update the user-role **'dev-switch'** and assign it also to the **'mc-employee'** role on the MC.

4. Remove the previous role (this ensures no previous settings are maintained), define the role again, and redirect the traffic to the MC.

```
ICX-Tx-Access1-vsa(config)# no port-access role dev-switch
ICX-Tx-Access1-vsa(config)# port-access role dev-switch
ICX-Tx-Access1-vsa(config-pa-role)# gateway-zone zone mc gateway-role mc-employee
ICX-T12-Access1-vsa(config-pa-role)# exit
```

5. Enable port 1/1/27 to Access2 again.

```
ICX-Tx-Access1-vsa(config)# interface 1/1/27
ICX-Tx-Access1-vsa(config-if)# no shutdown
ICX-Tx-Access1-vsa(config-if)# exit
```

Access2

6. Access the terminal of Access2 and enter the configuration mode.
7. Disable the port to PC4 to prevent unnecessary authentications.

```
ICX-Tx-Access2(config)# interface 1/1/4
ICX-Tx-Access2(config-if)# shut
ICX-Tx-Access2(config-if)# exit
```

8. Bounce the uplink port 1/1/27 to trigger VLAN 1 to request a DHCP address

```
ICX-Tx-Access2(config)# interface 1/1/27
ICX-Tx-Access2(config-if)# shutdown
ICX-Tx-Access2(config-if)# no shutdown
ICX-Tx-Access2(config-if)# exit
```

9. After about 30 seconds, the Access2 should have received a DHCP address on its VLAN 1 interface. The assigned IP address may be different in your setup.

```
ICX-Tx-Access2(config)# show ip interface brief
Interface         IP Address             Interface Status
                                           link/admin
vlan1            10.x.11.52/24            up/up
```

Take note of this DHCP IP address:

_____

10. Ping to the Core switch to generate some traffic.

```
ICX-Tx-Access2(config)# do ping 10.x.1.1
PING 10.x.1.1 (10.x.1.1) 100(128) bytes of data.
108 bytes from 10.x.1.1: icmp_seq=1 ttl=64 time=0.277 ms
108 bytes from 10.x.1.1: icmp_seq=2 ttl=64 time=0.280 ms
108 bytes from 10.x.1.1: icmp_seq=3 ttl=64 time=0.327 ms
108 bytes from 10.x.1.1: icmp_seq=4 ttl=64 time=0.270 ms
108 bytes from 10.x.1.1: icmp_seq=5 ttl=64 time=0.222 ms

--- 10.x.1.1 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4100ms
rtt min/avg/max/mdev = 0.222/0.275/0.327/0.034 ms
ICX-Tx-Access2(config)#
```

Verify MAC-auth

Access1

11. Verify the authentication status of port 1/1/27 (to Access2). The ClearPass RADIUS server returns the Aruba VSA User-role '**dev-switch'** for the mac-authentication for the switch.

```
ICX-Tx-Access1-vsa(config)# show aaa authentication port-access interface 1/1/27
client-status

Port Access Client Status Details

Client 88:3a:30:97:b6:00, 883a3097b600
============================
  Session Details
  ---------------
    Port         : 1/1/27
    Session Time : 113s

  Authentication Details
  ----------------------
    Status          : mac-auth Authenticated
    Auth Precedence : dot1x - Not attempted, mac-auth - Authenticated

  Authorization Details
  ----------------------
    Role   : dev-switch
    Status : Applied
```

12. About one minute should have passed by now, verify the authentication status of port 1/1/3 (PC3), it should have received the **'employee'** role.

```
ICX-Tx-Access1-vsa(config)# show aaa authentication port-access interface 1/1/3
client-status

Port Access Client Status Details

Client 00:50:56:b1:7a:37, 005056b17a37
============================
  Session Details
  ---------------
    Port         : 1/1/3
    Session Time : 16s

  Authentication Details
  ----------------------
    Status          : mac-auth Authenticated
    Auth Precedence : dot1x - Unauthenticated, mac-auth - Authenticated

  Authorization Details
  ----------------------
    Role   : employee
    Status : Applied
```

NOTE: Both user-roles have been configured to redirect traffic to the MC using the same **'mc-employee'** gateway role, so they are using the same firewall policy.

MC

13. On the MC, review the active users and their IP Address.

NOTE: Although both clients are using mac-auth at this moment, the MC may still report the original 802.1X client as 'Tunneled-User-802.1X' at this point.

```
(ICX-Tx-MC) [mynode] #show user
This operation can take a while depending on number of users. Please be patient
....

Users
-----
    IP            MAC            Name              Role         Age(d:h:m)  Auth
VPN link  AP name    Roaming   Essid/Bssid/Phy                     Profile
Forward mode  Type    Host Name  User Type
----------   ------------       ------             ----         ----------  ----
--------  -------    -------   ---------------                     -------
------------  ----    ---------  ---------
10.x.11.51  00:50:56:b1:7a:37  00:50:56:b1:7a:37  mc-employee  00:01:20
Tunneled-User-802.1X            10.x.1.4  Tunneled  tunnel
10/88:3a:30:98:30:c0/1/1/1   default-tunneled-user  tunnel        Win 10
TUNNELED USER

10.x.11.52  88:3a:30:97:b6:00  88:3a:30:97:b6:00  mc-employee  00:00:23
Tunneled-User-MAC               10.x.1.4  Tunneled  tunnel
11/88:3a:30:98:30:c0/1/1/27  default-tunneled-user  tunnel
TUNNELED USER

User Entries: 2/2
 Curr/Cum Alloc:2/25 Free:1/23 Dyn:3 AllocErr:0 FreeErr:0
```

One of the listed IP should match the IP of Access2 (MAC starts with 88:3a:30).

14. On PC3, attempt to ping to this IP. This should succeed, since the mc-employee firewall role has an 'allowall' policy.

Example screenshot from Table12. Make sure to use your own lab IP.

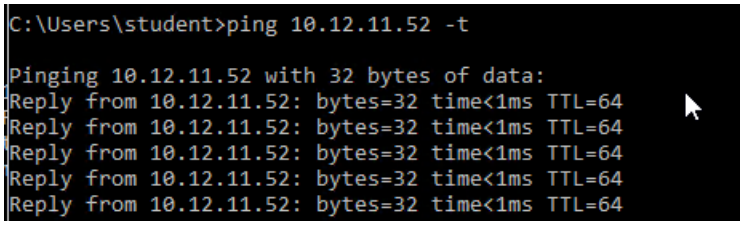

Disable intra-VLAN traffic in the Firewall Role

15. On the MC, define a new session policy that denies ICMP traffic to the local VLAN 11.

```
(ICX-Tx-MC) [mynode] #cd /mm
(ICX-Tx-MC) [mm] #configure terminal
Enter Configuration commands, one per line. End with CNTL/Z

(ICX-Tx-MC) [mm] (config) #ip access-list session no-iot-intravlan
(ICX-Tx-MC) ^[mm] (config-submode)#user network 10.x.11.0 255.255.255.0 icmp echo
deny
(ICX-Tx-MC) ^[mm] (config-submode)#exit
```

16. Review the current mc-employee role. The 'allowall' access-list is currently in position 3.

```
(ICX-Tx-MC) ^[mm] (config) #show rights mc-employee

...

access-list List
----------------
Position  Name                    Type     Location
--------  ----                    ----     --------
1         global-sacl             session
2         apprf-mc-employee-sacl  session
3         allowall                session
```

17. Add the "no-iot-intravlan" access-list to the firewall role mc-employee.

NOTE: Make sure to use 'position 3' in the command, it will insert the access-list before the 'allowall'.

```
(ICX-Tx-MC) ^[mm] (config) #user-role mc-employee
(ICX-Tx-MC) ^[mm] (config-submode)#access-list session no-iot-intravlan position
3
(ICX-Tx-MC) ^[mm] (config-submode)#exit
```

18. Save the configuration to make it effective.

```
(ICX-Tx-MC) ^[mm] (config) #write mem

Saving Configuration...

Configuration Saved.
```

19. Review the updated mc-employee role. The **'no-iot-intravlan'** access-list is now in position 3.

```
(ICX-Tx-MC) [mm] (config) #show rights mc-employee

...

access-list List
----------------
Position  Name                    Type     Location
--------  ----                    ----     --------
1         global-sacl             session
2         apprf-mc-employee-sacl  session
3         no-iot-intravlan        session
4         allowall                session
```

20. On the client PC3, the ping will no longer be possible to Access2.

```
Reply from 10.12.11.52: bytes=32 time<1ms TTL=64
Reply from 10.12.11.52: bytes=32 time<1ms TTL=64
Reply from 10.12.11.52: bytes=32 time<1ms TTL=64
Reply from 10.12.11.52: bytes=32 time<1ms TTL=64
Reply from 10.12.11.52: bytes=32 time<1ms TTL=64
Reply from 10.12.11.52: bytes=32 time<1ms TTL=64
Request timed out.
Request timed out.
Request timed out.
```

This demonstrates how the Aruba dynamic segmentation solution can offer complete network control, even for intra-VLAN traffic.

You have completed Lab 12.2!

Lab 13: Quality of Service

Lab Diagram

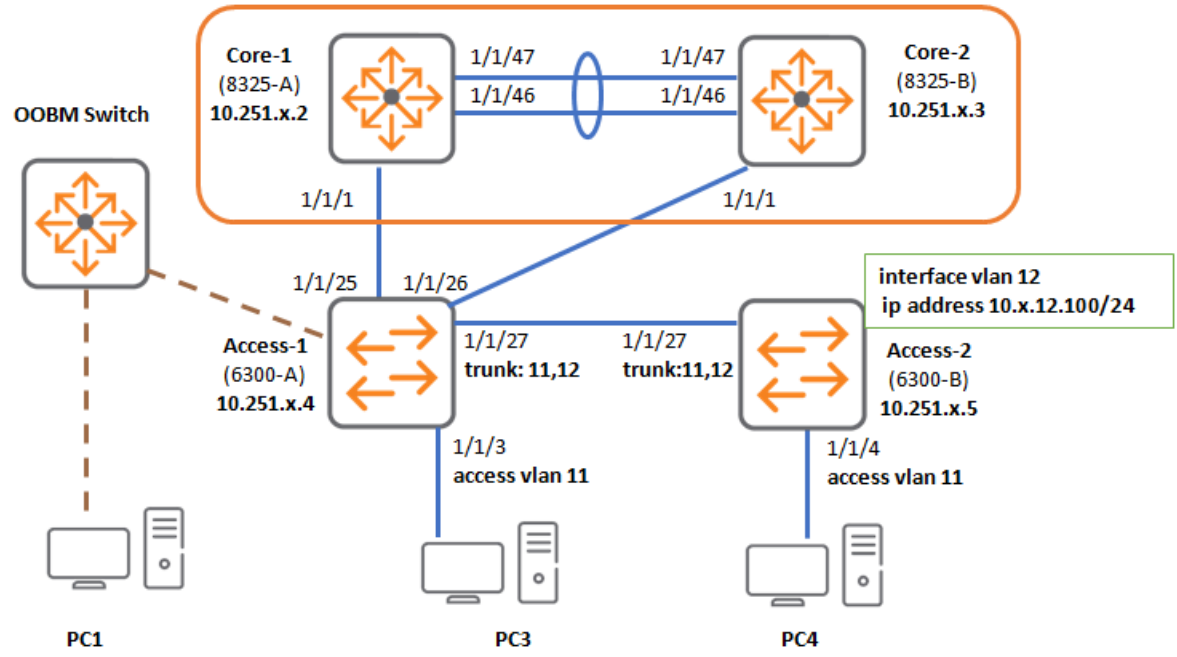

Overview

In this lab activity, the QOS trust and policy configuration will be explained. The first part of the lab will cover the port and global trust options. Next, the LLDP device profile based trust will be demonstrated.

The second part of the lab will explain how classifiers, policies and actions can be used to mark and prioritize specific traffic.

The lab will also show options for the Queue configuration and how to configure an LLDP-MED voice VLAN.

Objectives

- Configure and understand QOS trust modes on the switch
- Understand device profiles
- Configure QOS classifiers
- Understand queue configuration options
- Apply a voice VLAN configuration for LLDP-MED.

Task 1: Prepare the Lab Start Configuration

Diagram

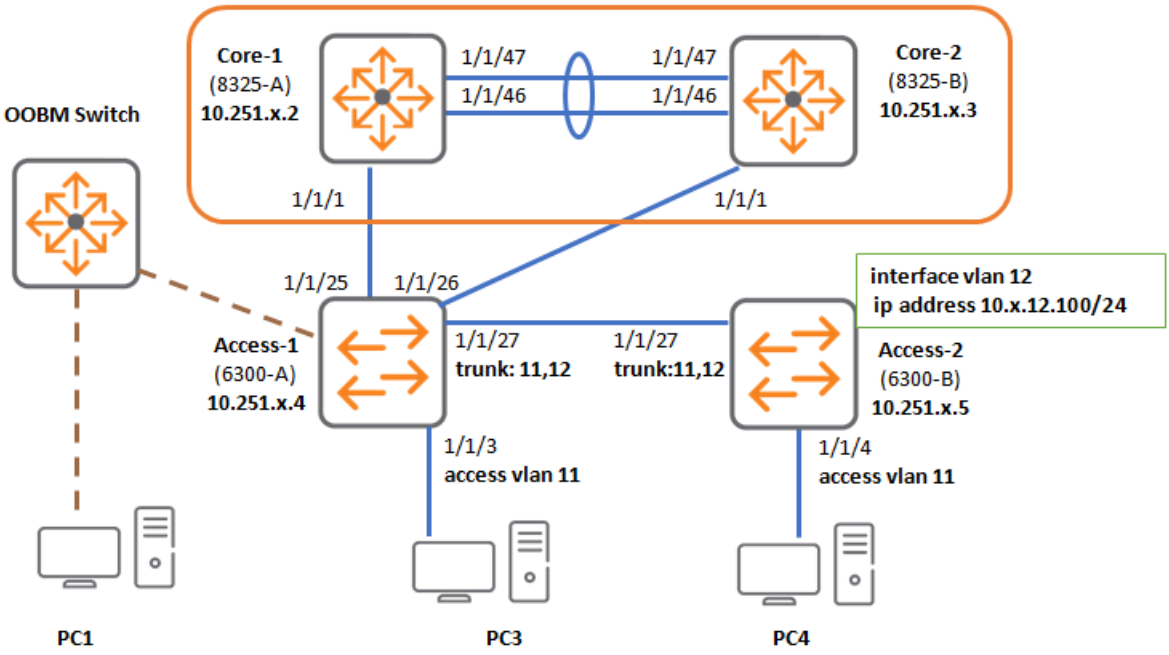

Objectives

- This lab is built on the base VSX topology.
- Make sure to complete these steps to get the base VSX checkpoint configuration on the devices.
- Adjust topology for QOS
- Force traffic from PC4 and Access2 over Access1 to Core

Steps (Required)

1. Open a console connection to the 6300A. Login using **admin**, password **aruba123**.

```
ICX-Tx-Access1# copy checkpoint icx-lab04-vsx running-config
Configuration changes will take time to process, please be patient.
ICX-Tx-Access1#
```

2. Open a console connection to the 6300B. Login using **admin**, password **aruba123**.

```
ICX-Tx-Access2# copy checkpoint icx-lab04-vsx running-config
Configuration changes will take time to process, please be patient.
ICX-Tx-Access2#
```

3. Open a console connection to the 8325A. Login using **admin**, password **aruba123**.

```
ICX-Tx-Core1# copy checkpoint icx-lab04-vsx running-config
Configuration changes will take time to process, please be patient.
ICX-Tx-Core1#
```

4. Open a console connection to the 8325B. Login using **admin**, password **aruba123**.

```
ICX-Tx-Core2# copy checkpoint icx-lab04-vsx running-config
Configuration changes will take time to process, please be patient.
ICX-Tx-Core2#
```

Adjust Topology for the QOS Lab

Access1

5. On Access1, enter the configuration mode.
6. Enable the peer link to Access2, configure it as VLAN trunk using VLAN 11 as native (untagged) VLAN id, and allow VLAN 11 and 12.

VLAN 11 can be used to simulate PC/data traffic.

VLAN 12 is used to simulate Voice traffic (tagged).

```
ICX-Tx-Access1(config)# interface 1/1/27
ICX-Tx-Access1(config-if)# vlan trunk native 11
ICX-Tx-Access1(config-if)# vlan trunk allowed 11,12
ICX-Tx-Access1(config-if)# no shutdown
ICX-Tx-Access1(config-if)# exit
```

Access2

7. Open a terminal to Access2, enter the configuration mode.
8. Disable the uplinks to VSX Core1 and Core2.

```
ICX-Tx-Access2(config)# interface 1/1/25,1/1/26
ICX-Tx-Access2(config-if-<1/1/25,1/1/26>)# shutdown
ICX-Tx-Access2(config-if-<1/1/25,1/1/26>)# exit
```

9. Enable the peer link to Access1, configure it as VLAN trunk, Define an IP on the voice VLAN 12.

```
ICX-Tx-Access2(config)# interface 1/1/27
ICX-Tx-Access2(config-if)# no shutdown
ICX-Tx-Access2(config-if)# vlan trunk native 11
ICX-Tx-Access2(config-if)# vlan trunk allowed 11,12
ICX-Tx-Access2(config-if)# exit
```

10. Define IP address 10.x.12.100 on VLAN 12 and set the default route.

```
ICX-Tx-Access2(config)# interface vlan 12
ICX-Tx-Access2(config-if-vlan)# ip address 10.x.12.100/24
ICX-Tx-Access2(config-if-vlan)# exit
```

```
ICX-Tx-Access2(config)# ip route 0.0.0.0/0 10.x.12.1
```

11. Verify the IP and default route with a ping to a VLAN 11 ip address.

```
ICX-Tx-Access2(config)# do ping 10.x.11.1
PING 10.x.11.1 (10.x.11.1) 100(128) bytes of data.
108 bytes from 10.x.11.1: icmp_seq=1 ttl=64 time=0.207 ms
...
```

12. Enable the port to PC4 connected on Access2. Make it a member of VLAN11 so that PC4 simulates a PC connected to a phone (simulated by Access2) (untagged VLAN 11 PC traffic passing Access2 to Access1).

```
ICX-Tx-Access2(config)# int 1/1/4
ICX-Tx-Access2(config-if)# vlan access 11
ICX-Tx-Access2(config-if)# no shutdown
ICX-Tx-Access2(config-if)# exit
```

13. On PC4, bounce the **'Lab NIC'** to refresh the IP address.

Task 2: Port Classification - Trust Configuration

Objectives

- Access2 is used as test device to send DSCP marked traffic
- Access1 is where the QOS configuration is applied to deal with the incoming marks
- Review the default QOS trust mode on Access1
- Adjust the default global trust mode and verify the result

Steps

Send normal and marked packets from Access2 to PC3 and verify these using Wireshark.

Default QOS Behavior with Marked Traffic

Diagram

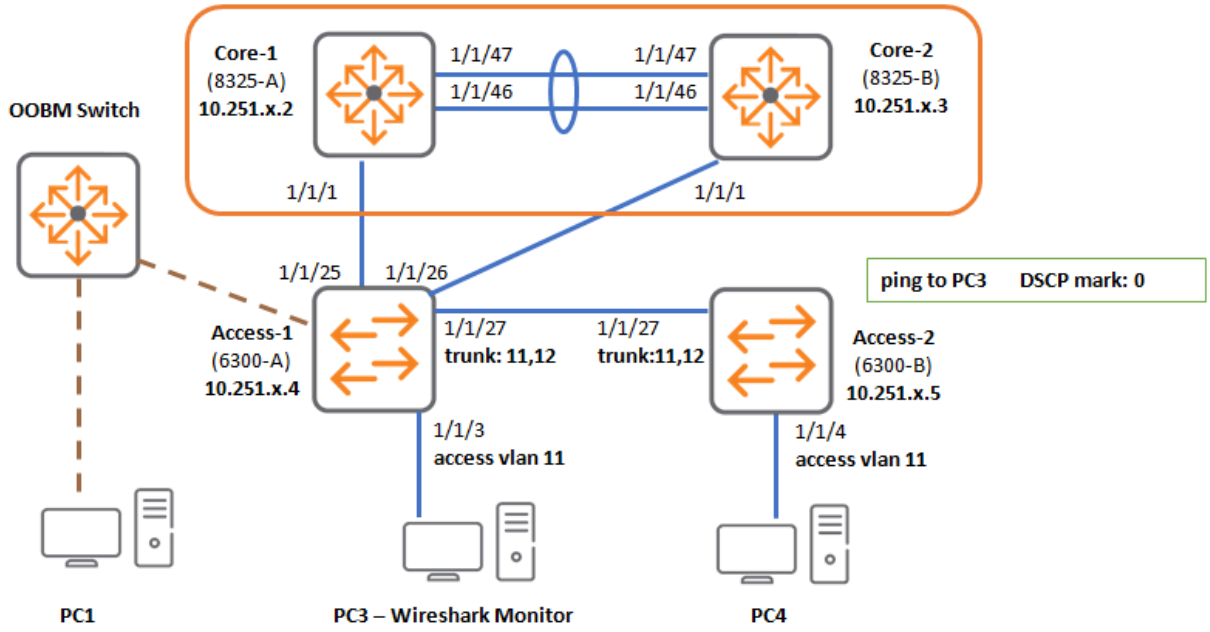

Access1

1. On Access1, enter the configuration mode and clear the interface statistics on ports 1/1/3 (PC) and LAG256 (uplink to the VSX core).

```
ICX-Tx-Access1(config)# do clear interface statistics
```

NOTE: The lab is using this generic clear interface statistics for simplicity, a more selective clear is available per interface, e.g.: *do clear interface 1/1/3 statistics.*

NOTE: In the current release, the 'clear interface statistics' command may need to be executed again to be effective.

Core1

2. On Core1, enter the configuration mode and clear the interface statistics.

```
ICX-Tx-Core1(config)# do clear interface statistics
```

Core2

3. Repeat this on Core2.

```
ICX-Tx-Core2(config)# do clear interface statistics
```

Access1

4. Review the current Queue statistics on e.g. port 1/1/3, this is the port connected to PC3. Most queues should have statistics that are close to 0.

```
ICX-Tx-Access1(config)# show interface 1/1/3 queues
Interface 1/1/3 is up
 Admin state is up
                  Tx Bytes      Tx Packets       Tx Errors
 Q0                     0               0               0
 Q1                     0               0               0
 Q2                     0               0               0
 Q3                     0               0               0
 Q4                     0               0               0
 Q5                     0               0               0
 Q6                     0               0               0
 Q7                   123               1               0
```

NOTE: If the statistics are not close to 0, you should repeat the 'clear interface statistics' command and verify the queue counters again.

NOTE: Clearing the queue statistics makes is easier to troubleshoot and understand what traffic is sent via which queue. Feel free to repeat the clear statistics command in case tests need to be repeated.

NOTE: The switch will automatically send L2 control protocols via Q7, normal user frames will be sent on Q1 by default.

5. Open a connection to PC3 (connected to Access1 port 1/1/3).

6. Check and take note of the IP address (it should have an IP address in VLAN 11), you will need this IP address in the next steps on PC-Access2.

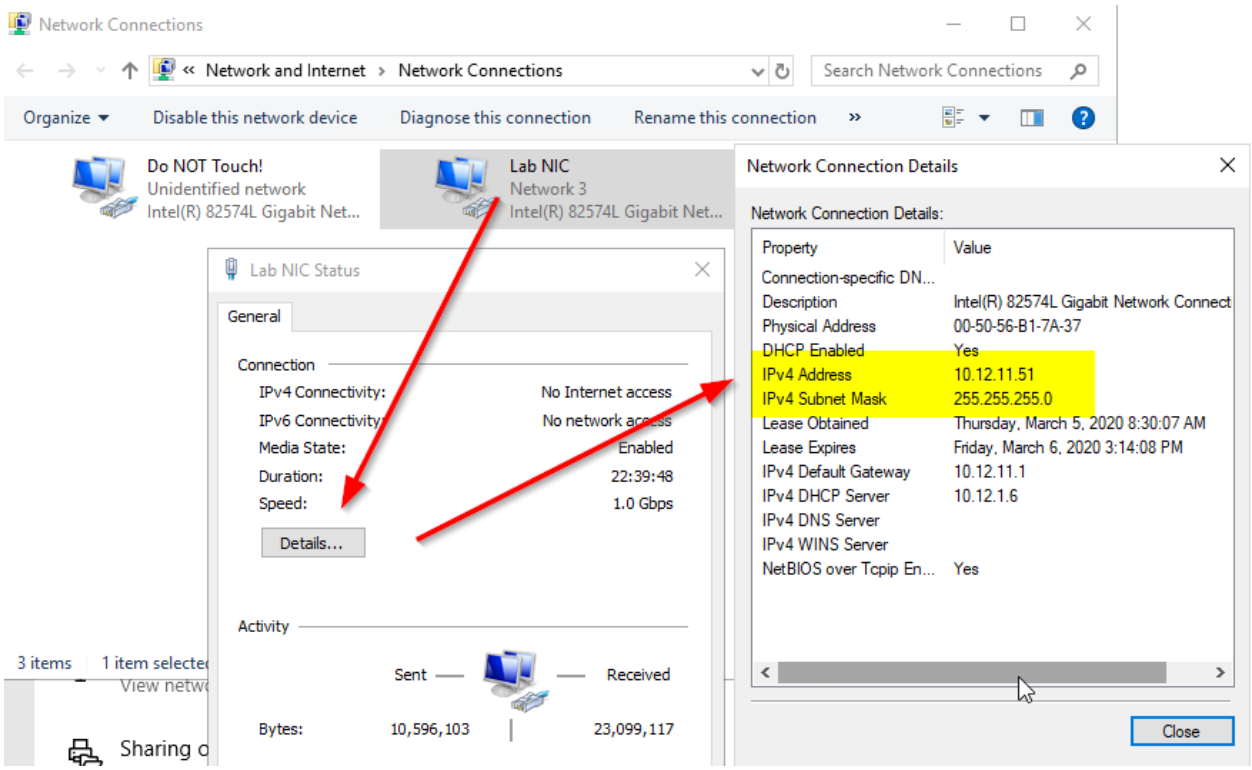

IP address of PC3:

_____

7. Start Wireshark.

8. Double-click the '**Lab NIC**' interface to start a new trace.

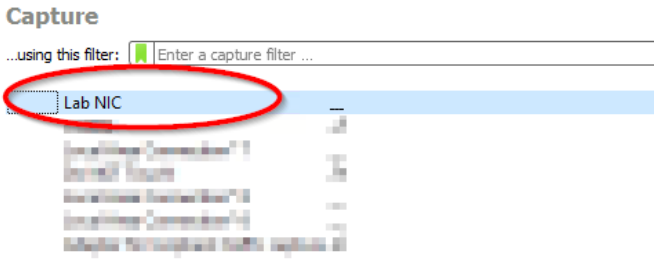

Apply a Display Filter for ICMP Traffic

9. Click in the 'Apply display filter' field.

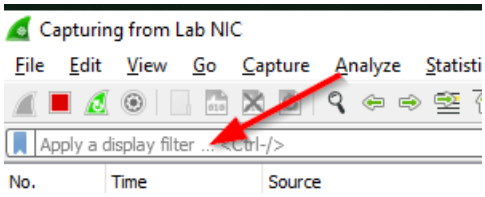

10. Type **'icmp'** and press <**ENTER**>.

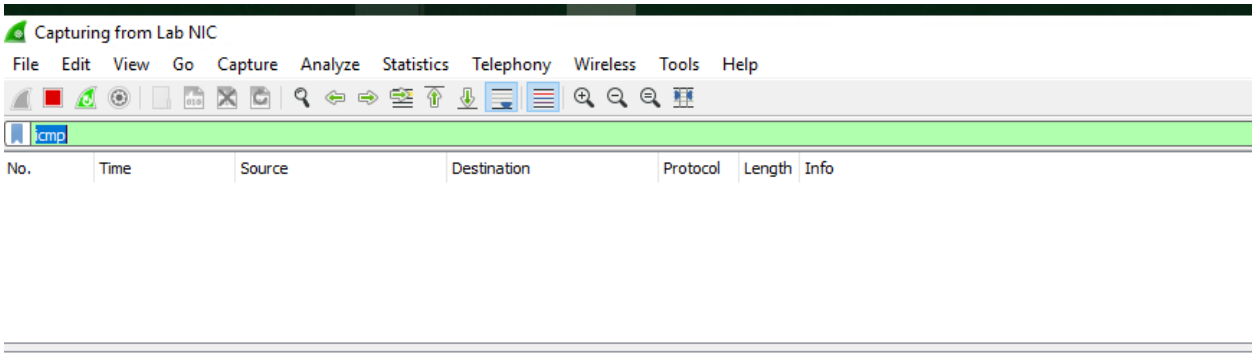

Access2

11. On Access2, perform a ping to PC3 on Access1. Access2 has an IP in VLAN12, so this ping will be routed via the Core VSX. The ping should be successful.

```
ICX-Tx-Access2(config)# do ping 10.x.11.y datagram-size 14000 repetitions 10
```

> **NOTE:** Any ping is fine for this lab, the above ping, using 10 requests of size 14000 generates about 100 outbound packets (due to fragmentation), so it is easy to recognize these test packets in the statistics.

12. On PC3 (connected to Access1), stop the Wireshark capture.

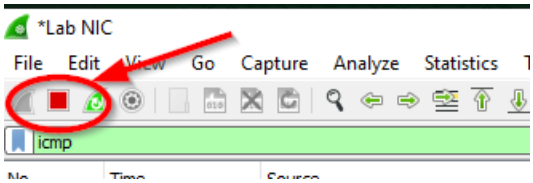

13. Click on packet with source IP of the Access2 (10.x.12.100), open the **'Internet Protocol'** and **'Differentiated Services Codepoint'** sections. Check the DSCP value. Since no configuration was done, this should be Default (0) at this point.

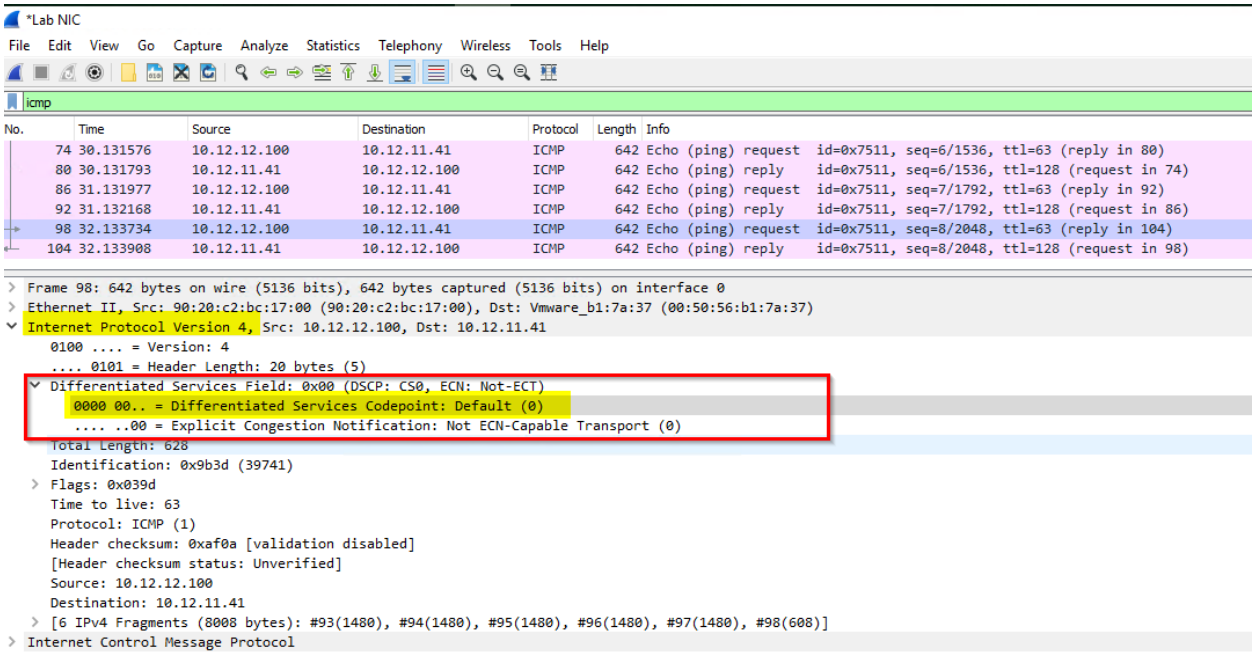

Access1

14. On Access1, review the interface 1/1/3 statistics, this traffic should have been sent out using the 'normal' queue. There should be about 100 packets in the Q1 at this moment.

```
ICX-Tx-Access1(config)# show interface 1/1/3 queues
Interface 1/1/3 is up
 Admin state is up
                  Tx Bytes      Tx Packets       Tx Errors
 Q0                      0               0               0
 Q1                 144420             106               0
 Q2                      0               0               0
 Q3                      0               0               0
 Q4                      0               0               0
 Q5                      0               0               0
```

```
Q6                      0                0                0
Q7                   1599               13                0
```

NOTE: The PC may be generating some other traffic, so the statistics will typically have some variation.

Marked Voice Traffic

Diagram

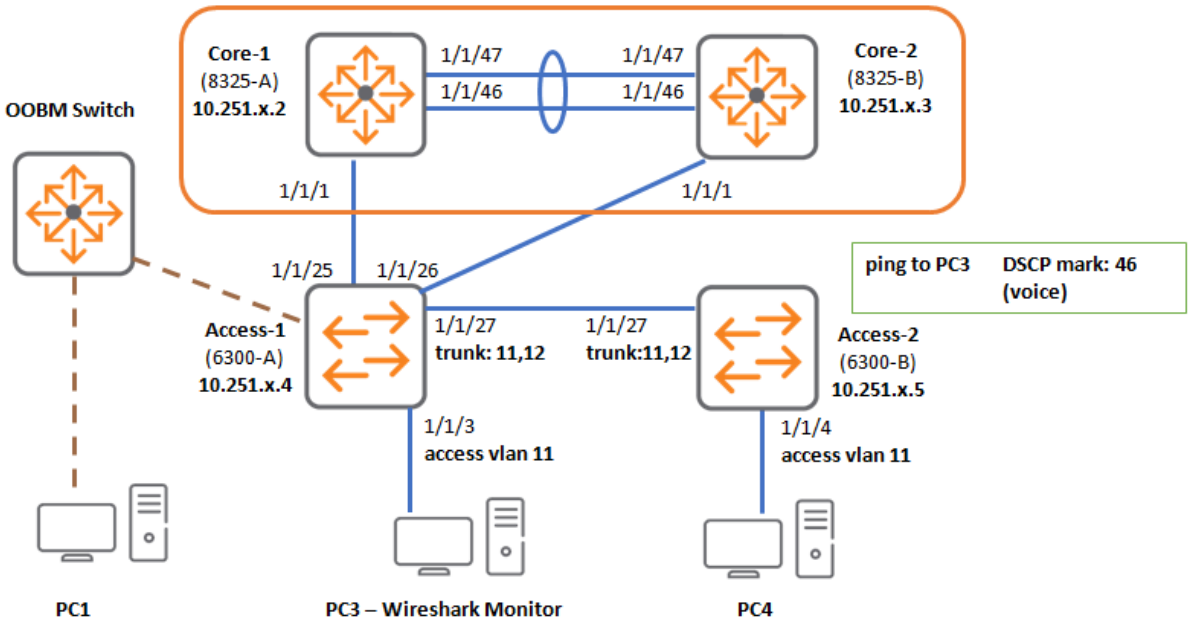

15. On PC3 (connected to Access1), start the Wireshark trace. Click **'Continue without saving'** to start the trace.

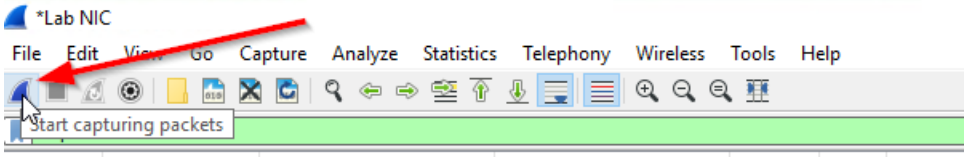

Access2

16. Now use Access2 to send traffic from VLAN 12 marked with voice DSCP (46).

TIP: On an AOS-CX switch, the administrator can enter the TOS value that should be used for the outgoing ICMP packets.

The binary value for DSCP 46 is '101110'.

The IP TOS field has 2 extra bits at the end, so the complete value is '**10111000**'

Converted into decimal, this results in value 184.

```
ICX-Tx-Access2(config)# do ping 10.x.11.y datagram-size 17000 repetitions 8 tos
184
```

17. On the receiving PC, verify the received marked traffic. In Wireshark, stop the trace and verify the incoming DSCP value of the ICMP request (make sure to select a request, not the reply).

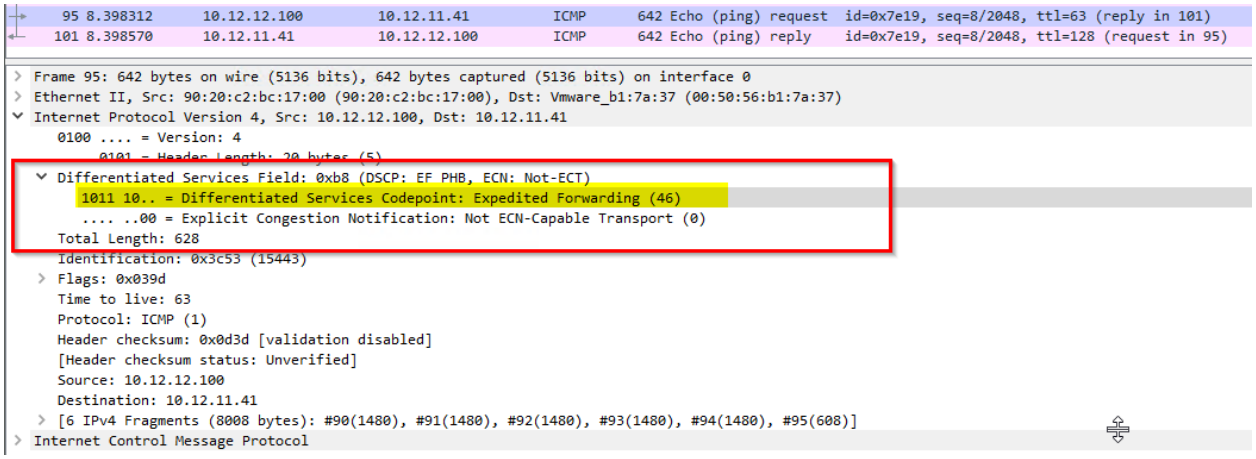

Access1

18. On Access1, Review statistics on the queues, it should still be normal traffic, even when marked with the DSCP value for voice.

```
ICX-Tx-Access1(config)# show interface 1/1/3 queues
Interface 1/1/3 is up
 Admin state is up
                  Tx Bytes      Tx Packets       Tx Errors
 Q0                      0               0               0
 Q1                 283510             238               0
 Q2                      0               0               0
 Q3                      0               0               0
 Q4                      0               0               0
 Q5                      0               0               0
 Q6                      0               0               0
 Q7                   1860              14               0
```

Enable Global DSCP Trust

Diagram

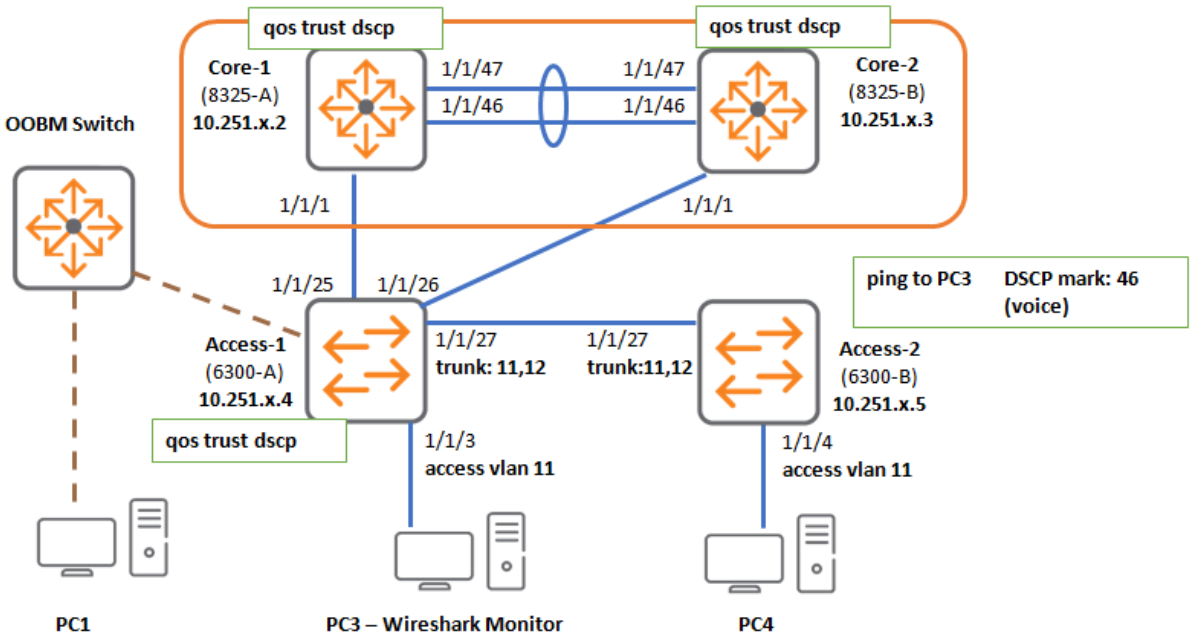

In the next steps, QOS trust mode on the switches will be set to globally trust the incoming DSCP value. This will ensure that the switch will inspect every incoming packet for a DSCP value, and apply it to the correct local-mapping based on the DSCP-map.

19. On Access1, enable the global DSCP trust option.

```
ICX-Tx-Access1(config)# qos trust dscp
```

Access1

20. Clear the port counters on Access1.

```
ICX-Tx-Access1(config)# do clear interface statistics
```

Access2

21. On Access2, repeat the ping (use the up arrow to get the previous command).

```
do ping 10.x.11.y datagram-size 14000 repetitions 10 tos 184
```

Access1

22. Once the ping has completed, review the queue statistics on the Access1 switch.

```
ICX-Tx-Access1(config)# show interface 1/1/3 queues
Interface 1/1/3 is up
```

```
 Admin state is up
                  Tx Bytes      Tx Packets       Tx Errors
 Q0                      0               0               0
 Q1                    476               7               0
 Q2                      0               0               0
 Q3                      0               0               0
 Q4                      0               0               0
 Q5                 144280             100               0
 Q6                      0               0               0
 Q7                   2475              20               0

ICX-Tx-Access1(config)#
```

The DSCP marking of the traffic did not change, but now the switch is using this mark in the incoming packet to assign them to the correct local-priority, which is assigned to queue 5.

Trust End-to-End

The Core VSX also needs to trust DSCP to ensure packets are assigned to the correct queue. It is a best practice to enable global **qos trust dscp** for VSX on Core 1 and Core2 and review the traffic assignment.

Core1

23. Enable the global trust on Core1.

```
ICX-Tx-Core1(config)# qos trust dscp
```

Core2

24. Repeat this step on Core2.

```
ICX-Tx-Core2(config)# qos trust dscp
```

Access2

25. On Access2, repeat the ping (use the up arrow to get the previous command).

```
do ping 10.x.11.y datagram-size 14000 repetitions 10 tos 184
```

Core1

Review the lag1 (LAG to Access1) queue statistics. You should see that some traffic was sent into Queue 5.

```
ICX-Tx-Core1# show interface lag1 queues
 Aggregate-name lag1
 Aggregated-interfaces :
 1/1/1
 Speed 10000 Mb/s
                  Tx Bytes      Tx Packets       Tx Errors
 Q0                      0               0               0
 Q1                  84480             331               0
```

```
 Q2                    0                 0                 0
 Q3                    0                 0                 0
 Q4                    0                 0                 0
 Q5               132160               100                 0
 Q6                  992                 8                 0
 Q7                 3670                29                 0
```

In case there is no traffic stats in queue5, check the command for the peer VSX system, since the link-aggregation on Access1 may have sent the traffic to Core2.

```
ICX-Tx-Core1# show interface lag1 queues vsx-peer
 Aggregate-name lag1
 Aggregated-interfaces :
 1/1/1
 Speed 10000 Mb/s
                  Tx Bytes      Tx Packets       Tx Errors
 Q0                     0               0               0
 Q1               4566671           29284               0
 Q2                     0               0               0
 Q3                     0               0               0
 Q4                     0               0               0
 Q5                132160             100               0
 Q6                 89836            1266               0
 Q7               1318752           10425               0
```

Task 3: LLDP Device Profile for QOS Trust

Diagram

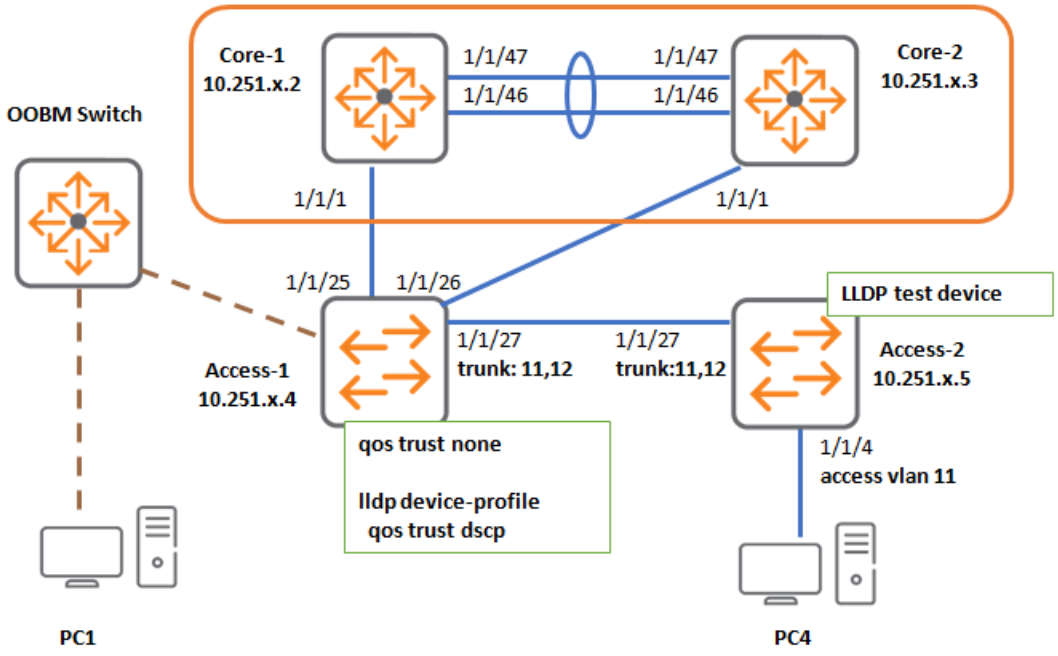

Objectives

Some customers may prefer to control the DSCP and queue assignment using classifiers and policies, so they may not want to enable the global qos trust DSCP on the access switches.

However, an Aruba controlled AP (this is an AP connected to a Mobility Controller [MC]), can use this DSCP mark as well. It will send wireless traffic over an IP GRE tunnel to the MC and, based on the configuration, the wireless QOS may be copied into the IP DSCP mark. This ensures a high-priority wireless packet can also be recognized as a high-priority packet on the path between the AP and the MC. In this case, the switch should trust the DSCP mark if the traffic comes from a port with an Aruba AP connected.

AOS-CX switches support this scenario using LLDP based device profiles. The LLDP device trust allows the administrator to leave the global trust level to 'none', since the switch will automatically set the port trust to DSCP when a matching LLDP device is discovered on a switch port.

Steps

Configured a device profile with LLDP trust settings.

Use the Access2 Switch to simulate an LLDP device profile to test the device profile.

Access1

1. On Access1, remove the global trust DSCP.

```
ICX-Tx-Access1(config)# qos trust none
```

2. Clear statistics and verify incoming marked DSCP traffic is not trusted anymore. Ping test from Access2 should arrive in normal queue (Q1) again.

```
ICX-Tx-Access1(config)# do clear interface statistics
```

Access2

3. Repeat the test ping on Access2.

```
do ping 10.x.11.y datagram-size 14000 repetitions 10 tos 184
```

Access1

4. Verify on Access1 that the traffic is handled in the normal queue again.

```
ICX-Tx-Access1(config)# show interface lag255 queues
Aggregate-name lag255
 Aggregated-interfaces :
 1/1/25  1/1/26
 Speed 20000 Mb/s
                  Tx Bytes      Tx Packets       Tx Errors
 Q0                     0               0               0
 Q1                317172             223               0
 Q2                     0               0               0
 Q3                     0               0               0
 Q4                     0               0               0
 Q5                     0               0               0
 Q6                     0               0               0
 Q7                  1240              10               0
```

5. Review the default port trust mode.

```
ICX-Tx-Access1(config)# show interface 1/1/27

Interface 1/1/27 is up
 Admin state is up
 Link transitions: 5
 Description:
 Hardware: Ethernet, MAC Address: 88:3a:30:98:30:ce
 MTU 1500
 Type SFP+DAC1
 Full-duplex
 qos trust none
 Speed 10000 Mb/s
 Auto-negotiation is off
 Flow-control: off
```

Configure the Device Profile

6. On Access1, configure the LLDP device profile role with DSCP trust.

```
ICX-Tx-Access1(config)# port-access role demo-lldp-switch
ICX-Tx-Access1(config-pa-role)# trust-mode dscp
ICX-Tx-Access1(config-pa-role)# exit
```

7. Configure the LLDP device profile with an LLDP group for ArubaOS-CX switches.

In a real deployment, this would be an Aruba AP, but the Access2 switch is used to simulate the behavior.

> **NOTE:** All students should verify the model number of Access1 and Access2 – For example, in **Pod11** they are model **JL660A**.

```
ICX-Tx-Access1(config)# port-access lldp-group demo-switch
ICX-Tx-Access1(config-lldp-group)# match sys-desc JL668A
ICX-Tx-Access1(config-lldp-group)# exit
```

```
ICX-Tx-Access1(config)# port-access device-profile dev-aoscx-switch
ICX-Tx-Access1(config-device-profile)# associate role demo-lldp-switch
ICX-Tx-Access1(config-device-profile)# associate lldp-group demo-switch
ICX-Tx-Access1(config-device-profile)# enable
ICX-Tx-Access1(config-device-profile)# exit
```

8. Verify the LLDP profile has been applied to the port 1/1/27.

```
ICX-Tx-Access1(config)# show port-access device-profile interface all
  Port 1/1/27, Neighbor-Mac  88:3a:30:97:b6:00
    Profile Name:           : dev-aoscx-switch
    LLDP Group:             : demo-switch
    CDP Group:              :
    Role:                   : demo-lldp-switch
    State:                  : applied
    Failure Reason:         :
```

9. Verify the changed port QOS trust mode.

```
ICX-Tx-Access1(config)# show interface 1/1/27

Interface 1/1/27 is up
 Admin state is up
 Link transitions: 1
 Description:
 Hardware: Ethernet, MAC Address: 88:3a:30:98:30:ce
 MTU 1500
 Type SFP+DAC1
 Full-duplex
 qos trust dscp
```

```
 Speed 10000 Mb/s
 Auto-negotiation is off
 Flow-control: off
 Error-control: off
 VLAN Mode: native-untagged
 Native VLAN: 11
 Allowed VLAN List: 11-12
```

Access2

10. Repeat the test ping on Access2.

```
do ping 10.x.11.y datagram-size 14000 repetitions 10 tos 184
```

Access1

11. Check queue statistics of interface lag255 on Access1, the traffic should now appear in Queue 5 again.

```
ICX-Tx-Access1(config)# show interface lag255 queues
 Aggregate-name lag256
 Aggregated-interfaces :
 1/1/25  1/1/26
 Speed 20000 Mb/s
                 Tx Bytes      Tx Packets       Tx Errors
 Q0                     0               0               0
 Q1                201464             165               0
 Q2                     0               0               0
 Q3                     0               0               0
 Q4                     0               0               0
 Q5                144280             100               0
 Q6                     0               0               0
 Q7                  2680              20               0
```

This demonstrates how the AOS-CX switch can dynamically enable QOS DSCP trust based on the LLDP neighbor, such as an Aruba AP.

12. Lab cleanup: Disable the LLDP profile.

```
ICX-Tx-Access1(config)# port-access device-profile dev-aoscx-switch
ICX-Tx-Access1(config-device-profile)# disable
ICX-Tx-Access1(config-device-profile)# exit
```

Task 4: QOS Classification

Diagram

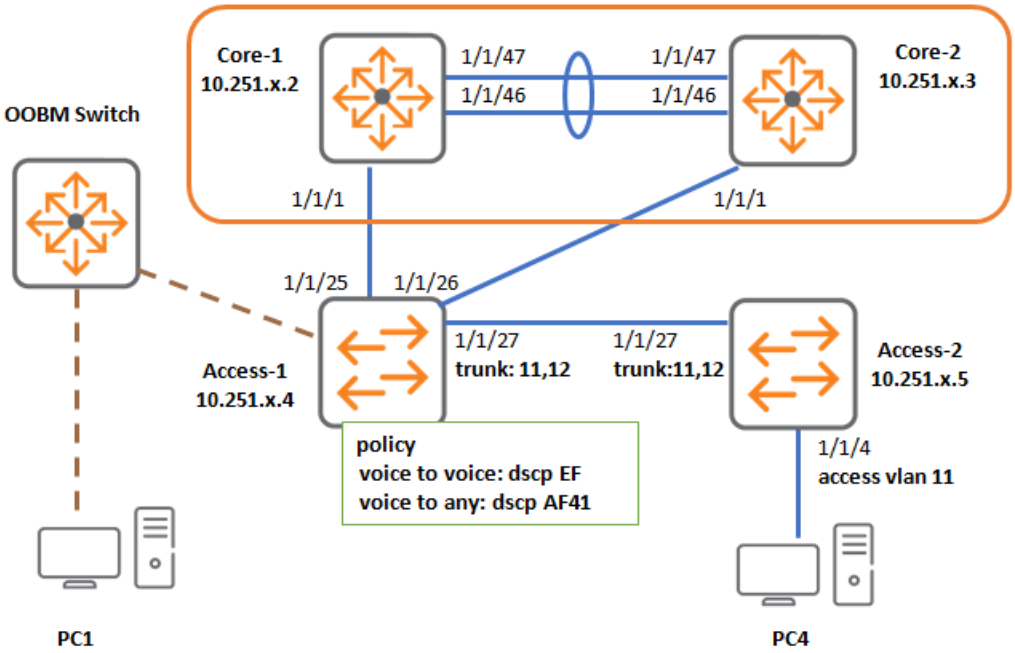

Objectives

In this task, the classifier-based QOS will be demonstrated. The purpose of the classifier is to provide the administrator a mechanism to remark and classify traffic in case the endpoint cannot perform the marking by itself, or the endpoint is not trusted to perform the marking.

In this task, traffic from the Voice VLAN 12 will be changed for these two scenarios:

- **Voice to Voice**: Traffic coming from the simulated voice VLAN 12 to the other VLAN 12 hosts will be marked as DSCP 46, so even when an unmarked ping is done, the traffic will be marked with value 46.
- **Voice to other**: Traffic coming from the simulated voice VLAN 12 to any other destination will be marked as DSCP AF41. This is just an example mark to demonstrate the feature.

Steps

To use the classifier-based QOS, the administrator should complete these steps:

- Define the classes
- Assign the classes into a policy and set actions
- Bind the policy to an interface or VLAN

Define the Classes

Access1

1. On Access1, define a new class for the voice-to-voice traffic, set the count option as well to get class statistics on matched traffic. The 'VLAN' keyword applies to the given VLAN ID as an additional condition.

Wildcard Mask

In the IP class definition, wildcard masks are supported. On some platforms, wildcard masks are inversed bits. For AOC-CX, the wildcard mask is using the same logic as a normal subnet mask, so a 1 bit to match, 0 bit to ignore the value.

Therefore, 10.0.12.0/0.255.0.255 means traffic:

- matching '10' in first byte
- with any value in second byte
- matching '12' in the third byte
- with any value in the fourth byte

```
ICX-Tx-Access1(config)# class ip voice-dst-voice
ICX-Tx-Access1(config-class-ip)# 10 match any any 10.0.12.0/255.0.255.0 vlan 12
count
ICX-Tx-Access1(config-class-ip)# exit
```

2. Define a new class for the voice to any traffic with the count option.

```
ICX-Tx-Access1(config)# class ip voice-dst-any
ICX-Tx-Access1(config-class-ip)# 10 match any any any vlan 12 count
ICX-Tx-Access1(config-class-ip)# exit
```

Define the Policy and Assign Actions to the Classes

Now combine the traffic classes into a policy with actions.

For the voice to voice traffic, a remark action of DSCP value 46 (EF) will be set.

For the voice to any traffic, a remark action of DSCP value 34 (AF41) will be set.

3. Define the policy.

```
ICX-Tx-Access1(config)# policy access
ICX-Tx-Access1(config-policy)# 10 class ip voice-dst-voice action dscp EF
ICX-Tx-Access1(config-policy)# 20 class ip voice-dst-any action dscp AF41
ICX-Tx-Access1(config-policy)# exit
```

Apply the Policy

A policy can be applied at various levels (global/VLAN/interface). In this lab, the policy will be applied at the interface level.

4. On the port connected to Access2, apply the QOS policy on the inbound traffic.

```
ICX-Tx-Access1(config)# interface 1/1/27
ICX-Tx-Access1(config-if)# apply policy access in
ICX-Tx-Access1(config-if)# exit
```

5. Review the applied policy on port 1/1/27.

```
ICX-Tx-Access1(config)# show policy interface 1/1/27
Direction
           Name
           Additional Policy Parameters
  Sequence Comment
           Class Type
                    action
------------------------------------------------------------------------------
Inbound
           access
        10
           voice-dst-voice ipv4
                    dscp EF

        20
           voice-dst-any ipv4
                    dscp AF41
------------------------------------------------------------------------------
```

Verify Marking and Queuing for Voice to Data Scenario

In the next steps, the result of the configuration will be verified.

6. Prepare the test:

- On the PC3, connected to Access1, start a Wireshark trace
- On Access1, clear the interface statistics (run the command twice)

```
do clear interface statistics
```

Access2

7. On the Access2, run the ping to this PC (voice to any) without marking the traffic (so **no tos** in the command) The unmarked traffic should be remarked by the policy.

```
ICX-Tx-Access2(config)# do ping 10.x.11.y datagram-size 14000 repetition 10
```

Access1

8. On Access1, verify that the traffic has matched the policy by checking the hit counts of the policy. The number **100** in this example output is the number of packets that have matched the class in the policy.

```
ICX-Tx-Access1(config)# show policy hitcounts access
Statistics for Policy access:
Interface 1/1/27* (in):
           Hit Count  Configuration
10 class ip voice-dst-voice action dscp EF
                    0  10 match any any 10.0.12.0/0.255.0.255 vlan 12 count
20 class ip voice-dst-any action dscp AF41
                  100  10 match any any any vlan 12 count
* policy statistics are shared among each context type (interface, VLAN).
  For routed ingress, they are only shared within the same VRF.
  Use 'policy NAME copy' to create a new policy for separate statistics.
ICX-Tx-Access1(config)#
```

9. On the PC, stop the Wireshark trace and check the inbound DSCP value. This should be AF41.

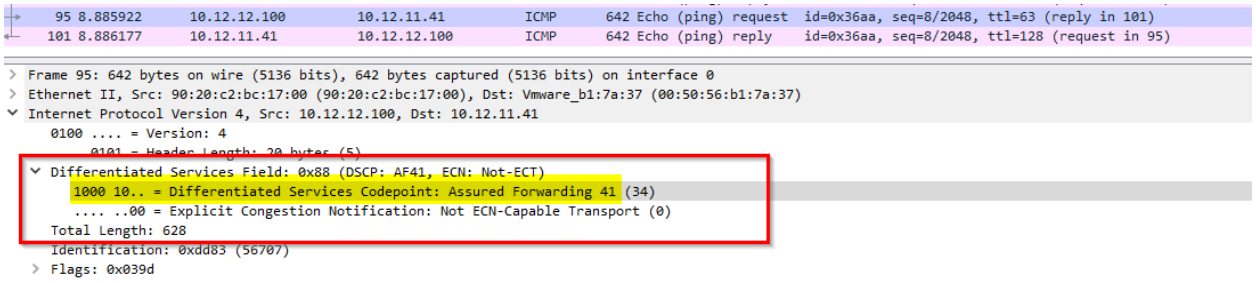

10. Now verify if the traffic was handled by the correct Queue on Access1.

```
ICX-Tx-Access1(config)# show interface 1/1/3 queues
Interface 1/1/3 is up
 Admin state is up
                  Tx Bytes      Tx Packets       Tx Errors
 Q0                      0               0               0
 Q1                 144824             108               0
 Q2                      0               0               0
 Q3                      0               0               0
 Q4                      0               0               0
 Q5                      0               0               0
 Q6                      0               0               0
 Q7                   3966              32               0
```

Q: Why is the traffic sent out in the 'normal' queue (Q1), instead of Q4 (the expected Queue for AF41 traffic).

A: Check the global qos trust option.

11. Review the global trust option.

```
ICX-Tx-Access1(config)# show qos trust
qos trust none
```

12. Adjust the global trust to DSCP.

```
ICX-Tx-Access1(config)# qos trust dscp
ICX-Tx-Access1(config)# show qos trust
qos trust dscp
```

13. Repeat the ping test from Access2 to the PC.

14. On Access1, the traffic should now be visible in Q4.

```
ICX-Tx-Access1(config)# show interface 1/1/3 queues
Interface 1/1/3 is up
Admin state is up
                  Tx Bytes      Tx Packets       Tx Errors
 Q0                     0               0               0
 Q1                160626             332               0
 Q2                     0               0               0
 Q3                     0               0               0
 Q4                144280             100               0
 Q5                     0               0               0
 Q6                     0               0               0
 Q7                122144             988               0
```

This concludes the first part of the validation.

Verify Marking and Queuing for Voice to Voice Scenario

In the next section, traffic from voice to voice will be verified.

This will be tested with a basic ping to the default gateway in the voice VLAN (10.x.12.1).

Access1

15. Prepare the test: on Access1, clear the interface statistics.

```
ICX-Tx-Access1(config)# do clear interface statistics
```

Access2

16. On Access2, ping to the voice default gateway (10.x.12.1) without marking the traffic.

```
ICX-Tx-Access2(config)# do ping 10.x.12.1 datagram-size 8000 repetitions 8
```

Access1

17. On Access1, verify the traffic matched by checking the hitcounts of the policy.

```
ICX-Tx-Access1(config)# show policy hitcounts access
Statistics for Policy access:
Interface 1/1/27* (in):
          Hit Count  Configuration
10 class ip voice-dst-voice action dscp EF
                  100  10 match any any 10.0.12.0/255.0.255.0 vlan 12 count
20 class ip voice-dst-any action dscp AF41
                  100  10 match any any any vlan 12 count
* policy statistics are shared among each context type (interface, VLAN).
  For routed ingress, they are only shared within the same VRF.
  Use 'policy NAME copy' to create a new policy for separate statistics.
```

18. Now verify that the traffic was assigned to the correct queue on the uplink interface.

```
ICX-Tx-Access1(config)# show interface lag255 queues
 Aggregate-name lag255
 Aggregated-interfaces :
 1/1/25  1/1/26
 Speed 20000 Mb/s
                  Tx Bytes      Tx Packets       Tx Errors
 Q0                     0               0               0
 Q1                335639             305               0
 Q2                     0               0               0
 Q3                     0               0               0
 Q4                     0               0               0
 Q5                     0               0               0
 Q6                     0               0               0
 Q7                  6432              48               0
```

Q: Why would the traffic be assigned to the normal queue?

_____

A: In the previous section, the QOS global trust was enabled, however, that does not seem to affect the current setup.

This is because any traffic that matches a policy will not be trusted, so it is not automatically assigned to a queue based on the trust map.

In the previous section, the trust did work, since the traffic was first sent to the VSX core (to be routed between the voice and data VLAN), so when the traffic was sent back from the Core to the Access1 switch, the global 'qos trust dscp' ensured that the incoming DSCP mark was used to assign the packet to the correct queue.

In the current example, the incoming traffic on 1/1/27 matches a policy, the policy remarks the DSCP, but the policy did not apply a local-priority to the traffic. Therefore, although the DSCP remark was done, the packet was still considered 'normal' inside the switch.

In the next section, the policy will be adjusted to apply both the DSCP mark and the local priority.

> **NOTE:** This also applied to the previous section, if the uplink LAG255 statistics would have been checked, it would have showed that the traffic was sent in the normal queue, even with the 'qos trust dscp' enabled.

To summarize:

Any match on an entry in a policy overrules the trust, so the administrator should assign a local-priority as well.

Adjust the policy to mark and local policy

Diagram

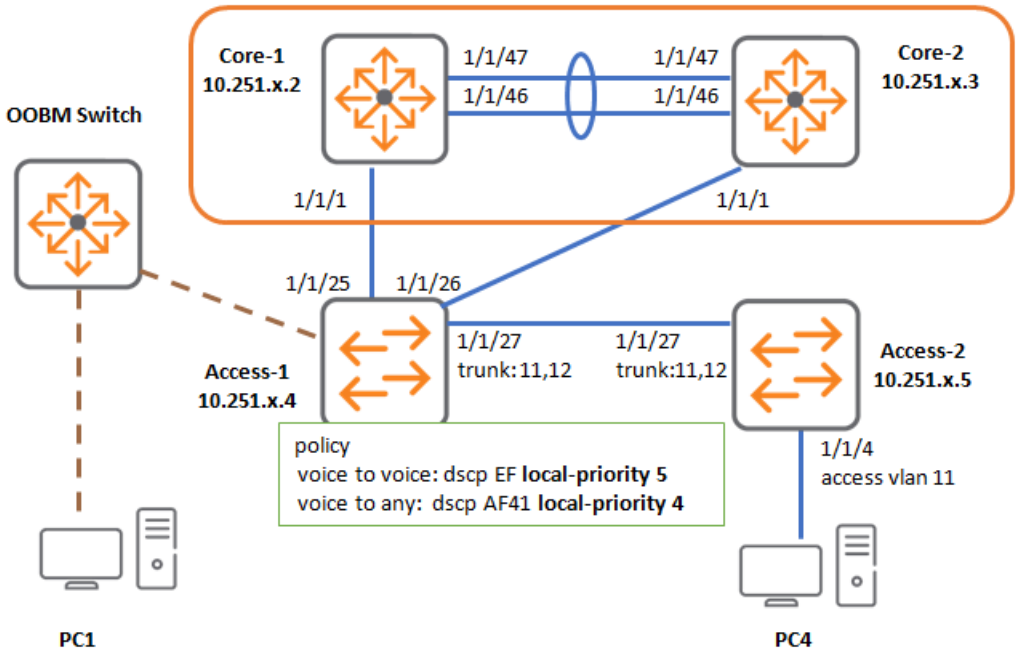

19. Enter the policy context, review the current configuration.

```
ICX-Tx-Access1(config)# policy access
ICX-Tx-Access1(config-policy)# show running-config current-context
policy access
    10 class ip voice-dst-voice action dscp EF
    20 class ip voice-dst-any action dscp AF41
```

20. Update the policy to include the local-priority action. Make sure to use the line numbers (10 and 20) in the commands. Otherwise, additional configuration lines would be added and the traffic would still match the original first line.

```
ICX-Tx-Access1(config-policy)#  10 class ip voice-dst-voice action dscp EF action
local-priority 5
ICX-Tx-Access1(config-policy)#  20 class ip voice-dst-any action dscp AF41 action
local-priority 4
```

21. Review the updated configuration. Make sure there are only.

```
ICX-Tx-Access1(config-policy)# show running-config current-context
policy access
    10 class ip voice-dst-voice action local-priority 5 action dscp EF
    20 class ip voice-dst-any action local-priority 4 action dscp AF41
ICX-Tx-Access1(config-policy)# exit
```

NOTE: If you would have added additional lines, these can be remove using 'no' + the line number. Example:

```
ICX-Tx-Access1(config-policy)# no 30
```

22. On Access1, check the LAG255 queue statistics.

```
ICX-Tx-Access1(config)# show interface lag255 queues
 Aggregate-name lag255
 Aggregated-interfaces :
 1/1/25  1/1/26
 Speed 20000 Mb/s
                 Tx Bytes      Tx Packets       Tx Errors
 Q0                     0               0               0
 Q1                146504             197               0
 Q2                     0               0               0
 Q3                     0               0               0
 Q4                   384               4               0
 Q5                     0               0               0
 Q6                     0               0               0
 Q7                 23584             176               0
```

23. and repeat the command.

```
ICX-Tx-Access1(config)# do repeat
```

The output will now automatically repeat every 2 seconds

Access2

24. On Access2, repeat the verification test with a ping without mark to the default gateway.

```
ICX-Tx-Access2(config)# do ping 10.x.12.1 datagram-size 8000 repetitions 8
```

Access1

25. On Access1, check the LAG255 queue statistics, traffic should pass over Q5 now.

```
************************************
Iteration : 6 Command : show interface lag256 queues
************************************
 Aggregate-name lag256
 Aggregated-interfaces :
 1/1/25  1/1/26
 Speed 20000 Mb/s
                 Tx Bytes      Tx Packets       Tx Errors
 Q0                     0               0               0
 Q1                149021             211               0
 Q2                     0               0               0
 Q3                     0               0               0
 Q4                   384               4               0
 Q5                144280             100               0
 Q6                     0               0               0
 Q7                 27336             204               0
```

26. Use <CTRL>-C to stop the repeat.

This demonstrates the use of classifiers, policies and actions.

Task 5: Queue configuration

Objectives

In this task, the queue configuration and priority mappings will be reviewed.

The priority mapping tables allow the administrator to customize the trust maps:

- DSCP to COS mapping
- COS to local-priority mapping

With the queue configuration, the administrator can control to which queue each of the local-priority values will be mapped.

With a schedule profile, the administrator can control the relative bandwidth of each of the queues and the type of scheduling, such as strict priority or DWRR.

Steps

Review default statistics

Change queue profile

Apply queue profile

Change schedule profile

Apply schedule profile

Review default DSCP to local-priority mapping

Access1

1. On Access1, review the default QOS DSCP map. Notice how each DSCP code point is mapped to a local-priority on the switch. This local-priority will be used to assign the packet to a queue. This map is used for any incoming traffic on ports with '**qos trust dscp'** (based on interface, global or dynamic with LLDP or port-access) .

```
ICX-Tx-Access1(config)# show qos dscp-map
DSCP     code_point local_priority cos color   name
-------- ---------- -------------- --- ------- ----
000000   0          1                  green   CS0
000001   1          1                  green
000010   2          1                  green
000011   3          1                  green
000100   4          1                  green
000101   5          1                  green
000110   6          1                  green
000111   7          1                  green
001000   8          0                  green   CS1
001001   9          0                  green
001010   10         0                  green   AF11
```

```
001011   11          0                   green
001100   12          0                   yellow  AF12
001101   13          0                   green
001110   14          0                   yellow  AF13
001111   15          0                   green
010000   16          2                   green   CS2
010001   17          2                   green
010010   18          2                   green   AF21
010011   19          2                   green
010100   20          2                   yellow  AF22
010101   21          2                   green
010110   22          2                   yellow  AF23
010111   23          2                   green
011000   24          3                   green   CS3
011001   25          3                   green
011010   26          3                   green   AF31
011011   27          3                   green
011100   28          3                   yellow  AF32
011101   29          3                   green
011110   30          3                   yellow  AF33
011111   31          3                   green
100000   32          4                   green   CS4
100001   33          4                   green
100010   34          4                   green   AF41
100011   35          4                   green
100100   36          4                   yellow  AF42
100101   37          4                   green
100110   38          4                   yellow  AF43
100111   39          4                   green
101000   40          5                   green   CS5
101001   41          5                   green
101010   42          5                   green
101011   43          5                   green
101100   44          5                   green
101101   45          5                   green
101110   46          5                   green   EF
101111   47          5                   green
110000   48          6                   green   CS6
110001   49          6                   green
110010   50          6                   green
110011   51          6                   green
110100   52          6                   green
110101   53          6                   green
110110   54          6                   green
110111   55          6                   green
111000   56          7                   green   CS7
111001   57          7                   green
111010   58          7                   green
111011   59          7                   green
111100   60          7                   green
111101   61          7                   green
111110   62          7                   green
111111   63          7                   green
```

Q: What is the local-priority that would be assigned to traffic marked with DSCP 46 (EF), typically used for voice?

_____

A: DSCP 46 will be mapped to local-priority 5.

Q: What is the local-priority that would be assigned to traffic without DSCP mark (0)?

_____

A: DSCP CS0 will be mapped to local-priority 1.

> **NOTE:** The DSCP map also shows the drop precedence level option as color code (green/yellow), however, this is beyond the scope of this course.

2. Review the default COS to local-priority map. This map would be used for ports that have the '**qos trust cos**' command. Most deployments will opt for the DSCP trust instead.

```
ICX-Tx-Access1(config)# show qos cos-map
code_point local_priority color   name
---------- -------------- ------- ----
0          1              green   Best_Effort
1          0              green   Background
2          2              green   Excellent_Effort
3          3              green   Critical_Applications
4          4              green   Video
5          5              green   Voice
6          6              green   Internetwork_Control
7          7              green   Network_Control
```

Q: What is the local-priority that would be assigned to traffic without COS mark (0)?

_____

A: COS 0 will be mapped to local-priority 1.

Review the queue configuration

3. Review the available queue-profiles. Only 1 queue-profile can be applied on the switch, this is a switch global option.

```
ICX-Tx-Access1(config)# show qos queue-profile
profile_status profile_name
-------------- ------------
applied        factory-default
```

4. Check the details of this queue profile.

```
ICX-Tx-Access1(config)# show qos queue-profile factory-default
queue_num local_priorities name
--------- ---------------- ----
0         0                Scavenger_and_backup_data
1         1
2         2
3         3
4         4
5         5
6         6
7         7
```

Q: To which queue will default traffic (local-priority 1) be assigned?

_____

A: Default (Best effort) traffic is assigned to queue1 based on this table. This is consistent with the queue statistics that were observed in previous tasks.

Review the scheduling profile

In the next steps, the scheduling of the queues will be reviewed.

The scheduling profile controls how each queue will be able to send out traffic, relative to other traffic in the queues.

5. Review the current scheduling profiles.

```
ICX-Tx-Access1(config)# show qos schedule-profile
profile_status profile_name
-------------- ------------
applied        factory-default
complete       strict
```

6. Check the details of the 'factory-default' profile.

```
ICX-Tx-Access1(config)# show qos schedule-profile factory-default
queue_num algorithm weight max-bandwidth_kbps
```

```
--------- --------- ------ -----------------
0         dwrr      1
1         dwrr      1
2         dwrr      1
3         dwrr      1
4         dwrr      1
5         dwrr      1
6         dwrr      1
7         dwrr      1
```

Q: What does the weight indicate?

_____

A: The weight is used to measure the 'weight' of the queue when multiple queues need to deliver traffic.

For example, when there is only traffic waiting in Q1 and Q5, the total weight at that moment is 1+1, so both Q1 and Q5 would have 1 out of 2, so up to 50% each.

In case only Q1 has traffic waiting, the total weight at that moment is 1, so Q1 would have 1 out of 1, so up to 100% of the bandwidth.

The result is that each Q has an equal share of the potential bandwidth when required.

In case the administrator wants to ensure that a queue needs more (or less) than this equal share, the weight can be adjusted.

Example adjustment of queue profile

In this example the schedule profile will be adjusted, so in case of traffic in various queues, these will be the minimum share values of the bandwidth.

To make the calculation easy, use a total of 100 for all the queues, this is not a requirement however.

Each individual queue weight can have a value between 1-1023.

Queue	**Weight**
0	1
1	39
2	10
3	10
4	10
5	10

6	10
7	10

This means:

If there is guest traffic (assuming this is assigned to queue 0) with weight 1, and normal traffic, with weight 39, guest would get 1 out of 40 (39+1), so 2,5%.

7. Define the new profile, assign each queue a value.

```
ICX-Tx-Access1(config)# qos schedule-profile icx
ICX-Tx-Access1(config-schedule)# dwrr queue 0 weight 1
ICX-Tx-Access1(config-schedule)# dwrr queue 1 weight 39
ICX-Tx-Access1(config-schedule)# dwrr queue 2 weight 10
ICX-Tx-Access1(config-schedule)# dwrr queue 3 weight 10
ICX-Tx-Access1(config-schedule)# dwrr queue 4 weight 10
ICX-Tx-Access1(config-schedule)# dwrr queue 5 weight 10
ICX-Tx-Access1(config-schedule)# dwrr queue 6 weight 10
ICX-Tx-Access1(config-schedule)# dwrr queue 7 weight 10
```

8. Review the configuration.

```
ICX-Tx-Access1(config-schedule)# show running-config current-context
qos schedule-profile icx
    dwrr queue 0 weight 1
    dwrr queue 1 weight 39
    dwrr queue 2 weight 10
    dwrr queue 3 weight 10
    dwrr queue 4 weight 10
    dwrr queue 5 weight 10
    dwrr queue 6 weight 10
    dwrr queue 7 weight 10
```

9. Apply the new schedule profile to queue 1/1/27.

```
ICX-Tx-Access1(config)# interface 1/1/27
ICX-Tx-Access1(config-if)# apply qos schedule-profile icx
ICX-Tx-Access1(config-if)# exit
```

10. Review the schedule profile list.

```
ICX-Tx-Access1(config-if)# show qos schedule-profile
profile_status profile_name
-------------- ------------
applied        icx
applied        factory-default
complete       strict
```

This concludes the scheduling and mapping profiles task.

Task 6: LLDP-MED and Voice VLAN configuration

Diagram

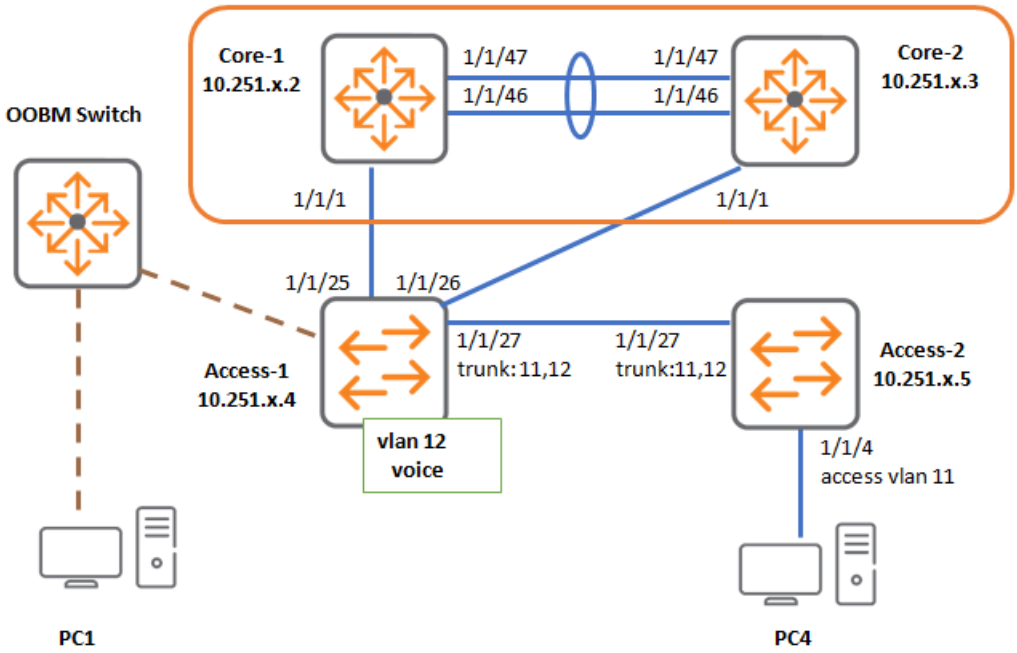

Objectives

When phones are connected to a switch that supports LLDP-MED, the switch should be configured with the correct voice VLAN. This will ensure that the switch announces this voice VLAN ID using LLDP-MED to the phone, so the phone will automatically send its own traffic with the voice VLAN tag.

The switch can also be configured to control which LLDP TLV (type/length/value) information is sent out. This may be required for security purposes, so the access ports cannot see the firmware version of the device.

Steps

Configure a voice VLAN

Access1

1. On Access1, configure VLAN 12 as the voice VLAN.

```
ICX-Tx-Access1(config)# vlan 12
ICX-Tx-Access1(config-vlan-12)# voice
ICX-Tx-Access1(config-vlan-12)# exit
```

2. Review the active voice VLAN.

```
ICX-Tx-Access1(config)# show vlan voice

------------------------------------------------------------------------------
-----------------------------
VLAN  Name                                Status  Reason               Type
Interfaces
------------------------------------------------------------------------------
-----------------------------
12    VLAN12                              up      ok                   static
1/1/27,lag255
```

No further configuration is required. The LLDP-MED TLVs are automatically enabled on all interfaces. When an LLDP-MED capable device comes online, the switch will announce the VLAN ID 12 as the voice VLAN.

Review the advertised TLV information

3. Review the advertised TLVs.

```
ICX-Tx-Access1(config)# show lldp tlv

TLVs Advertised
===============

Management Address
Port Description
Port VLAN-ID
System Capabilities
System Description
System Name
OUI
```

Access2

4. On Access2, review the details of the LLDP neighbor on port 1/1/27 (Access1) .

   NOTE: All students should verify the model number of Access1 and Access2 – For example, in **Pod11** they are model **JL660A**.

```
ICX-Tx-Access2(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 Chassis-Name         : ICX-Tx-Access1
Neighbor Chassis-Description  : Aruba JL668A  FL.10.04.0030
Neighbor Chassis-ID           : 88:3a:30:98:30:c0
Neighbor Management-Address   : 10.251.x.4
Chassis Capabilities Available : Bridge, Router
Chassis Capabilities Enabled  : Bridge, Router
```

```
Neighbor Port-ID                : 1/1/27
Neighbor Port-Desc              : 1/1/27
Neighbor Port VLAN ID           : 11
TTL                             : 120

Neighbor Mac-Phy details
Neighbor Auto-neg Supported     : true
Neighbor Auto-Neg Enabled       : false
Neighbor Auto-Neg Advertised    : Other
Neighbor MAU type               : 10 GIGBASEER
```

Q: What is the firmware version of the peer switch?

A: Based on the output, the version is FL.10.04.0030.

Q: What is the management IP address?

A: Based on the output, the management IP is 10.251.x.4.

Access1

5. On Access1, change the LLDP configuration, so the Description and management IP are no longer advertised.

```
ICX-Tx-Access1(config)# no lldp select-tlv system-description
ICX-Tx-Access1(config)# no lldp select-tlv management-address
```

Access2

6. Wait up to *30* seconds, then check the LLDP neighbors again on Access2.

```
ICX-Tx-Access2(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 Chassis-Name           : ICX-Tx-Access1
Neighbor Chassis-Description    :
Neighbor Chassis-ID             : 88:3a:30:98:30:c0
Neighbor Management-Address     :
Chassis Capabilities Available  : Bridge, Router
Chassis Capabilities Enabled    : Bridge, Router
Neighbor Port-ID                : 1/1/27
Neighbor Port-Desc              : 1/1/27
Neighbor Port VLAN ID           : 11
TTL                             : 120
```

```
Neighbor Mac-Phy details
Neighbor Auto-neg Supported    : true
Neighbor Auto-Neg Enabled      : false
Neighbor Auto-Neg Advertised   : Other
Neighbor MAU type              : 10 GIGBASEER
```

You've completed Lab 13!

Lab 14 Virtual Routing and Forwarding

Lab Diagram

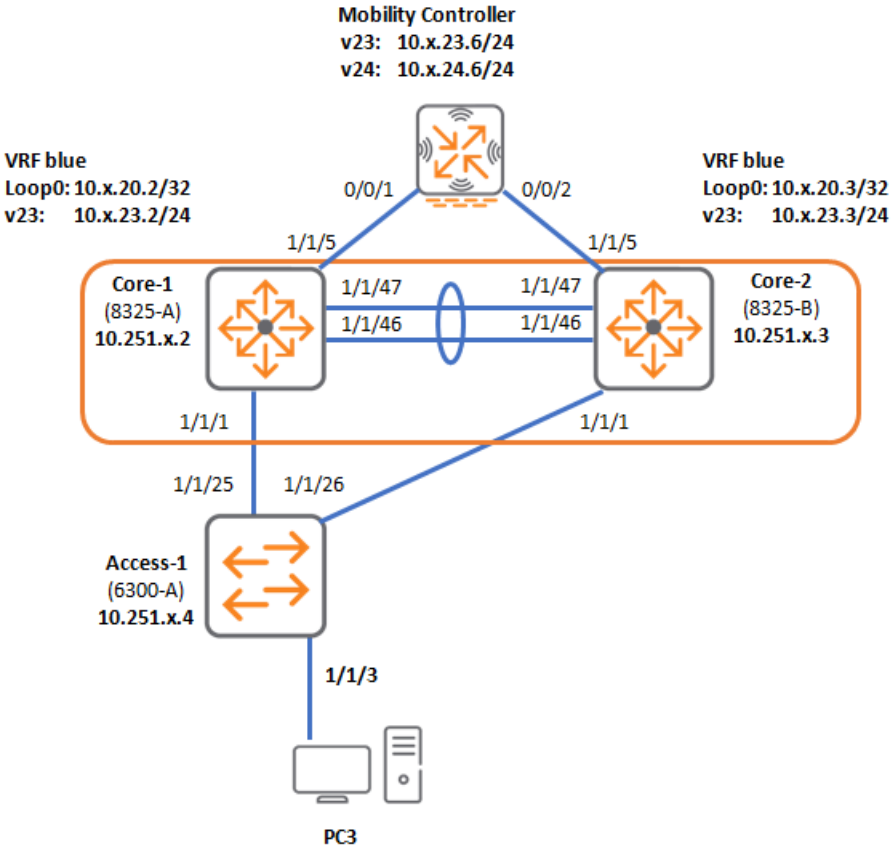

Overview

In this lab activity, the VRF feature will be explored. A new routing table will be defined on the VSX Core switches for an example customer 'blue'.

After defining a new routing context, the lab will show how L3 interfaces can be attached to a VRF, this will be done with Loopback and VLAN interfaces.

The lab then shows how a static route can be defined within the VRF context, how user subnets in the VRF can be configured and DHCP helper can be enabled.

The last section will show how a routing protocol, such as OSPF, can be enabled and bound to a VRF context.

Objectives

- Define a new VRF
- Attach L3 interfaces to a VRF
- Define static routes in a VRF
- Configure OSPF in a VRF context

Task 1: Prepare the lab start configuration

This lab is built on the base VSX topology.

Make sure to complete these steps to get the base VSX checkpoint configuration on the devices.

Steps (Required)

1. Open a console connection to the 6300A switch. Login using **admin** and a password of **aruba123**.

```
ICX-Tx-Access1# copy checkpoint icx-lab04-vsx running-config
Configuration changes will take time to process, please be patient.
ICX-Tx-Access1#
```

2. Open a console connection to the 6300B. Login using **admin** and a password of **aruba123**.

```
ICX-Tx-Access2# copy checkpoint icx-lab04-vsx running-config
Configuration changes will take time to process, please be patient.
ICX-Tx-Access2#
```

3. Open a console connection to the 8325A. Login using **admin** and a password of **aruba123**.

```
ICX-Tx-Core1# copy checkpoint icx-lab04-vsx running-config
Configuration changes will take time to process, please be patient.
ICX-Tx-Core1#
```

4. Open a console connection to the 8325B. Login using **admin** and a password of **aruba123**.

```
ICX-Tx-Core2# copy checkpoint icx-lab04-vsx running-config
Configuration changes will take time to process, please be patient.
ICX-Tx-Core2#
```

Task 2: Add a New Routing VRF

Diagram

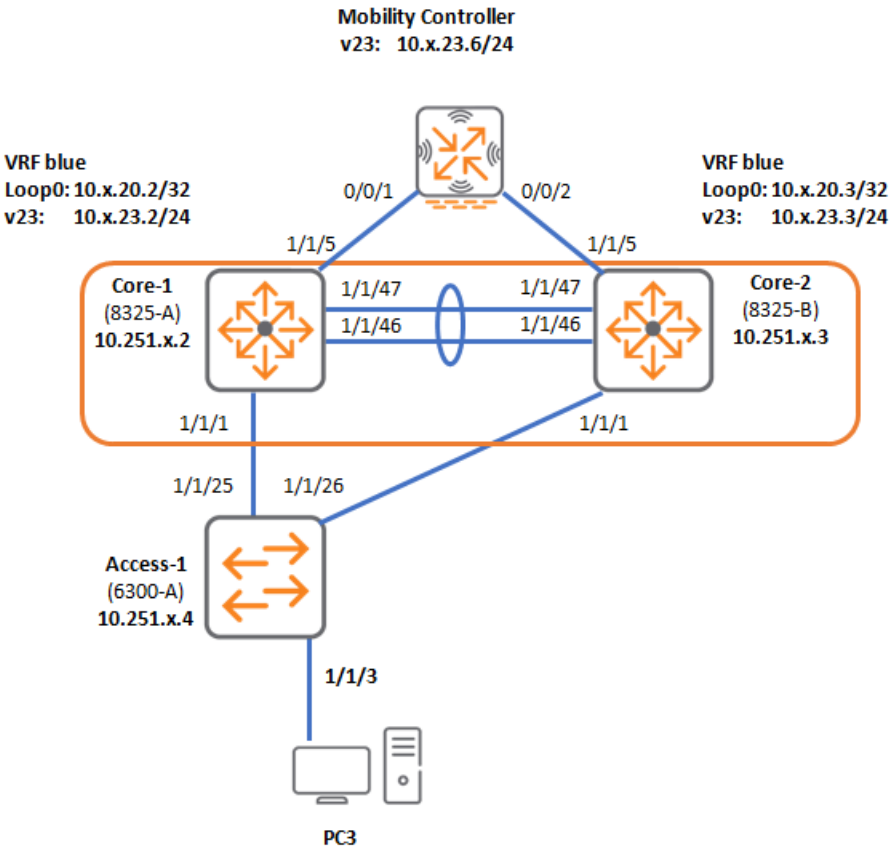

Objectives

In this task, a new VRF will be defined for a customer 'blue'. Next, L3 IP interfaces, such as Loopback and VLAN interfaces, will be attached to this VRF.

On the VSX Core, a routed VLAN with a VSX active-gateway will be defined to the upstream Aruba MC, that will be operating as the upstream firewall in this setup. With this routed connection in place, an example static route will be configured.

The next step will be to define a user subnet in the customer VRF and verify the PC has access to the customer resources.

Steps

Review existing VRFs

Core1

1. Open a terminal connection to Core1, enter the configuration mode.

2. Review the currently defined VRFs on the system.

```
ICX-Tx-Core1(config)# show vrf
```

Q: What are the names of the VRFs shown in this output?

_____

A: **default** and **KA** (the keep-alive VRF for VSX)

Q: Which VRF is not shown in this output?

_____

A: The VRF 'mgmt' is a default, built-in VRF for the management interface.

3. Review the VRF mgmt and the mgmt interface.

```
ICX-Tx-Core1(config)# show vrf mgmt
VRF Configuration:
------------------
VRF Name   : mgmt
use "show interface mgmt" for mgmt interfaces
```

```
ICX-Tx-Core1(config)# show interface mgmt
  Address Mode                  : static
  Admin State                   : up
  Mac Address                   : 90:20:c2:bc:17:01
  IPv4 address/subnet-mask      : 10.251.12.2/24
  Default gateway IPv4          : 10.251.12.254
  IPv6 address/prefix           :
  IPv6 link local address/prefix: fe80::9220:c2ff:febc:1701/64
  Default gateway IPv6          :
  Primary Nameserver            :
  Secondary Nameserver          :
```

Create new VRF

4. Define a new VRF named 'blue'.

```
ICX-Tx-Core1(config)# vrf blue
ICX-Tx-Core1(config-vrf)# exit
```

5. Define a new loopback interface with id '20' and assign it IP address 10.x.20.2/32.

```
ICX-Tx-Core1(config)# interface loopback 20
ICX-Tx-Core1(config-loopback-if)# ip address 10.x.20.2/32
```

6. Use the 'show vrf' command to review the interfaces and the VRF assignment.

```
ICX-Tx-Core1(config-loopback-if)# show vrf
VRF Configuration:
------------------
VRF Name   : default
        Interfaces              Status
        ----------------------------
        1/1/3                     down
...
        1/1/56                    down
        loopback20                up
        vlan11                    up
        vlan12                    up

VRF Name   : KA
        Interfaces              Status
        ----------------------------
        1/1/45                    up

VRF Name   : blue
        Interfaces              Status
        ----------------------------
```

Q: To which VRF is a L3 IP interface attached by default?

_____

A: The VRF named 'default', this is the switch global routing table.

7. Review the current configuration on the Loopback 20 interface.

```
ICX-Tx-Core1(config-loopback-if)# show run interface loopback 20
interface loopback 20
    ip address 10.x.20.2/32
    exit
```

8. Attach the Loopback 20 interface to the VRF blue. This will automatically remove the interface from the 'default' routing table.

```
ICX-Tx-Core1(config-loopback-if)# vrf attach blue
```

9. Review the current configuration on the loopback 20 interface again.

```
ICX-Tx-Core1(config-loopback-if)# show run interface loopback 20
interface loopback 20
    vrf attach blue
    exit
```

Q: What happens to the existing IP configuration when an interface is moved to a different VRF?

_____

A: The existing L3 IP configuration is removed. This is how the switch prevents the IP configuration of one customer to not be accidently inserted into another customer or security environment.

10. Assign the IP address again to the loopback 20 interface (use the up arrow to get the previous commands).

```
ICX-Tx-Core1(config-loopback-if)# ip address 10.x.20.2/32
ICX-Tx-Core1(config-loopback-if)# exit
```

Review the global routing table versus the VRF blue routing table

11. Review the global routing table using **'show ip route'**.

```
ICX-Tx-Core1(config)# show ip route

Displaying ipv4 routes selected for forwarding

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

10.x.3.0/24, vrf default
        via  1/1/7,  [0/0],  connected
10.x.3.2/32, vrf default
        via  1/1/7,  [0/0],  local
10.x.11.0/24, vrf default
        via  vlan11,  [0/0],  connected
10.x.11.2/32, vrf default
        via  vlan11,  [0/0],  local
10.x.12.0/24, vrf default
        via  vlan12,  [0/0],  connected
10.x.12.2/32, vrf default
        via  vlan12,  [0/0],  local
10.255.101.0/24, vrf default
        via  1/1/7,  [0/0],  connected
10.255.101.2/32, vrf default
        via  1/1/7,  [0/0],  local
```

Q: Is the 10.x.20.2/32 route available in the global routing table?

_____

A: No, a VRF provides complete L3 routing isolation.

12. Review the VRF blue routing table using 'show ip route vrf blue' command. Notice that none of the routes of the default routing table are available here.

```
ICX-Tx-Core1(config)# show ip route vrf blue

Displaying ipv4 routes selected for forwarding

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

10.x.20.2/32, vrf blue
        via  loopback20,  [0/0],  local
```

Configure a VLAN in the VRF Blue

In the next steps, a VLAN will be defined to make a routed connection to an upstream firewall. In this lab environment, the Aruba MC has the role of the upstream firewall, it will also operate as the DHCP server for hosts in the VRF blue.

Define a new VLAN for the upstream routed connection to the MC

The VSX core already has a VSX LAG (LAG5) to the MC, so only an extra VLAN and VLAN interface is required for this communication. This will be the routed connection between the VSX Core and the Aruba MC. VLAN 23 will be used for this purpose.

13. Define VLAN 23, use vsx-sync to have it created on Core2 as well.

```
ICX-Tx-Core1(config)# vlan 23
ICX-Tx-Core1(config-vlan-23)# vsx-sync
ICX-Tx-Core1(config-vlan-23)# exit
```

14. Configure the L3 interface and configure the IP settings for the new VLAN. Replace “Tx” with your table number, where Table 1 would be “01” and Table 12 would be “12”.

```
ICX-Tx-Core1(config)# interface vlan23
ICX-Tx-Core1(config-if-vlan)# vsx-sync active-gateways
ICX-Tx-Core1(config-if-vlan)# vrf attach blue
ICX-Tx-Core1(config-if-vlan)# ip address 10.x.23.2/24
ICX-Tx-Core1(config-if-vlan)# active-gateway ip mac 02:02:00:00:Tx:00
ICX-Tx-Core1(config-if-vlan)# active-gateway ip 10.x.23.1
```

```
ICX-Tx-Core1(config-if-vlan)# exit
ICX-Tx-Core1(config)#
```

Configure Core2 to Support the VRF blue

The previous steps were only executed on Core1. In the next steps, Core2 will be configured to support the VRF blue as well.

Core2

15. Open a terminal connection to Core2, enter the configuration mode.
16. Define the VRF blue and the Loopback interface 20.

```
ICX-Tx-Core2(config)# vrf blue
ICX-Tx-Core2(config-vrf)# exit
ICX-Tx-Core2(config)# interface loopback 20
ICX-Tx-Core2(config-loopback-if)# vrf attach blue
ICX-Tx-Core2(config-loopback-if)# ip address 10.x.20.3/32
ICX-Tx-Core2(config-loopback-if)# exit
```

17. Define the VLAN 23 L3 interface and verify the VSX active-gateway configuration was synchronized.

```
ICX-Tx-Core2(config)# interface vlan23
ICX-Tx-Core2(config-if-vlan)# vrf attach blue
ICX-Tx-Core2(config-if-vlan)# ip address 10.x.23.3/24
ICX-Tx-Core2(config-if-vlan)# show run current
interface vlan23
    vsx-sync active-gateways
    vrf attach blue
    ip address 10.x.23.3/24
    active-gateway ip mac 02:02:00:00:12:00
    active-gateway ip 10.x.23.1
ICX-Tx-Core2(config-if-vlan)# exit
```

Core1

18. On Core1, verify connectivity to the upstream MC with a ping from within the VRF 'blue' context.

```
ICX-Tx-Core1(config)# do ping 10.x.23.6 vrf blue
PING 10.x.23.6 (10.x.23.6) 100(128) bytes of data.
108 bytes from 10.x.23.6: icmp_seq=1 ttl=64 time=7.71 ms
108 bytes from 10.x.23.6: icmp_seq=2 ttl=64 time=0.335 ms
```

19. Review the ARP table on Core1.

```
ICX-Tx-Core1(config)# show arp
```

Q: Do you see an ARP entry for the 10.x.23.6 IP address? Why?

A: ARP is also processed per VRF, since a VRF may have overlapping IP addresses with another VRF, so isolation is required.

20. Review the ARP table of the VRF blue. Verify that the 10.x.23.6 address is listed in this ARP table.

```
ICX-Tx-Core1(config)# show arp vrf blue
IPv4 Address     MAC                Port         Physical Port    State      VRF

--------------------------------------------------------------------------------
10.x.23.6        20:4c:03:5f:a0:c2  vlan23       lag5             reachable  blue

Total Number Of ARP Entries Listed- 1.
--------------------------------------------------------------------------------
ICX-Tx-Core1(config)#
```

Configure Static Route in a VRF Context

Diagram

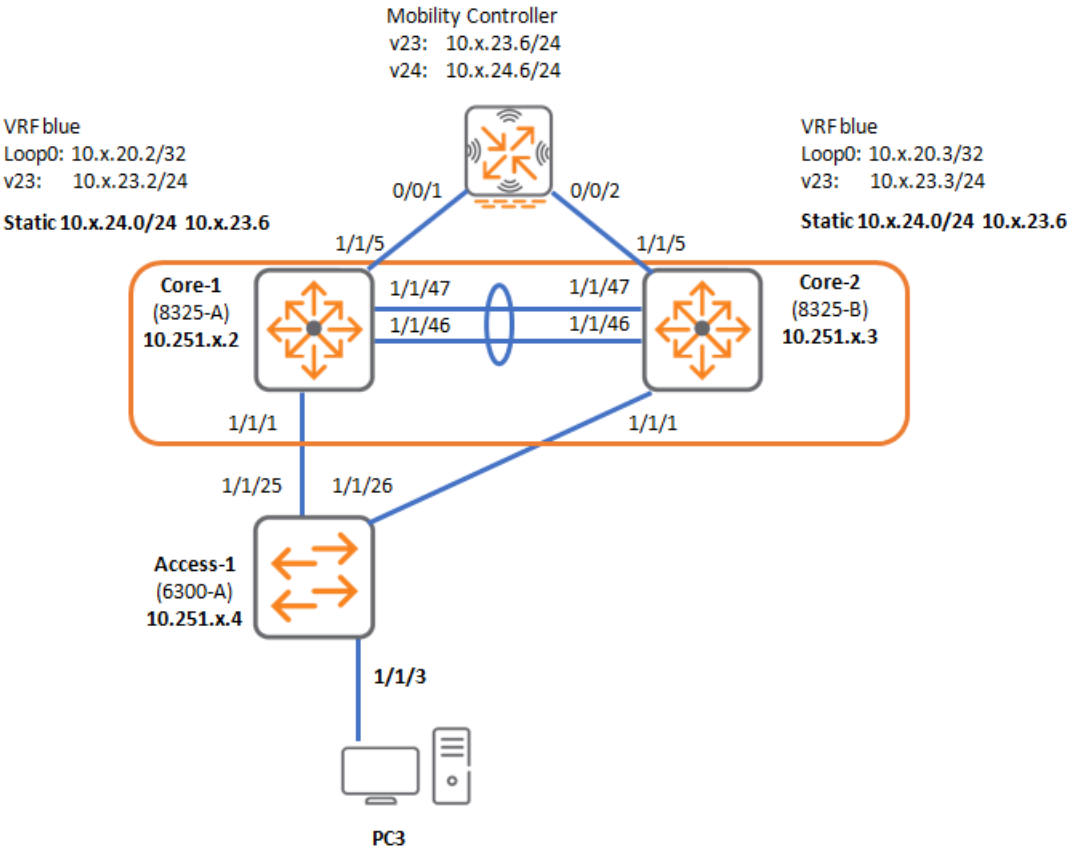

A VRF is just like a regular routing table, so it can be configured with static routes and dynamic routing protocols.

In the next steps, a static route will be added and verified in the VRF blue. The VSX Core currently has a direct link to the MC on the VLAN23. On the MC, there is an additional IP subnet 10.x.24.0/24, where the MC has IP 10.x.24.6. Currently, the VSX Core switches cannot reach the 10.x.24.0/24 subnet. By adding a static route to this subnet on the Core switches, the core switches will be able to reach this IP 10.x.24.6.

Core1

21. On Core1, review the current routing table of the VRF blue.

```
ICX-Tx-Core1(config)# show ip route vrf blue

Displaying ipv4 routes selected for forwarding

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

10.x.20.2/32, vrf blue
        via  loopback20,  [0/0],  local
10.x.23.0/24, vrf blue
```

```
        via  vlan23,  [0/0],  connected
10.x.23.2/32, vrf blue
        via  vlan23,  [0/0],  local
```

22. Add a static route for the 10.x.24.0/24 subnet to the MC IP.

```
ICX-Tx-Core1(config)# ip route 10.x.24.0/24 10.x.23.6 vrf blue
```

23. Verify the routing table again.

```
ICX-Tx-Core1(config)# show ip route vrf blue

Displaying ipv4 routes selected for forwarding

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

10.x.20.2/32, vrf blue
        via  loopback20,  [0/0],  local
10.x.23.0/24, vrf blue
        via  vlan23,  [0/0],  connected
10.x.23.2/32, vrf blue
        via  vlan23,  [0/0],  local
10.x.24.0/24, vrf blue
        via  10.x.23.6,  [1/0],  static
```

Core2

24. On Core2, add the static route as well.

```
ICX-Tx-Core2(config)# ip route 10.x.24.0/24 10.x.23.6 vrf blue
```

NOTE: It is also possible to enable the 'vsx-sync static-routes' under the VSX context on Core1. In this case the VSX synchronization would define the static route on Core2.

25. Verify that the routing table in the VRF blue now lists a route to the 10.x.24.0 subnet.

```
ICX-Tx-Core2(config)# show ip route 10.x.24.0 vrf blue

Displaying ipv4 routes selected for forwarding

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

10.10.24.0/24, vrf blue
via  10.10.23.6,  [1/0],  static

ICX-Tx-Core2(config)#
```

Test the Static Route

The MC has been pre-configured with IP 10.x.24.6 in the lab.

Attempt to ping this IP, this should succeed from both Core1 and Core2.

Core1

26. Verify the ping to 10.x.24.6 on Core1.

```
ICX-Tx-Core1(config)# do ping 10.x.24.6 vrf blue
PING 10.x.24.6 (10.x.24.6) 100(128) bytes of data.
108 bytes from 10.x.24.6: icmp_seq=1 ttl=64 time=0.468 ms
108 bytes from 10.x.24.6: icmp_seq=2 ttl=64 time=0.373 ms
```

Core2

27. Verify the ping to 10.x.24.6 on Core2.

```
ICX-Tx-Core2(config)# do ping 10.x.24.6 vrf blue
PING 10.x.24.6 (10.x.24.6) 100(128) bytes of data.
108 bytes from 10.x.24.6: icmp_seq=1 ttl=64 time=0.468 ms
108 bytes from 10.x.24.6: icmp_seq=2 ttl=64 time=0.373 ms
```

This demonstrates how a static route can be defined within a VRF context.

Configure User Subnet in the VRF Blue

Diagram

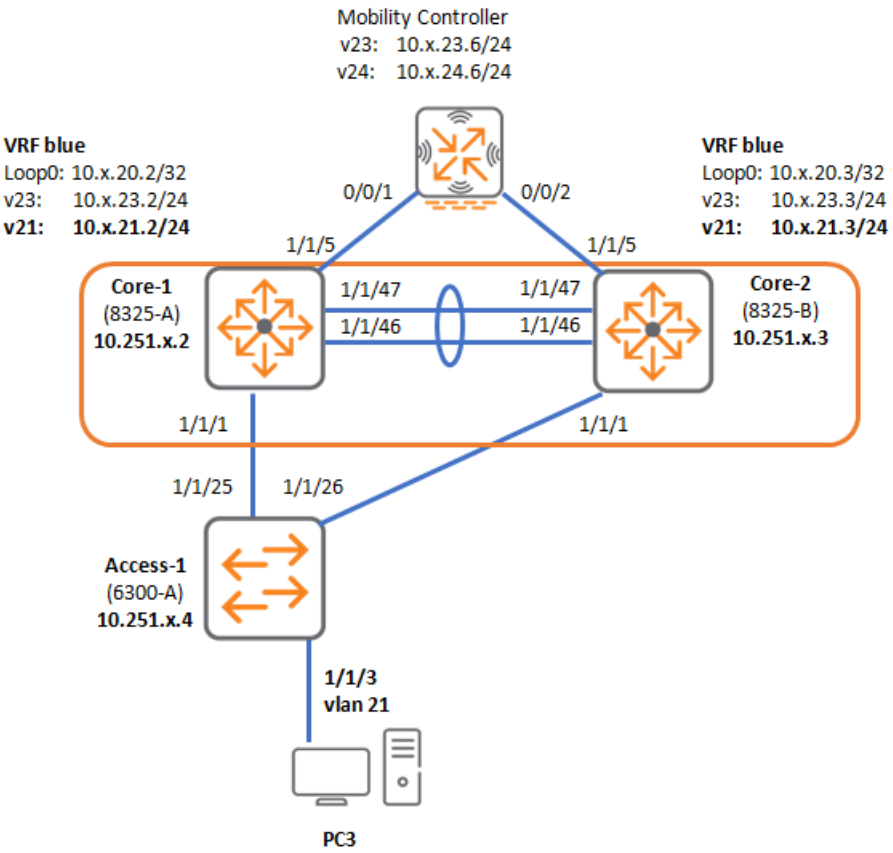

In these steps, a new user VLAN will be defined in the VRF blue for the PC connected to Access1. The VSX Core will be the default gateway for this subnet in the VRF blue, so active gateway and DHCP relay will be configured.

The new user VLAN is VLAN 21.

Core1

28. On Core1, define the new VLAN and IP interface.

```
ICX-Tx-Core1(config)# vlan 21
ICX-Tx-Core1(config-vlan-21)# vsx-sync
ICX-Tx-Core1(config-vlan-21)# exit
```

29. Define the IP Interface. Replace “Tx” with your table number, where Table 1 would be “01” and Table 12 would be “12”.

```
ICX-Tx-Core1(config)# interface vlan21
ICX-Tx-Core1(config-if-vlan)# vrf attach blue
ICX-Tx-Core1(config-if-vlan)# ip address 10.x.21.2/24
ICX-Tx-Core1(config-if-vlan)# vsx-sync active-gateways
ICX-Tx-Core1(config-if-vlan)# active-gateway ip mac 02:02:00:00:Tx:00
```

```
ICX-Tx-Core1(config-if-vlan)# active-gateway ip 10.x.21.1
```

30. Configure DHCP relay on the VLAN, use the MC as the DHCP server.

```
ICX-Tx-Core1(config-if-vlan)# ip helper-address 10.x.23.6
ICX-Tx-Core1(config-if-vlan)# exit
```

31. Allow the VLAN 21 on the LAG1 to Access1 and verify the configuration.

```
ICX-Tx-Core1(config)# interface lag 1
ICX-Tx-Core1(config-lag-if)# vlan trunk allowed 21
ICX-Tx-Core1(config-lag-if)# show run current
interface lag 1 multi-chassis
    no shutdown
    description access1
    no routing
    vlan trunk native 1
    vlan trunk allowed 1,11-13,21
    lacp mode active
ICX-Tx-Core1(config-lag-if)# exit
ICX-Tx-Core1(config)#
```

Core2

On Core2, several settings have already been synchronized by VSX:

- L2 VLAN
- L2 LAG allowed VLAN list
- L3 VLAN Active gateway settings
- L3 IP helper

The only remaining configuration steps are:

- Define the L3 interface
- Attach it to VRF
- Assign the interface an IP address

Core2

Define the L3 interface, attach it to the VRF blue and set the IP address.

```
ICX-Tx-Core2(config)# interface vlan 21
ICX-Tx-Core2(config-if-vlan)# vrf attach blue
ICX-Tx-Core2(config-if-vlan)# ip address 10.x.21.3/24
```

32. Verify the VSX active-gateway settings have been synchronized correctly.

```
ICX-Tx-Core2(config-if-vlan)# show running-config current-context
```

```
interface vlan21
    vsx-sync active-gateways
    vrf attach blue
    ip address 10.x.21.3/24
    active-gateway ip mac 02:02:00:00:XX:00
    active-gateway ip 10.x.21.1
    ip helper-address 10.x.23.6
ICX-Tx-Core2(config-if-vlan)# exit
```

33. Verify VLAN 21 exists and is allowed on LAG1 (to Access1).

```
ICX-Tx-Core2(config)# show vlan 21

--------------------------------------------------------------------------------
VLAN  Name              Status  Reason                Type      Interfaces
--------------------------------------------------------------------------------
21    VLAN21            up      ok                    static    lag1,lag5,lag256
ICX-Tx-Core2(config)#
```

Access1

34. Open a terminal connection on Access1 and enter configuration mode.
35. Define VLAN 21, verify it is allowed on the uplink.

```
ICX-TxAccess1(config)# vlan 21
ICX-TxAccess1(config-vlan-21)# exit
ICX-TxAccess1(config)# show vlan 21

--------------------------------------------------------------------------------
VLAN  Name            Status  Reason                Type      Interfaces
--------------------------------------------------------------------------------
21    VLAN21          up      ok                    static    lag255
```

36. Assign the port connected to PC3 (1/1/3) as access port in VLAN 21.

```
ICX-TxAccess1(config)# int 1/1/3
ICX-TxAccess1(config-if)# vlan access 21
ICX-TxAccess1(config-if)# exit
```

37. On the PC3, connected to Access1 port 1/1/3, reset (disable/enable) the 'Lab NIC' interface and verify the IP address received via DHCP. This should be in the 10.x.21.0/24 range.
38. On PC3, run a ping and traceroute to 10.x.24.6 (the MC) to test the routing and the static route.

```
C:\Users\student>ping 10.x.24.6

Pinging 10.x.24.6 with 32 bytes of data:
Reply from 10.x.24.6: bytes=32 time<1ms TTL=63
```

```
Reply from 10.x.24.6: bytes=32 time<1ms TTL=63

Ping statistics for 10.x.24.6:
    Packets: Sent = 2, Received = 2, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 0ms, Maximum = 0ms, Average = 0ms
Control-C
^C

C:\Users\student>tracert -d 10.x.24.6

Tracing route to 10.x.24.6 over a maximum of 30 hops

  1    <1 ms    <1 ms    <1 ms  10.x.21.3
  2    <1 ms    <1 ms    <1 ms  10.x.24.6

Trace complete.
```

NOTE: The traffic of PC3 can be hashed to either Core1 or Core2, so the first response in the traceroute could come from either Core1 (10.x.21.2) or Core2(10.x.21.3).

39. Close the command prompt using the 'exit' command.

This demonstrates how customer devices and subnets can be configured in an isolated routing context.

Task 3: OSPF Routing Protocols Inside the VRF

Diagram

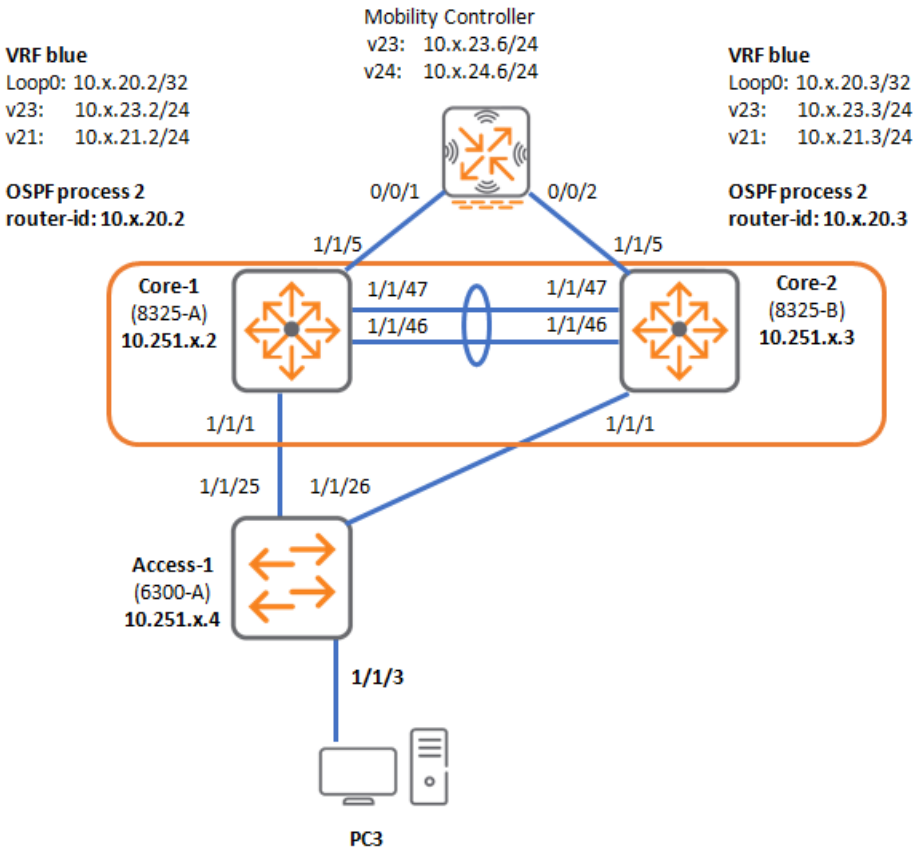

Objectives

In this task, OSPF routing will be configured inside the VRF blue. The only main difference will be to bind the OSPF process to the VRF. The other steps, such as the configuration of router-id, areas and interfaces, are the same as the OSPF of the default routing table.

Steps

Core1

1. On Core1, define a new OSPF process for VRF blue. Typically, OSPF process id '1' would be used for OSPF in the default routing table, so process id '2' is used here.

NOTE: Without the '**vrf blue**' command option, the OSPF process would operate in the default routing table.

```
ICX-Tx-Core1(config)# router ospf 2 vrf blue
```

2. Configure the loopback IP that was defined in the VRF blue as the router-id.

```
ICX-Tx-Core1(config-ospf-2)# router-id 10.x.20.2
```

3. Define OSPF area 0.

```
ICX-Tx-Core1(config-ospf-2)# area 0
ICX-Tx-Core1(config-ospf-2)# exit
```

4. Enable OSPF on the Loopback and upstream routed VLAN 23. Notice that the 'ip ospf' command requires the OSPF process number. Typically, that would be '1' for the OSPF process that is running in the global routing table. In this example, OSPF process id 2 was bound to the VRF blue. This is how the administrator can control to which OSPF process the interface belongs.

NOTE: This is different from the VRF attach command, since a single routing table (VRF) may contain multiple OSPF processes!

```
ICX-Tx-Core1(config)# interface loopback 20
ICX-Tx-Core1(config-loopback-if)# ip ospf 2 area 0
ICX-Tx-Core1(config-loopback-if)# exit

ICX-Tx-Core1(config)# interface vlan 23
ICX-Tx-Core1(config-if-vlan)# ip ospf 2 area 0
ICX-Tx-Core1(config-if-vlan)# exit
ICX-Tx-Core1(config)#
```

Verify OSPF operation

5. Review the OSPF settings.

```
ICX-Tx-Core1(config)# show ip ospf
OSPF Process is not running on VRF default.
```

Q: Why is there no information shown about OSPF?

A: The default “show ip ospf” command will look for an OSPF process that is bound to the 'default' VRF, but there is no OSPF in the 'default' routing context.

6. Repeat the command, but now use the 'vrf' command option, now the expected OSPF status should be displayed.

```
ICX-Tx-Core1(config)# show ip ospf vrf blue
Routing Process 2 with ID : 10.x.20.2 VRF blue
-----------------------------------------------

OSPFv2 Protocol is enabled
Graceful-restart is configured
Restart Interval: 120, State: inactive
Last Graceful Restart Exit Status: none
SPF: Start Time: 200ms, Hold Time: 1000ms, Max Wait Time: 5000ms
Maximum Paths to Destination: 4
Number of external LSAs 0, checksum sum 0
Number of areas is 1, 1 normal, 0 stub, 0 NSSA
Number of active areas is 1, 1 normal, 0 stub, 0 NSSA
BFD is disabled
Reference Bandwidth: 100000 Mbps
Area (0.0.0.0) (Active)
  Interfaces in this Area: 2 Active Interfaces: 2
  Passive Interfaces: 0 Loopback Interfaces: 1
  SPF calculation has run 8 times
  Area ranges:
  Number of LSAs: 3, checksum sum 75629
```

Core2

7. On Core2, VSX synchronization should have completed most of the configuration. Review the OSPF configuration.

```
ICX-Tx-Core2(config)# show run ospf
router ospf 2 vrf blue
   area 0.0.0.0
interface loopback 20
    ip ospf 2 area 0.0.0.0
interface vlan23
    ip ospf 2 area 0.0.0.0
```

NOTE: The only missing configuration command is the 'router-id', but an existing IP will be selected by default.

8. Review the OSPF status, notice the router ID 10.x.20.3 was automatically selected.

```
ICX-Tx-Core2(config)# show ip ospf vrf blue
Routing Process 2 with ID : 10.x.20.3 VRF blue
-----------------------------------------------

OSPFv2 Protocol is enabled
Graceful-restart is configured
Restart Interval: 120, State: inactive
Last Graceful Restart Exit Status: none
SPF: Start Time: 200ms, Hold Time: 1000ms, Max Wait Time: 5000ms
Maximum Paths to Destination: 4
Number of external LSAs 0, checksum sum 0
```

```
Number of areas is 1, 1 normal, 0 stub, 0 NSSA
Number of active areas is 1, 1 normal, 0 stub, 0 NSSA
BFD is disabled
Reference Bandwidth: 100000 Mbps
Area (0.0.0.0) (Active)
  Interfaces in this Area: 2 Active Interfaces: 2
  Passive Interfaces: 0 Loopback Interfaces: 1
  SPF calculation has run 8 times
  Area ranges:
  Number of LSAs: 3, checksum sum 75629
```

9. Review the OSPF neighbors, this should be Core1.

```
ICX-Tx-Core2(config)# show ip ospf neighbors vrf blue
OSPF Process ID 2 VRF blue
==========================

Total Number of Neighbors: 1

Neighbor ID      Priority  State             Nbr Address       Interface
-------------------------------------------------------------------------
10.x.20.2       1         FULL/BDR          10.x.23.2         vlan23
```

10. Review the IP routing table of the VRF blue, this should show an OSPF route to the loopback IP of the Core1 switch.

```
ICX-Tx-Core2(config)# show ip route vrf blue

Displaying ipv4 routes selected for forwarding

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

10.x.20.3/32, vrf blue
        via  loopback20,  [0/0],  local
10.x.20.2/32, vrf blue
        via  10.x.23.2,  [110/100],  ospf
10.x.21.0/24, vrf blue
        via  vlan21,  [0/0],  connected
10.x.21.3/32, vrf blue
        via  vlan21,  [0/0],  local
10.x.23.0/24, vrf blue
        via  vlan23,  [0/0],  connected
10.x.23.3/32, vrf blue
        via  vlan23,  [0/0],  local
10.x.24.0/24, vrf blue
        via  10.x.23.6,  [1/0],  static
```

11. Attempt to ping this remote loopback IP, using the Core2 loopback IP as the source. This should succeed, since Core1 should also have a route to the Core2 loopback IP.

```
ICX-Tx-Core2(config)# do ping 10.x.20.2 source 10.x.20.3 vrf blue
PING 10.x.20.2 (10.x.20.2) from 10.x.20.3 : 100(128) bytes of data.
108 bytes from 10.x.20.2: icmp_seq=1 ttl=64 time=0.233 ms
108 bytes from 10.x.20.2: icmp_seq=2 ttl=64 time=0.195 ms
108 bytes from 10.x.20.2: icmp_seq=3 ttl=64 time=0.185 ms
108 bytes from 10.x.20.2: icmp_seq=4 ttl=64 time=0.131 ms
108 bytes from 10.x.20.2: icmp_seq=5 ttl=64 time=0.133 ms

--- 10.x.20.2 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4092ms
rtt min/avg/max/mdev = 0.131/0.175/0.233/0.040 ms
ICX-Tx-Core2(config)#
```

This demonstrates how OSPF can be configured to operate within a VRF context.

You've completed Lab 14!

Lab 15: Switch Captive Portal- Captive Portal Authentication

Lab Diagram

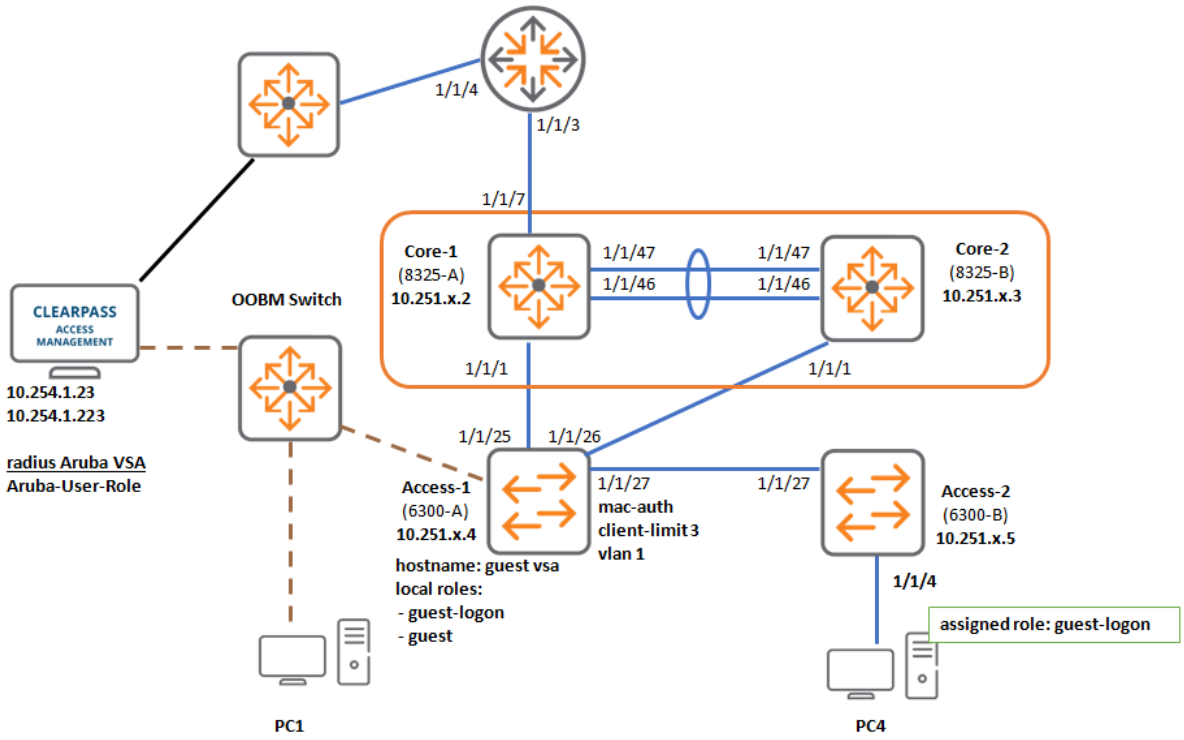

Overview

In this lab activity, the captive portal function will be configured on the on the Access1 switch. Captive portal authentication relies on MAC-authentication, which was previously configured on the port 1/1/27 of the Access1.

First, the actual captive portal (HTTP intercept) function will be configured using the 'guest-logon' user-role. Once this process has been configured, the lab guest account will connect to the network and the final 'guest' user-role will be assigned.

During this process, the RADIUS server will send a Change of Authorization (CoA) request to the switch to trigger the re-authentication with the resulting 'guest' user-role.

Requirements

This lab builds on the MAC-authentication lab configuration. The saved checkpoints will be loaded in the first task.

Task 1: Prepare the lab start configuration

Overview

- This lab is build on the MAC Authentication topology. This checkpoint was only saved on the Access1 and Access2 switches. The Core switches will be loaded with the VSX checkpoint configuration.
- Make sure to complete these steps to get the required configuration on the devices.

Steps (Required)

Access1 and Access2: MAC Authentication checkpoint

1. Open a console connection to the 6300A. Login using admin, password aruba123

```
ICX-Tx-Access1# copy checkpoint icx-lab11-mac running-config
Configuration changes will take time to process, please be patient.
ICX-Tx-Access1#
```

2. Open a console connection to the 6300B. Login using admin, password aruba123

```
ICX-Tx-Access2# copy checkpoint icx-lab11-mac running-config
Configuration changes will take time to process, please be patient.
ICX-Tx-Access2#
```

Core1 and Core2: VSX configuration checkpoint

3. Open a console connection to the 8325A. Login using admin, password aruba123

```
ICX-Tx-Core1# copy checkpoint icx-lab04-vsx running-config
Configuration changes will take time to process, please be patient.
ICX-Tx-Core1#
```

4. Open a console connection to the 8325B. Login using admin, password aruba123

```
ICX-Tx-Core2# copy checkpoint icx-lab04-vsx running-config
Configuration changes will take time to process, please be patient.
ICX-Tx-Core2#
```

RouterA: OSPF checkpoint

5. Open a console connection to RouterA. Login using admin, password aruba123

```
ICX-RouterA# copy checkpoint ospf running-config
```

Task 2: Define Captive Portal Role

Objectives

In this task, the guest-logon user-role will be defined and tested. This role will allow limited access to the network, so the user can get an IP address through DHCP and perform name resolution through DNS.

Traffic sent to HTTP TCP port 80 will be handled by the switch captive portal process and it will be redirected to the captive portal page on the ClearPass server. Any other traffic will be dropped for this role.

Steps

Access1

1. Open a terminal connection to Access1, enter the configuration mode.
2. For this lab, the ClearPass radius server has a number of lab-specific rules that are triggered based on the hostname of the switch. Rename the hostname of the Access1 to include the words **'guest-vsa'**. Replace “Tx” with your table information; for example, Table 1 would be “T1” and Table 12 would be “T12”.

> **NOTE:** The hostname is used as the NAS-identifier in the access-request to ClearPass. ClearPass has been pre-configured to look for the values **'guest** and' **'vsa**' in the NAS-identifier and to return various Aruba Vendor Specific Attributes (**VSA**) based on the **Aruba-User-Roles** that are needed for this lab.

```
ICX-Tx-Access1(config)# hostname ICX-Tx-Access1-guest-vsa
ICX-Tx-Access1-guest-vsa(config)#
```

Prepare Objects for Captive Portal Role

IP Classes

3. Define IP Class for web access to ClearPass.

```
ICX-Tx-Access1-guest-vsa(config)# class ip cppm-web
ICX-Tx-Access1-guest-vsa(config-class-ip)# match tcp any 10.254.1.23/32 eq 80
ICX-Tx-Access1-guest-vsa(config-class-ip)# match tcp any 10.254.1.23/32 eq 443
ICX-Tx-Access1-guest-vsa(config-class-ip)# exit
ICX-Tx-Access1-guest-vsa(config)#
```

4. Define IP Class for control protocols, such as DHCP, DNS and ICMP.

```
ICX-Tx-Access1-guest-vsa(config)# class ip dhcp-dns-icmp
ICX-Tx-Access1-guest-vsa(config-class-ip)#  match udp any any eq 67
ICX-Tx-Access1-guest-vsa(config-class-ip)#  match udp any any eq 68
ICX-Tx-Access1-guest-vsa(config-class-ip)#  match udp any any eq 53
ICX-Tx-Access1-guest-vsa(config-class-ip)#  match icmp any any
```

```
ICX-Tx-Access1-guest-vsa(config-class-ip)#  exit
```

5. Define IP Class for all web traffic.

```
ICX-Tx-Access1-guest-vsa(config)# class ip any-web
ICX-Tx-Access1-guest-vsa(config-class-ip)#  match tcp any any eq 80
ICX-Tx-Access1-guest-vsa(config-class-ip)#  match tcp any any eq 443
ICX-Tx-Access1-guest-vsa(config-class-ip)#  exit
```

Combine the IP Classes into a Port-access Policy

6. Define a new port-access policy. The logic for the policy is:
   - Allow the client device to get an IP address (DHCP), send DNS queries and ping (ICMP)
   - Allow access to the ClearPass host on port 80 and 443,but apply a rate limit to ensure the client cannot overload the ClearPass host.
   - Send all other web traffic to the local redirect process on the switch (CPU).

```
ICX-Tx-Access1-guest-vsa(config)# port-access policy cppm-redirect
ICX-Tx-Access1-guest-vsa(config-pa-policy)# class ip dhcp-dns-icmp
ICX-Tx-Access1-guest-vsa(config-pa-policy)# class ip cppm-web action cir kbps
1024 cbs 2048 exceed drop
ICX-Tx-Access1-guest-vsa(config-pa-policy)# class ip any-web action redirect
captive-portal
ICX-Tx-Access1-guest-vsa(config-pa-policy)# exit
```

Define Captive Portal Profile with URL to ClearPass

7. Define captive portal profile to ClearPass guest logon page. This will be used by the local redirect process of the switch.

   When the current client sends traffic on port 80, the port-access policy will apply the 'redirect captive-portal' action, so the traffic is delivered to the local redirect process. This redirect process will use the 'Captive Portal profile' to find out to what destination URL the client should be redirected. This URL can be unique for each captive portal role.

   First, the global captive portal profile must be defined with the URL, this profile will be linked in the user-role in the next step. This same URL profile could be used by multiple user-roles.

   - Profile name: **cppm-guest**
   - URL: **http://10.254.1.23/guest/guest-cx.php**

```
ICX-Tx-Access1-guest-vsa(config)# aaa authentication port-access captive-portal-
profile cppm-guest
```

```
ICX-Tx-Access1-guest-vsa(config-captive-portal)# url
http://10.254.1.23/guest/guest-cx.php
ICX-Tx-Access1-guest-vsa(config-captive-portal)# exit
```

Define New User-role: Associate the Policy and Captive Portal Profile

8. A new user-role 'guest-logon' is now defined. It represents the access the user will need in order to complete the captive portal logon to the network.

```
ICX-Tx-Access1-guest-vsa(config)# port-access role guest-logon
ICX-Tx-Access1-guest-vsa(config-pa-role)# associate captive-portal-profile cppm-
guest
ICX-Tx-Access1-guest-vsa(config-pa-role)# associate policy cppm-redirect
ICX-Tx-Access1-guest-vsa(config-pa-role)# vlan access 11
ICX-Tx-Access1-guest-vsa(config-pa-role)# exit
```

Test Redirect Role with PC4 (connected to Access2)

Guest captive portal is using mac-authentication to push the final network access to the switch. Since the current lab setup is using the PC4 on Access2 with mac-authentication, the guest lab will use this PC4 as well for the Captive Portal demonstration.

Access1

9. On the Access1 switch, verify the port configuration (no changes were made).

The switch port is still configured with mac-authentication, so any 'new' MAC Address will be sent to ClearPass for mac-authentication. Since ClearPass is now configured to return the 'guest-logon' Aruba VSA User-role for any MAC address that is not previously authenticated, the client should come online with this role.

```
ICX-Tx-Access1-guest-vsa(config)# show run int 1/1/27
interface 1/1/27
    no shutdown
    no routing
    vlan access 1
    aaa authentication port-access client-limit 3
    aaa authentication port-access mac-auth
        enable
    exit
```

10. Bounce the port on Access1 to clear any previously active authentications.

```
ICX-Tx-Access1-guest-vsa(config)# interface 1/1/27
ICX-Tx-Access1-guest-vsa(config-if)# shutdown
ICX-Tx-Access1-guest-vsa(config-if)# no shutdown
ICX-Tx-Access1-guest-vsa(config-if)# exit
```

Access2

11. Open a terminal connection to Access2 and enter the configuration mode

12. Enable port 1/1/4 (connected to PC4)

```
ICX-Tx-Access2(config)# interface 1/1/4
ICX-Tx-Access2(config-if)# no shutdown
ICX-Tx-Access2(config-if)# exit
```

PC4

13. On PC4, reset the **'Lab NIC'** network interface to trigger a DHCP request (disable/enable).

Access1

14. Verify the result of the initial mac-auth.

```
ICX-Tx-Access1-guest-vsa(config)# show aaa authentication port-access interface
1/1/27 client-status

Port Access Client Status Details

Client 00:50:56:b1:fc:9b, 005056b1fc9b
============================
  Session Details
  ---------------
    Port         : 1/1/27
    Session Time : 69s

  Authentication Details
  ----------------------
    Status          : mac-auth Authenticated
    Auth Precedence : dot1x - Not attempted, mac-auth - Authenticated

  Authorization Details
  ---------------------
    Role   : guest-logon
    Status : Applied
```

NOTE: It is possible that the MAC Address of the Access2 switch is also shown in the output. Since the Access2 switch cannot perform actual guest captive portal logon, the training guide is not showing it in the output and only shows the PC4 MAC Address.

15. Use PC1 (Management PC) to open a session to ClearPass and review your authentication entry in Access Tracker. You may see two mac-auth entries, one for the Access2 switch and one for the PC4. In this lab you will only use the PC4.

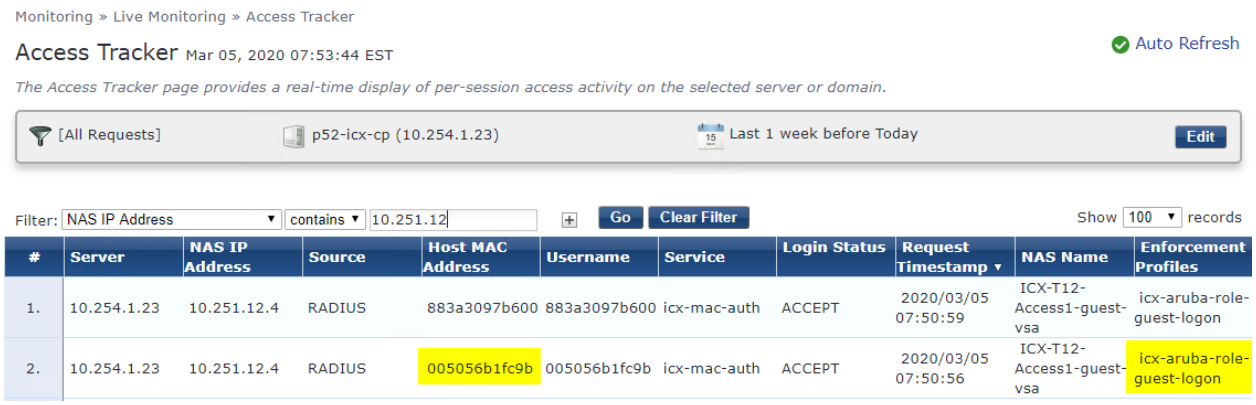

16. Open the entry for your PC4 (should start with MAC address **005056**). Review the **'Output'** tab, it should show the **'guest-logon'** Aruba VSA Aruba-User-Role.

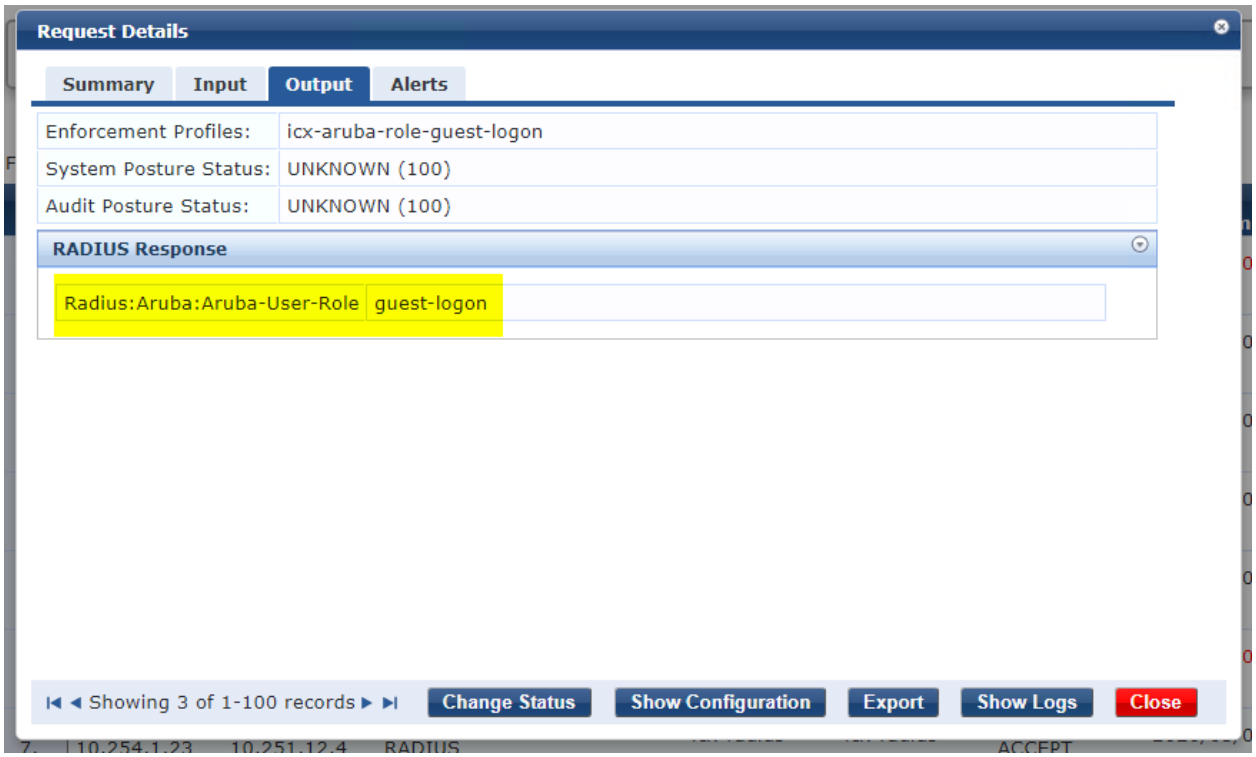

Test CoA

The Change of Authorization or Disconnect Message is critical in the Captive portal process.

Once the guest user has entered credentials on the ClearPass web page, ClearPass will mark the MAC address of this client as 'authenticated'. But that information will only be processed during the next authentication request.

After valid credentials have been entered on the Web page, the ClearPass server will send a CoA/DM to the switch to trigger the switch to perform re-authentication. The

guest user will see a 'please wait' page with a timer, so the re-authentication can take place at that moment.

In this lab, the 'guest' role and the 'guest-logon' role are in the same VLAN, so the guest user does not need to get a new IP Address, only a new access policy (different access rules in the 'guest-logon' and the 'guest' role). In case they would be assigned to different VLANs, the count-down timer should be large enough to handle a DHCP renew on the client side.

First, test the Disconnect Message for the client.

17. Click '**Change Status**'

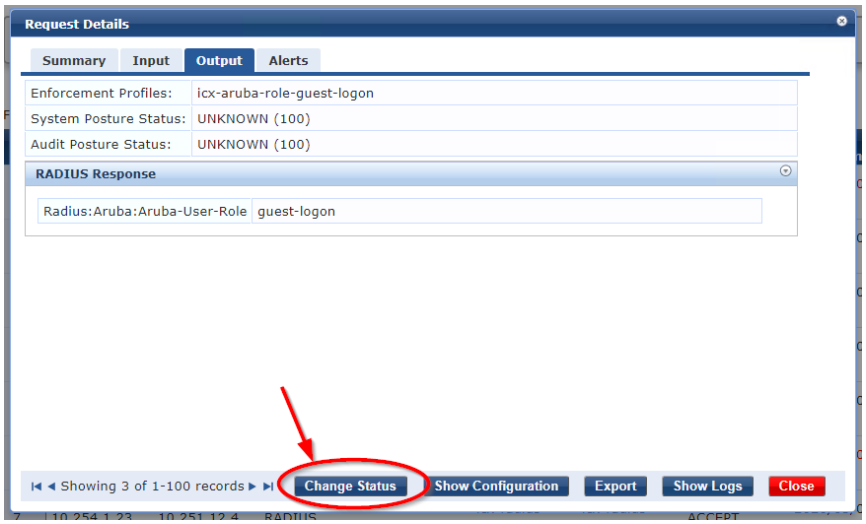

18. Click '**Submit**' to send the CoA.

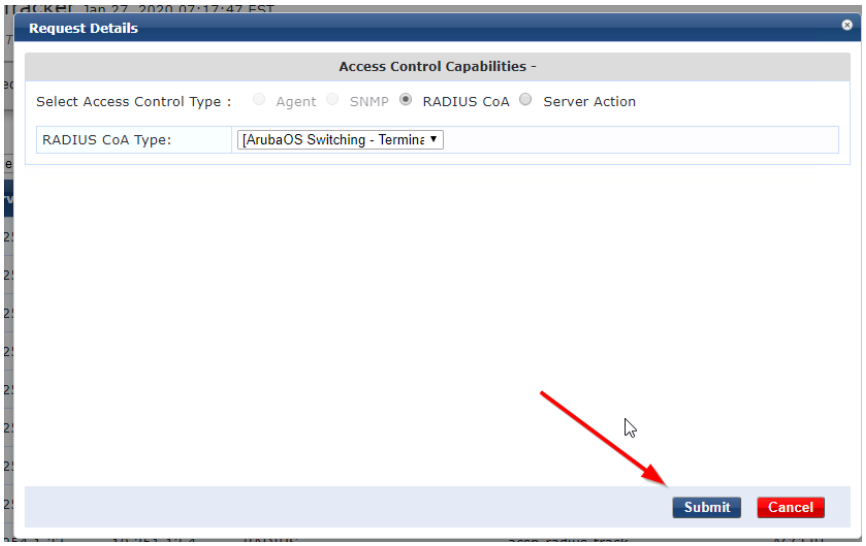

19. Make sure the CoA is 'successful', then click '**Close**' to close the window.

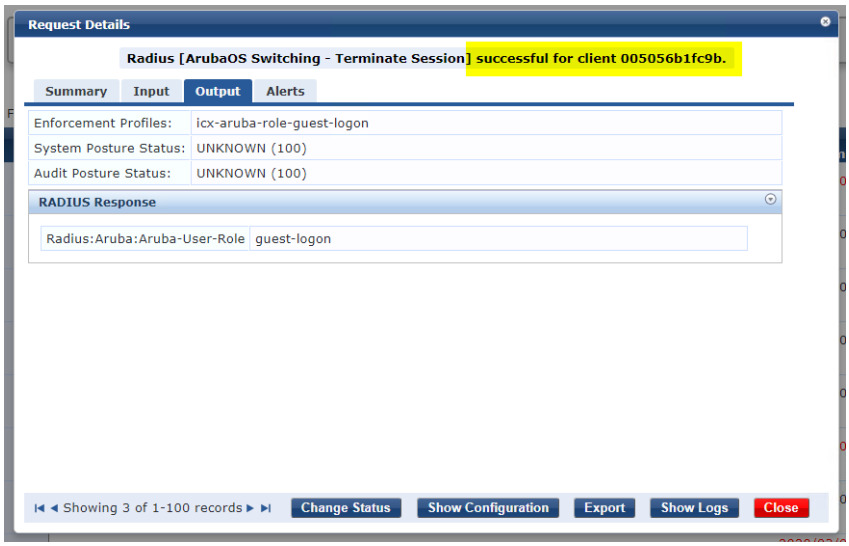

20. Take note of the last two characters of the client MAC address, you can use this in some ClearPass filters in the rest of this lab.

Test the Client

21. Open the PC4, start Firefox and attempt to navigate to http://10.254.1.223/.

.

NOTE: The lab may not be connected to the full internet, so internet DNS may or may not work. This is why the initial connection is made to an IP Address.

NOTE: Since there may not be a real internet host, the IP 10.254.1.223 is used to simulate a host on the internet. In reality, this is just a secondary address on the ClearPass server, but since this IP is not allowed by the guest-logon policy, it works fine for the captive portal redirect.

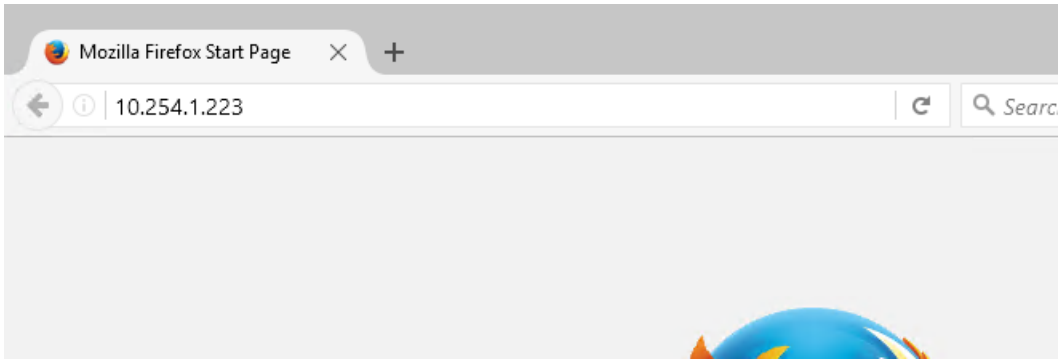

22. You should get redirected by the switch to the ClearPass guest logon page. Do not attempt to logon at this point.

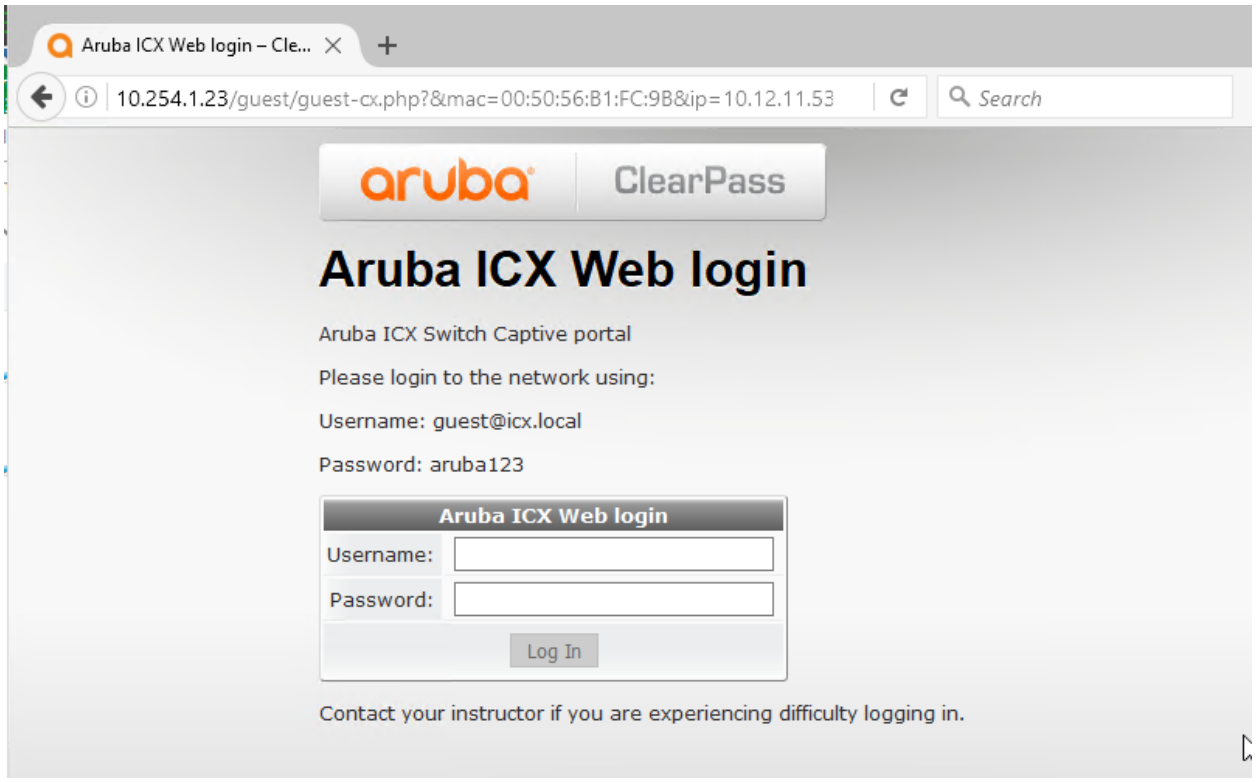

23. Verify that this 'guest-logon' role does not have access to the simulated 'internal' network.

Open another tab in the browser and attempt to open a web page to the Core CX switch on http://10.X.11.1. This should also be redirected to the ClearPass captive portal.

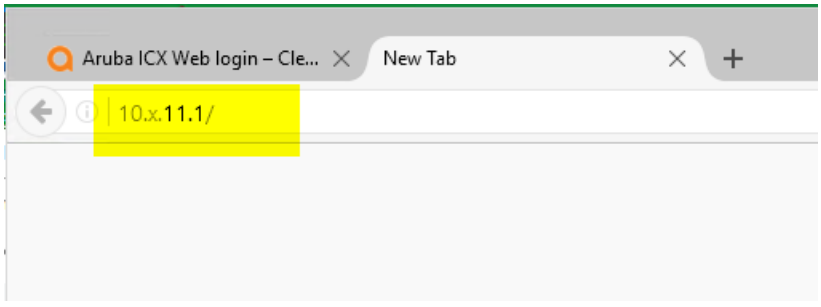

Task 3: Define the Guest Role

Diagram

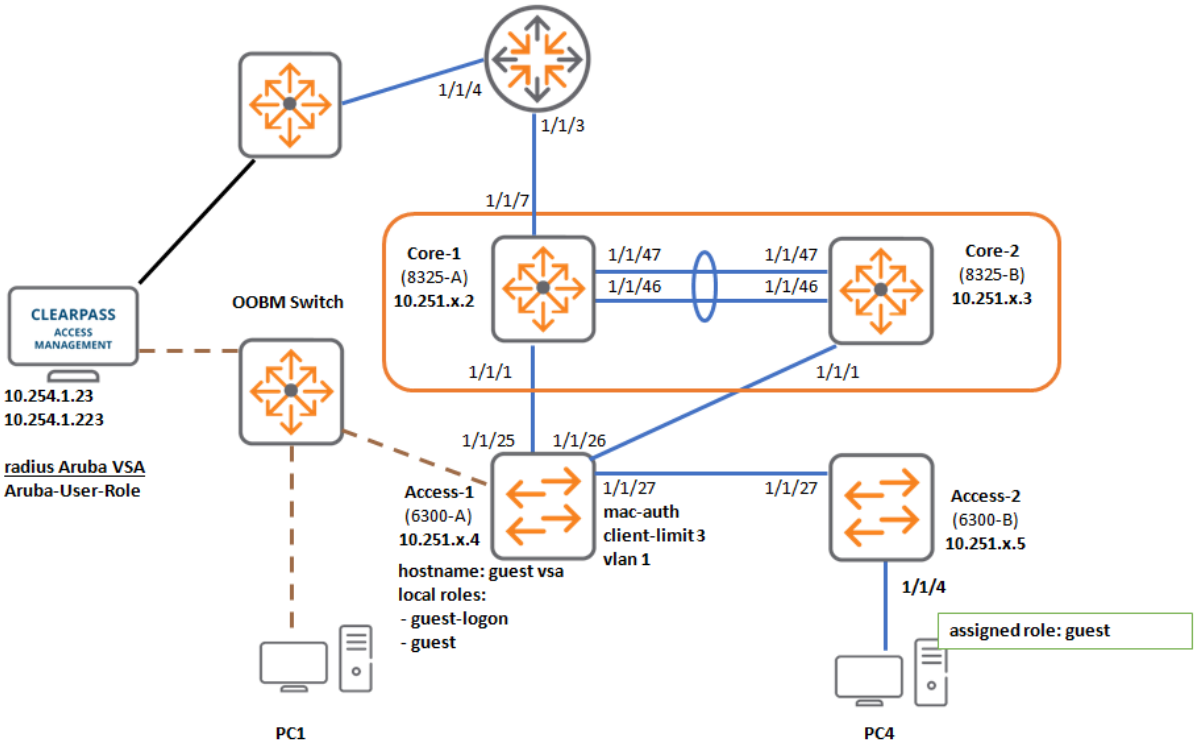

Objectives

In this task, the user-role that will be assigned to authenticated guests will be defined.

In the lab it is simulated that an authenticated guest cannot access the internal resources and has full http/https access to the internet.

In a real deployment, these rules should be adjusted based on the customer requirements.

Since the lab may not have real internet access for the test clients, the host 10.254.1.223 is used to verify access.

The 'internal resources' are defined as '10.0.0.0/8', except for the 10.254.1.0/24 range (since that range is used to test 'internet' access).

Steps

Access1

1. On the Access1 switch, define new class to describe the internal network.

```
ICX-Tx-Access1-guest-vsa(config)# class ip internal-networks
ICX-Tx-Access1-guest-vsa(config-class-ip)# ignore any any
10.254.1.0/255.255.255.0
```

```
ICX-Tx-Access1-guest-vsa(config-class-ip)# match any any 10.0.0.0/255.0.0.0
```

2. Review the configuration.

```
ICX-Tx–Access1-guest-vsa(config-class-ip)# show run cur
class ip internal-networks
    10 ignore any any 10.254.1.0/255.255.255.0
    20 match any any 10.0.0.0/255.0.0.0

ICX-Tx-Access1-guest-vsa(config-class-ip)# exit
```

3. Define a new policy using the new class and some of the previously defined classes.

```
ICX-Tx-Access1-guest-vsa(config)# port-access policy guest
ICX-Tx-Access1-guest-vsa(config-pa-policy)# class ip dhcp-dns-icmp
ICX-Tx-Access1-guest-vsa(config-pa-policy)# class ip cppm-web action cir kbps
1024 cbs 2048 exceed drop
ICX-Tx-Access1-guest-vsa(config-pa-policy)# class ip internal-networks action
drop
ICX-Tx-Access1-guest-vsa(config-pa-policy)# class ip any-web
```

4. Review the configuration.

```
ICX-Tx-Access1-guest-vsa(config-pa-policy)# show run cur
port-access policy guest
    10 class ip dhcp-dns-icmp
    20 class ip cppm-web action cir kbps 1025 cbs 2048 exceed drop
    30 class ip internal-networks action drop
    40 class ip any-web

ICX-Tx-Access1-guest-vsa(config-pa-policy)# exit
```

5. Define a new user-role, associate the policy to the user role. Set the re-authentication period to 5 min (300 seconds).

NOTE: In a real deployment, this timer would be much higher, based on the customer requirements. In this lab, the guest authentication is only valid for 5 minutes to be able to demonstrate the logon/logoff process and repeat the task if required.

```
ICX-Tx-Access1-guest-vsa(config)# port-access role guest
ICX-Tx-Access1-guest-vsa(config-pa-role)# associate policy guest
ICX-Tx-Access1-guest-vsa(config-pa-role)# vlan access 11
ICX-Tx-Access1-guest-vsa(config-pa-role)# reauth-period 300
ICX-Tx-Access1-guest-vsa(config-pa-role)# exit
```

6. On the client PC4, enter the credentials on the guest login page:

username **guest@icx.local**

password **aruba123**

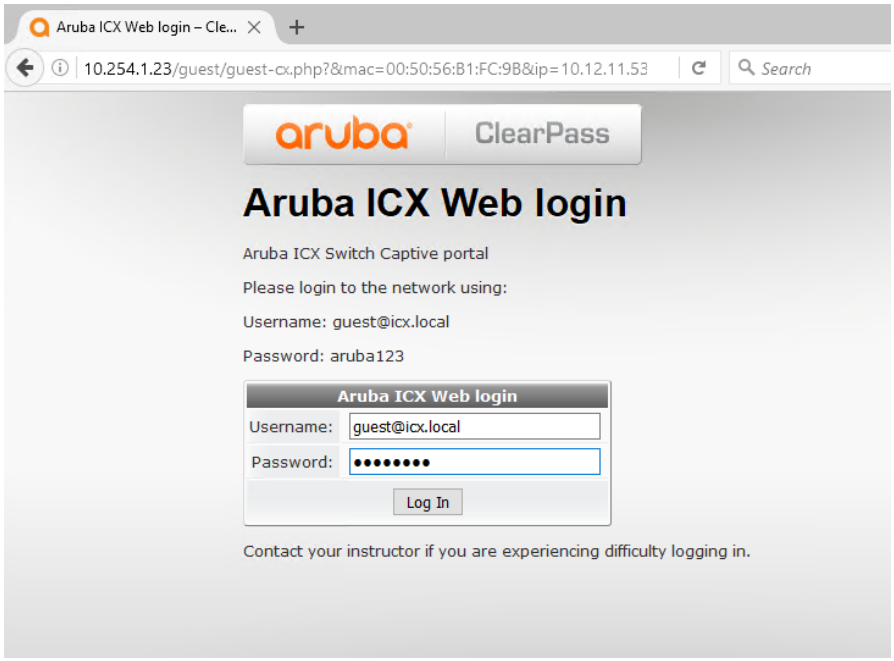

7. ClearPass has been configured to show a logon timer to the 'guest' user. This gives the network some time to complete the CoA and the re-authentication with the 'guest' role.

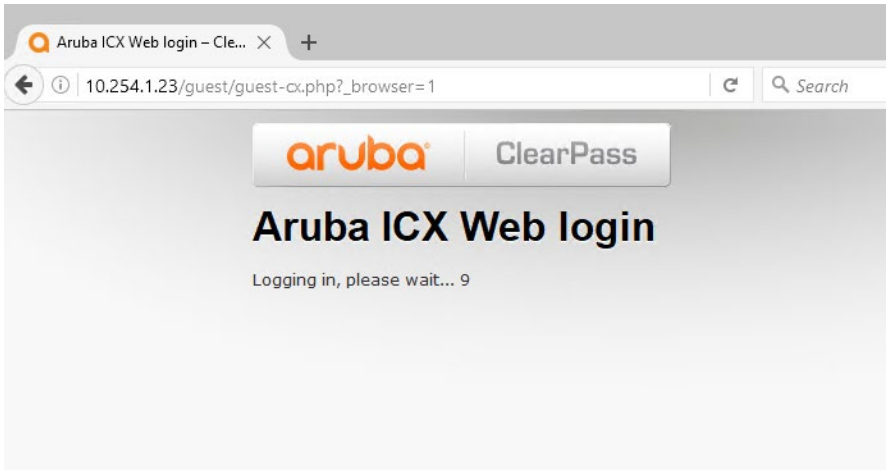

8. After the timer expires, ClearPass will redirect the client browser to a forced welcome page.

9. Verify that the client cannot access other resources in the 10.0.0.0/8 block. Open a new tab in the browser and attempt to connect to the Core switch on http://10.x.11.1/.

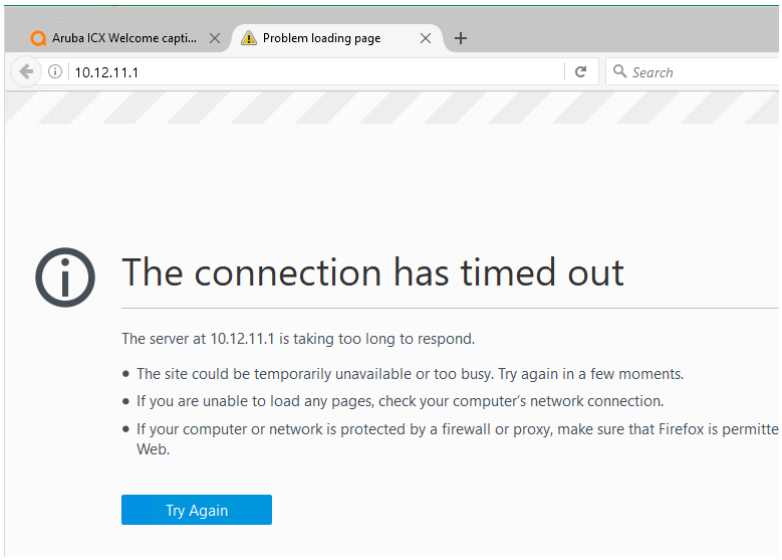

10. On the Access1 switch, review the authenticated devices on port 1/1/27.

```
ICX-Tx-Access1-guest-vsa(config)# show aaa authentication port-access interface
1/1/27 client-status

Port Access Client Status Details

Client 00:50:56:b1:fc:9b, guest@icx.local
============================
  Session Details
  ---------------
    Port         : 1/1/27
    Session Time : 149s

  Authentication Details
  ----------------------
    Status           : mac-auth Authenticated
```

```
  Auth Precedence : dot1x - Not attempted, mac-auth - Authenticated

Authorization Details
---------------------
  Role   : guest
  Status : Applied
```

Q: The switch is using mac-authentication on this port 1/1/27, how is it possible to see the username guest@icx.local in the authentication details?

A: While mac-auth is using the mac address as the authentication username sent from the switch to ClearPass, in the response from ClearPass, there is a 'user-name' AVP in the access-accept to the switch, so the switch has a meaningful username for the mac-auth. By default, this would just be the mac-address.

11. Check ClearPass Access Tracker. Use the Host MAC address filter and apply the last 2 characters of your client MAC address.

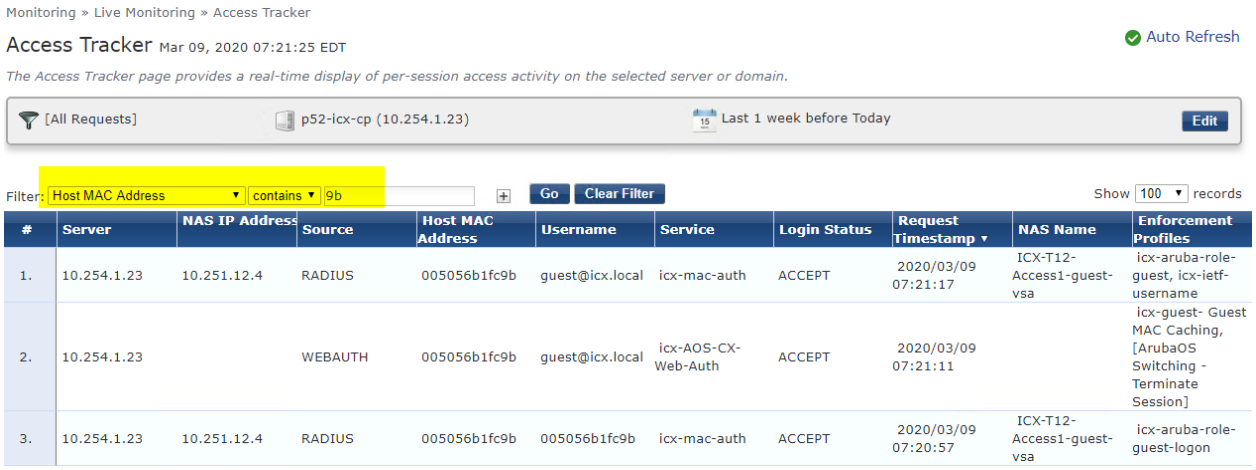

#	Server	NAS IP Address	Source	Host MAC Address	Username	Service	Login Status	Request Timestamp	NAS Name	Enforcement Profiles
1.	10.254.1.23	10.251.12.4	RADIUS	005056b1fc9b	guest@icx.local	icx-mac-auth	ACCEPT	2020/03/09 07:21:17	ICX-T12-Access1-guest-vsa	icx-aruba-role-guest, icx-ietf-username
2.	10.254.1.23		WEBAUTH	005056b1fc9b	guest@icx.local	icx-AOS-CX-Web-Auth	ACCEPT	2020/03/09 07:21:11		icx-guest- Guest MAC Caching, [ArubaOS Switching - Terminate Session]
3.	10.254.1.23	10.251.12.4	RADIUS	005056b1fc9b	005056b1fc9b	icx-mac-auth	ACCEPT	2020/03/09 07:20:57	ICX-T12-Access1-guest-vsa	icx-aruba-role-guest-logon

First, there is the WEBAUTH by the user, then the guest credentials are submitted to ClearPass.

On success authentication, ClearPass updates the endpoint with the correct attributes to mark the MAC address as 'valid'. The success authentication also includes a CoA instruction to trigger re-authentication on the switch.

12. Open this 'WEBAUTH' entry.

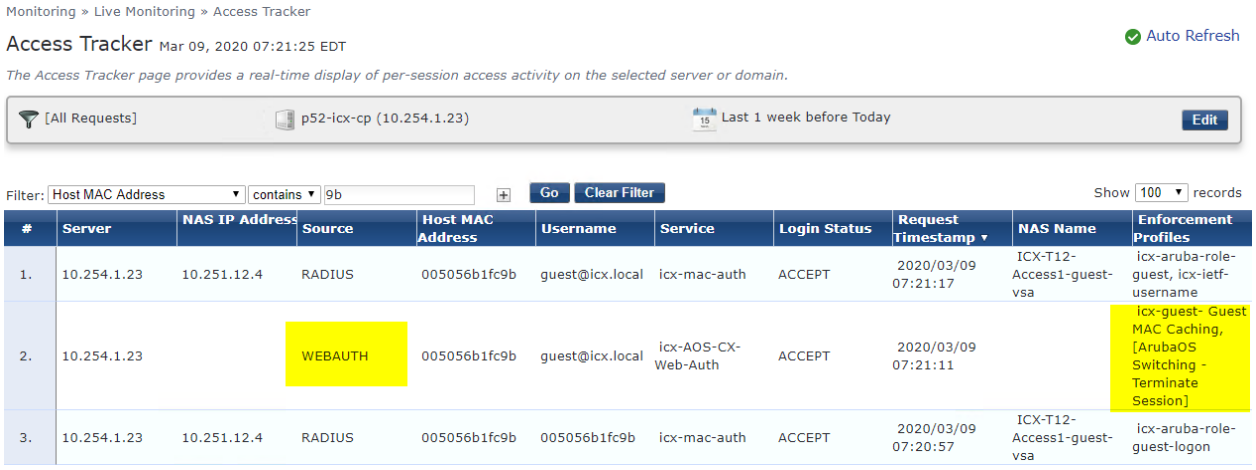

#	Server	NAS IP Address	Source	Host MAC Address	Username	Service	Login Status	Request Timestamp	NAS Name	Enforcement Profiles
1.	10.254.1.23	10.251.12.4	RADIUS	005056b1fc9b	guest@icx.local	icx-mac-auth	ACCEPT	2020/03/09 07:21:17	ICX-T12-Access1-guest-vsa	icx-aruba-role-guest, icx-ietf-username
2.	10.254.1.23		WEBAUTH	005056b1fc9b	guest@icx.local	icx-AOS-CX-Web-Auth	ACCEPT	2020/03/09 07:21:11		icx-guest- Guest MAC Caching, [ArubaOS Switching - Terminate Session]
3.	10.254.1.23	10.251.12.4	RADIUS	005056b1fc9b	005056b1fc9b	icx-mac-auth	ACCEPT	2020/03/09 07:20:57	ICX-T12-Access1-guest-vsa	icx-aruba-role-guest-logon

13. Click '**Output**' and verify it shows attributes that have been written on the endpoint. Close the entry.

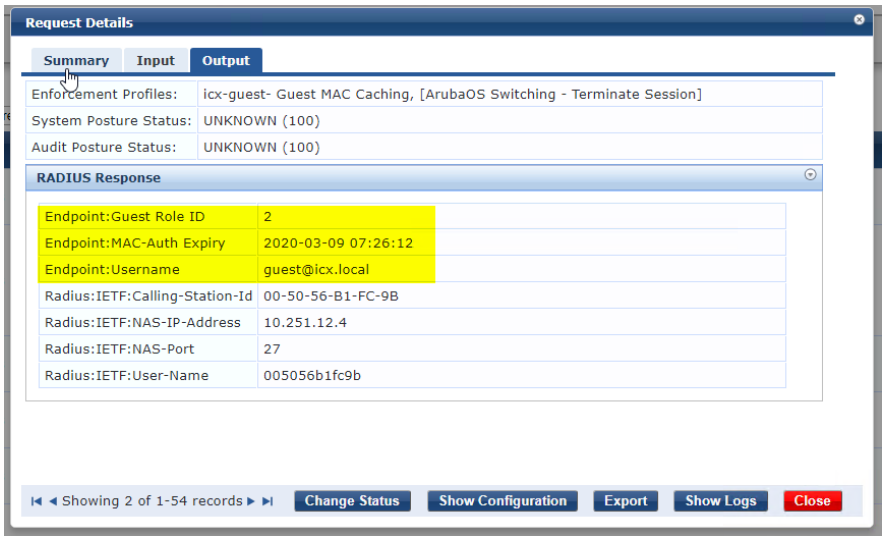

14. A few seconds later, a new RADIUS authentication request will be sent by the switch. This time, the mac-auth will be able to verify the attributes on the endpoint, and the **'guest'** role will be returned.

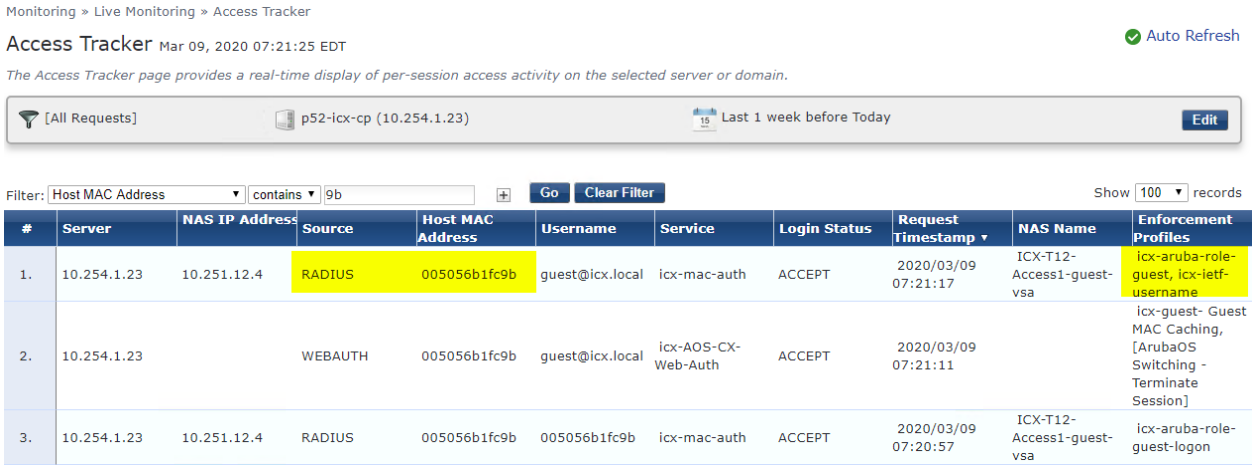

15. Open the entry, verify in the **'Output'** that the RADIUS access-accept contains the Aruba user-role **'guest'** and the **user-name** attribute. It is based on this attribute that the switch is able to show a 'username' for the mac-authenticated devices. Close the entry.

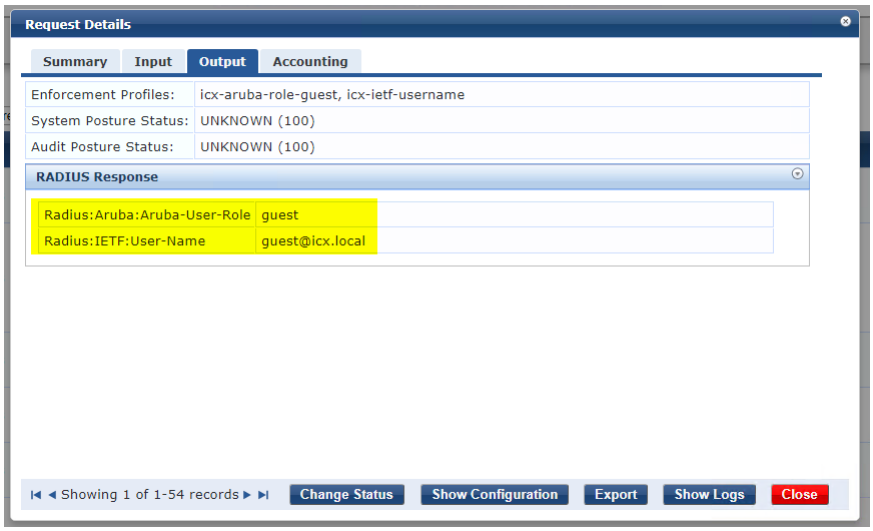

16. Open the **previous** **RADIUS** session, this is the session that returned the 'guest-logon' for the mac-authentication.

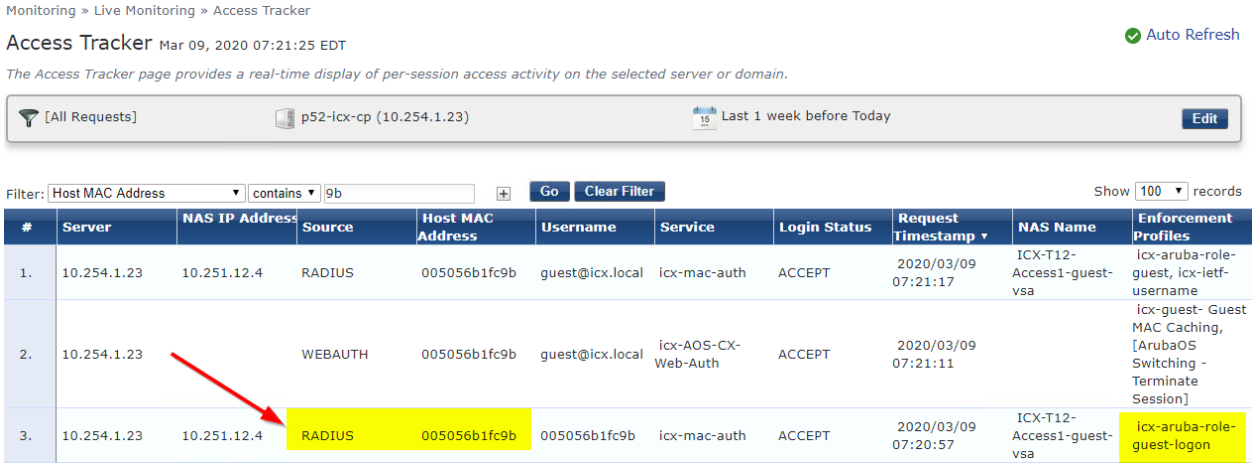

17. This should show a **'RADIUS CoA'** tab. Click on this tab and verify there is a **successful** CoA performed for that session .

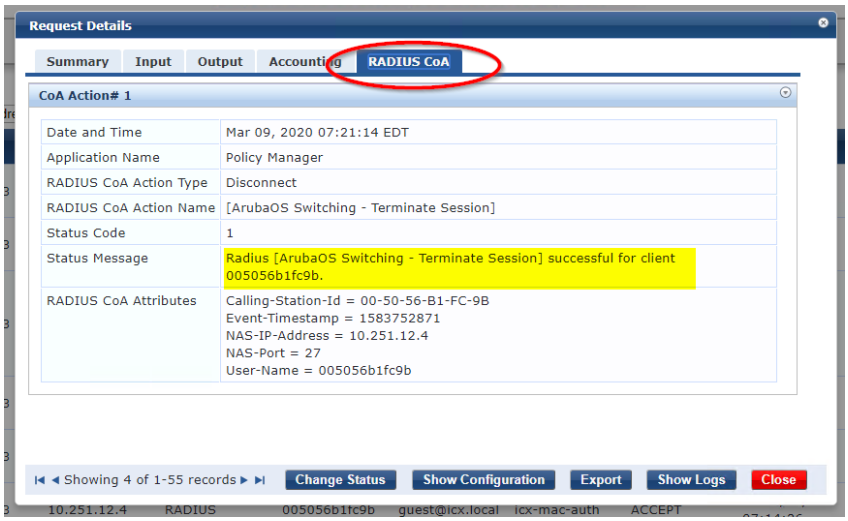

18. Wait about 5 minutes, since the guest user has successfully logged in.

> **NOTE:** In this lab environment, the mac-caching timer has been configured to 5 minutes only, to demonstrate the guest expiration feature as well. In a real deployment, the mac caching timer will be much longer.

19. On the client PC4, navigate to http://10.254.1.223 again (You may just 'refresh' the Welcome page if that is still open). Since the authentication has expired, the client has now to be authenticated again as 'guest-logon', so the captive portal page will be shown again.

This demonstrates the Captive portal configuration process.

Lab cleanup

20. Change the hostname of the switch to the default hostname, so remove the 'guest' and 'vsa' strings from the hostname.

```
ICX-Tx-Access1-guest-vsa(config)# hostname ICX-Tx-Access1
ICX-Tx-Access1(config)#
```

You've completed Lab 15!

6280 AMERICA CENTER DR| SAN JOSE CA 95002
TEL: 408.227.4500 | FAX: 408.227.4550
www.ARUBANETWORKS.com

EDU-ICX-RLABS-v20.211