



Remote Lab Guide



Aruba Advanced Switching Troubleshooting and Solutions

20.41

Lab Guide

December 2020

Aruba Advanced Switching Troubleshooting and Solutions Lab Guide

Copyright

© 2020 Aruba Networks, Inc. AirWave®, Aruba Networks®, Aruba Mobility Management System®, Bluescanner, For Wireless That Works®, Mobile Edge Architecture, People Move. Networks Must Follow., RFProtect, The All Wireless Workplace Is Now Open For Business, and The Mobile Edge Company® are trademarks of Aruba Networks, Inc. All rights reserved. All other trademarks are the property of their respective owners.

Open Source Code

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

http://www.arubanetworks.com/open_source

Legal Notice

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

Warranty

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

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

SKU: EDU-ASTS-RLABS-
v20.41 August 2020

Table of Contents

TABLE OF CONTENTS	I
LAB 0: TESTING LAB CONNECTIVITY (OPTIONAL)	1
OVERVIEW	1
Objectives	1
TASK 1: ARUBA TRAINING LAB ACCESS	2
Objectives	2
Steps	2
TASK 2: TESTING CONNECTIVITY	4
Objectives	4
6300-A and 6300-B	4
PC-1, PC-3 and PC-4	6
NetEdit (via PC-1)	6
CPPM Host (via PC-1)	7
LAB 01: INITIAL SETUP	9
TASK 1: FACTORY RESET THE SWITCHES	10
Objectives	10
Steps	10
TASK 2: LOAD THE INITIAL CONFIGURATION	12
Objectives	12
Steps	12
Load initial configuration to the router devices	12
OOBM Network Access	12
Push Base configuration to your devices	13
TASK 3: VERIFY DEVICE ACCESS	14
Objectives	14
Device IP Address Table	14
Steps	14
Verify in-band connectivity	15
OOBM Connection tests	15
Configuration file formats	15
Table configuration backup and restore on Management PC	16
Table checkpoint configuration (Optional)	17
Course labs 'known good' checkpoints	17
OPTIONAL LAB 02: TROUBLESHOOTING TOOLS	19
REQUIREMENTS	19
OBJECTIVES	19
TASK 1: INTRODUCTION TO SHOW COMMANDS AND SUPPORT FILES	20
Objectives	20
Steps	20
Show commands and support files review	20
Interface statistics	20
Support Logs	21

Accounting logs	21
Event log.....	22
Terminal Monitor.....	22
TASK 2: DIAGNOSTIC TOOLS	24
Objectives	24
Steps.....	24
Explore dropped frames - Jumbo frames	24
Explore dropped frames - VLAN mismatch	25
TASK 3: DEBUGGING OPTIONS	26
Objectives	26
Steps.....	26
Review default logging and logging destination.....	26
Review normal ARP operation.....	26
Review missing ARP reply operation.....	27
TASK 4: TRAFFIC MIRRORING AND PACKET CAPTURES	28
Objectives	28
Steps.....	28
Data plane mirror - Local mirror.....	28
Dataplane mirror - Remote mirror over GRE tunnel	28
Dataplane mirror - TSHARK - Mirror to cpu - Live	30
TSHARK - Mirror to cpu - PCAP.....	31
Class-based mirroring.....	32
TCPDUMP.....	32
LAB 03: MONITORING AND AUTOMATION TOOLS	34
REQUIREMENTS	34
OBJECTIVES	34
SCENARIO	35
NetEdit.....	35
RESTAPI	35
NAE	35
SNMP	35
Lab Diagram	36
TASK 1: NETEDIT	37
Objectives	37
Steps.....	37
Discovery issues.....	37
Diagnostics page	38
NetEdit Insights.....	38
NetEdit Debugging.....	39
Collect NetEdit logs	40
TASK 2: RESTAPI.....	41
Objectives	41
Steps.....	41
Using the built-in Swagger interface	41
Wildcard.....	41
Escape Characters	42
Filtering output.....	42
Configuration Depth.....	42
REST Sessions and clearing REST sessions on AOS-CX.....	43
TASK 3: NAE.....	45
Objectives	45
Steps.....	45

<i>Introduction</i>	45
<i>NAE Script Hierarchy</i>	45
<i>NAE Script Debugging</i>	46
<i>NAE Script deployment</i>	46
<i>Clearing NAE data</i>	46
TASK 4: SNMP CONFIGURATION	48
<i>Objectives</i>	48
<i>Steps</i>	48
<i>SNMP VRF Support</i>	48
<i>SNMPv2 communities</i>	48
<i>Configure SNMPv3 Support for Airwave</i>	49
<i>Troubleshooting SNMP packets</i>	49
<i>SNMP Events</i>	49
<i>Diagnostics SNMP</i>	49
LAB 04: VSX AND LAYER2 TECHNOLOGIES	50
REQUIREMENTS	50
OBJECTIVES	50
SCENARIO	51
VSX	51
<i>VSX LAG - Aruba Gateway</i>	51
<i>VSX LAG - Sw2</i>	51
<i>VSX - DHCP Relay</i>	51
<i>VSX - Virtual Active Gateway</i>	52
<i>Client in VLAN11</i>	52
<i>VSX Best Practices</i>	52
<i>VSX Linkup Delay</i>	52
<i>Spanning-Tree</i>	52
<i>Loop Protection</i>	53
LAB DIAGRAM	54
TASK 1: CONFIGURE VSX	55
<i>Objectives</i>	55
<i>Steps</i>	55
<i>VSX Configuration</i>	55
<i>VSX SVI Active Gateway - SVI 1</i>	55
<i>VSX SVI Active Gateway - SVI 12</i>	56
<i>VSX LAG to Aruba Gateway</i>	57
<i>VSX LAG to Sw2</i>	57
<i>VSX LAG LACP Status</i>	57
<i>VSX DHCP Relay Services</i>	58
<i>Impact of Primary node SVI Shutdown</i>	59
<i>Impact of Primary node SVI Shutdown - DHCP Relay</i>	59
<i>Impact of Primary node SVI Shutdown - Traffic</i>	59
<i>Virtual Active Gateway vs Active Gateway</i>	60
<i>Verify the configuration and troubleshoot the Layer2 connection</i>	61
TASK 2: APPLY SOME VSX BEST PRACTICES	63
<i>Objectives</i>	63
<i>Steps</i>	63
<i>MTU Configuration</i>	64
<i>Uplink delay-timer</i>	64
<i>Validate the linkup delay-timer</i>	64
<i>Review impact on the VSX LAG interfaces</i>	65
<i>Linkup-delay exclusion</i>	65

VSX LAG Hashing algorithm	67
TASK 3: CONFIGURING STP WITH VSX	68
Objectives	68
Steps.....	68
Configure STP	68
VSX STP BPDU Transmissions	68
STP features - Root-guard.....	69
Restore the configuration.....	69
VSX configuration consistency check.....	69
Configure and review STP on the Access switch Sw2.	70
STP Loop-guard	70
Validating STP features - Loop-guard	70
Configure the access ports	71
Topology Change Notifications.....	71
BPDU-Guard configuration	72
TASK 4: CONFIGURING LOOP-PROTECTION IN THE ACCESS LAYER	74
Objectives	74
Steps.....	74
Loop Setup	74
Configure the network loop without Loop Protect	74
Enable the loop.....	74
Loop Protection configuration	74
Restore the configuration.....	75
LAB 05: REVIEW OF LAYER3 BASICS	76
REQUIREMENTS	76
OBJECTIVES	76
SCENARIO	77
OSPF	77
VSX Active Forwarding.....	77
OSPF Key-chain and tuning (Optional)	77
Policy Based Routing (PBR).....	77
PBR Path selection.....	78
TASK 1: CONFIGURE VSX WITH OSPF.....	80
Objectives	80
Steps.....	80
OSPF Configuration on Agg1 and Agg2.....	80
Lab Diagram - Layer3 Topology	80
TASK 2: DEBUGGING OSPF.....	81
Objectives	81
Lab Diagram - Layer3 Topology	81
Steps.....	81
Adjacency Debugging.....	81
Link Troubleshooting	82
Review the normal OSPF adjacency setup	83
Invalid Timers	83
Invalid authentication	84
Router-id Conflict on Peer	85
OSPF Point to Point networks	85
TASK 3: VSX ACTIVE FORWARDING	86
Objectives	86
Lab Diagram - Layer3.....	86
Steps.....	86

Active Forwarding - Understand the need for Active Forwarding (optional)	88
Testing	88
Active Forwarding	89
OPTIONAL TASK 4: OSPF KEY-CHAIN AND MAX-METRIC ON STARTUP	91
Objectives	91
Steps	91
Configuration	91
Activation	92
Migration - Attempt 1	92
TCPDUMP on Agg1	92
Migration - Think about the migration steps	93
Migration - Attempt 2	93
Max-metric on Startup	94
TASK 5: POLICY BASED ROUTING (PBR)	96
Objectives	96
Steps	96
Prepare the setup	96
Option 1: Manual Configuration	96
Option2: Saved Script configuration	97
OSPF best path verification	97
Test assigned link and failover without PBR applied	98
Test link failover	98
Enhance the failover detection for the SVI using BFD	98
Test link failover	99
Policy Based Routing Configuration	99
Verification	100
Verify the 'other' traffic	100
Cleanup	100
Required: Post-lab checkpoint	100
LAB 06: BGP	102
REQUIREMENTS	102
SCENARIO	103
iBGP	103
eBGP	103
BGP Advertisements	103
BGP Peer Groups	104
Controlling Transit Traffic	104
Outbound Traffic Route Control	104
Inbound Traffic route control	105
AS Number filtering	106
iBGP peering	106
Lab Diagram	106
TASK 1: PREPARE THE SETUP	108
Objectives	108
Steps	108
Router Configurations	108
TASK 2: CONFIGURE IBGP AND EBGp PEERING	109
Objectives	109
Steps	109
Troubleshooting iBGP peering	109
Incorrect iBGP session IP address	109
Configure iBGP peering between the other switches.	110

Configure eBGP peering with routerA and routerB.....	110
Troubleshoot remote AS number.....	110
Next hop self.....	111
Solution: configure next-hop self.....	111
Solution: Reset the BGP peering.....	112
eBGP between Sw1 and Sw2 (acting as External BGP Partner)	112
Advertise default route into OSPF.....	114
Advertise a route to eBGP peer.....	114
TASK 3: USING BGP PEER GROUPS	117
Objectives.....	117
Steps.....	117
Prepare the Peer-group.....	117
Migrate the existing peer to the peer-group.....	117
TASK 4: CONTROLLING TRANSIT TRAFFIC	118
Objectives.....	118
Steps.....	118
Prevent transit AS.....	118
Allow selective transit AS.....	119
TASK 5: OUTBOUND TRAFFIC ROUTE CONTROL	120
Objectives.....	120
Steps.....	120
Apply default local preference	120
Apply selective local preference	122
AS multipath-relax	122
Optional: Test link failure and route failover	123
Using weight locally on a router.....	123
Soft reconfiguration inbound.....	125
Soft reconfiguration inbound impact on route refresh.....	126
TASK 6: INBOUND TRAFFIC ROUTE CONTROL	128
Objectives.....	128
Lab Diagram	128
Steps.....	130
Apply AS-Path length for outbound advertisements.....	130
Use MED for inbound route preference	131
Prepare the link between Agg2 and Sw2.....	131
Apply the MED values	132
Using Community Attributes for route control	133
Correct the iBGP peer configuration.....	134
OPTIONAL TASK 7: AS PATH LIST EXCLUSION	136
Objectives.....	136
Steps.....	136
OPTIONAL TASK 8: USING A ROUTE REFLECTOR FOR IBGP.....	137
Objectives.....	137
Steps.....	137
Break the iBGP full mesh topology.....	137
Review the result of the broken topology.....	137
Configure the route reflector	138
LAB 07: ROUTE REDISTRIBUTION	139
REQUIREMENTS	139
SCENARIO	140
Setup	140
Static route into OSPF	140

<i>Prefix Lists</i>	140
<i>Route Redistribution between OSPF and BGP</i>	141
<i>Single Link Redistribution</i>	141
<i>Dual Link Route Redistribution</i>	141
TASK 1: LOAD THE START CONFIGURATIONS	142
<i>Objectives</i>	142
<i>Steps</i>	142
TASK 2: ROUTE REDISTRIBUTION OF STATIC ROUTES INTO OSPF	143
<i>Objectives</i>	143
<i>Steps</i>	144
<i>Redistribute static routes into OSPF</i>	144
<i>Cost Manipulation of the redistributed routes</i>	144
<i>Granular control of the redistribution using route-maps</i>	145
<i>OSPF External Metric Type1</i>	146
<i>Test failover impact on the route metrics</i>	147
<i>Understanding Prefix Lists</i>	147
TASK 3: ROUTE REDISTRIBUTION BETWEEN OSPF AND BGP	149
<i>Introduction</i>	149
<i>Objectives</i>	149
<i>Lab Diagram</i>	150
<i>Steps</i>	151
<i>Prepare the setup</i>	151
<i>Verify current routes and OSPF neighbors</i>	151
<i>Configure eBGP peering between Agg1 and routerA</i>	151
<i>Redistribute local OSPF routes into BGP</i>	152
<i>Redistribute BGP routes into OSPF</i>	153
TASK 4: ROUTE REDISTRIBUTION WITH MULTIPLE LINKS	154
<i>Introduction</i>	154
<i>Objectives</i>	154
<i>Lab Diagram</i>	155
<i>Configure second link: Attempt 1</i>	155
<i>Review routing before the second link redistribution is configured</i>	156
<i>Agg2: Redistribute BGP into OSPF</i>	156
<i>Review routing after the redistribution</i>	156
<i>Configure Second Link: Attempt 2 using route tags</i>	157
<i>Customer Redistribution on Agg2</i>	157
<i>Partner network configuration</i>	158
LAB 08: VRF AND ROUTE LEAKING	159
OBJECTIVES	159
REQUIREMENTS	159
SCENARIO	160
<i>Static Route Leaking</i>	160
<i>Dynamic Route Leaking using MP-BGP</i>	160
<i>Route-map based import</i>	160
<i>Lab Diagram</i>	161
TASK 1: PREPARE THE BASE CONFIGURATION	162
<i>Objectives</i>	162
<i>Steps</i>	162
<i>IOT VRF Configuration</i>	162
TASK 2: EXPLORE THE ENVIRONMENT	163
<i>Objectives</i>	163
<i>Steps</i>	163

Connectivity Checks	163
Verify VRF Route Isolation	163
TASK 3: STATIC ROUTE LEAKING	164
Objectives	164
Steps.....	164
Manually Leak Directly Connected Interface	164
Make the leaked route available inside the VRF branch	165
Debug on Sw2	165
Prepare for MP-BGP leaking	166
TASK 4: MP-BGP	167
Objectives	167
Steps.....	167
Configure Route Distinguisher and import routes into BGP routing table.....	167
Route Targets	169
Simple method.....	169
Verify connectivity.....	170
Route-map based RT	171
LAB 09: MULTICAST	173
OBJECTIVES	173
REQUIREMENTS	173
Introduction	173
SCENARIO	174
Setup	174
Multicast routing.....	174
Lab Diagram	175
TASK 1: LOAD THE START CONFIGURATIONS	176
Objectives	176
Steps.....	176
Introduction	176
Explore the environment.....	176
TASK 2: CONFIGURE PIM SM	178
Objectives	178
Steps.....	178
Configure the C-BSR.....	178
Enable PIM-SM on all routing devices.....	178
Configure Sw1 as C-RP.....	179
TASK 3: CONFIGURE IGMP AND VERIFY MULTICAST OPERATION	180
Objectives	180
Steps.....	180
IGMP Querier.....	180
IGMP Snooping	181
Test Multicast and verify status on Sw1	181
Test without multicast receiver	182
Test with multicast receiver	183
Review the status of the multicast route	184
Test with mtrace	186
Optional: PIM Diagnostics Dump.....	187
TASK 4: CONFIGURE DISTRIBUTED RP	188
Objectives	188
Steps.....	188
Configure regional RP	188
Verify the registration traffic on the regional RP	188

OPTIONAL TASK 5: UNDERSTAND THE MAC - IP MULTICAST RELATION	189
Objectives	189
Steps.....	189
Explore the Multicast IP to Ethernet mapping.....	189
TASK 6: VSX MULTICAST FAILOVER	190
Objectives	190
Steps.....	190
VSX Active Active Configuration.....	190
Understand the VSX Multicast Routing failover.....	190
Verify the failover	192
TASK 7: IGMP ACL	193
Objectives	193
Steps.....	193
IGMP Querier ACL.....	193
IGMP Snooping ACL	193
IGMP debug.....	194
LAB 10: QUALITY OF SERVICE.....	195
OBJECTIVES	195
REQUIREMENTS	195
SCENARIO	196
Lab Diagram	197
TASK 1: QOS POLICIES.....	198
Objectives	198
Steps.....	198
Review default queuing and QOS maps.....	198
Configure QOS Policy on Sw2.....	199
Verify the Remark - last hop switch	199
Default Queue and Trust modes.....	200
Verify the impact of the QOS trust modes.	201
On Sw2, you will now review the local traffic handling.	201
TASK 2: RATE LIMITER AND TRAFFIC SHAPING	203
Objectives	203
Steps.....	203
Inbound Rate Limiters.....	203
Define the Rate Limiter on the Aggregation VSX	203
Troubleshoot the VSX-sync.....	203
Verify the rate limiter.....	204
Outbound Traffic shaping	205
Rate limit and burst size	205
LAB 11: DYNAMIC SEGMENTATION	206
OBJECTIVES	206
REQUIREMENTS	206
SCENARIO	207
Lab Diagram	210
TASK 1: PREPARE THE CONFIGURATION	211
Objectives	211
Steps.....	211
Configuration	211
Table with RADIUS test user accounts.....	211
TASK 2: CONFIGURE BASIC 802.1X AUTHENTICATION.....	212
Objectives	212

Steps.....	212
Basic RADIUS and 802.1X configuration	212
Statistics EAP and RADIUS.....	214
Review the debug output for a successful authentication.....	215
TASK 3: RADIUS TROUBLESHOOTING	216
Objectives	216
Steps.....	216
Invalid Aruba-User-Role	216
Cached re-authentication.....	217
Verify the operation without cached re-authentication.....	217
Critical Role for Authentication	218
Configure the switch with cached re-authentication	218
Understand the need for RADIUS tracking.....	220
Recover the link and verify the operation	221
TASK 4: ONBOARDING PRECEDENCE ORDER	222
Objectives	222
Steps.....	222
Device Profile with Local MAC Authentication (LMA)	222
Device Profile Fallback role	223
Device Profile with AAA enabled on the port: Onboarding Precedence	224
Reject RADIUS Response.....	224
Optional: Device Profile with RADIUS Failure	225
TASK 5: USER-BASED TUNNELING	226
Objectives	226
Steps.....	226
Configuring UBT	226
Review Normal UBT Authentication	227
Invalid Gateway Role.....	228
MTU for Tunneled users.....	228
Configure Jumbo support	230
Default role on the Aruba Gateway.....	231
TASK 6: USER-BASED TUNNELING QOS.....	232
Objectives	232
Steps.....	232
Apply a QOS Policy to the Switch user role.....	232
Optional: Trace traffic and compare inner DSCP to outer GRE DSCP markings.....	233
TASK 7: DOWNLOADABLE USER ROLES	235
Objectives	235
Steps.....	235
ClearPass DUR account configuration	235
ClearPass Certificate Validation	235
Install the Trust Anchor certificate on the Switch.....	238
Review normal DUR authentication and role.....	240
LAB 12: NETWORK SECURITY FEATURES	241
OBJECTIVES	241
REQUIREMENTS	241
SCENARIO	242
TCAM Resource Usage.....	242
Hitcounts.....	242
Control Plane Security.....	242
Control Plane Policing	243
DHCP Snooping	243

ARP Inspection	243
Lab Diagram	244
TASK 1: ACL AND RESOURCE USAGE	245
Objectives	245
Steps	245
Configuration	245
Prepare the a routed ACL	245
Update the ACL using object-groups	245
Hitcounts and Statistics	246
Optional: Unique Policy per interface for unique counters	247
TASK 2: CONTROL PLANE ACL	248
Objectives	248
Steps	248
COPP ACL on VRF mgmt	248
Prepare for failure	249
COPP ACL on VRF default	249
TASK 3: CONTROL PLANE POLICING	251
Objectives	251
Steps	251
VRF mgmt behavior	252
OPTIONAL TASK 4: MONITORING CONTROL PLANE WITH NAE	253
Objectives	253
Steps	253
TASK 5: DHCP SNOOPING	254
Objectives	254
Steps	254
Optional: Authorized Server	255
DHCP Bindings Database	255
TASK 6: ARP INSPECTION	256
Objectives	256
Steps	256
Configure ARP Inspection	256
Optional: Verify IP Binding database dependency	257
Verify ARP Inspection filtering	258
Endpoint with static IP	258
LAB 13: IPV6	260
OBJECTIVES	260
REQUIREMENTS	260
SCENARIO	261
OSPFv3	261
SLAAC	261
DHCPv6 - Stateless	261
DHCPv6 - Stateful	262
Lab Diagram	263
TASK 1: CONFIGURE IPV6 LINKS AND OSPFv3	264
Objectives	264
Steps	264
Agg1 - Agg2 Link	264
OSPFv3 on Agg1 and Agg2	265
Configure Sw1 as OSPFv3 router	265
TASK 2: CONFIGURE IPV6 ON CLIENT SUBNET USING SLAAC	267
Objectives	267

Steps.....	267
Aggregation configuration.....	267
Client verification	267
TASK 3: CONFIGURE IPV6 ON CLIENT SUBNET USING DHCP	268
Objectives.....	268
Steps.....	268
Configure PC4.....	268
Configure DHCPv6 Server.....	269
Configure SVI12 for Stateless DHCPv6	269
Verify Stateless DHCPv6.....	269
Configure SVI 12 for Stateful DHCPv6.....	270
Verify Stateful DHCPv6	270
LAB 14: TROUBLE TICKETS HINTS AND SOLUTIONS.....	272
TROUBLE TICKET HINTS AND SOLUTIONS.....	273
Trouble Ticket 01.01.....	273
Trouble Ticket 01.02.....	273
Trouble Ticket 02.01.....	274
Trouble Ticket 03.01.....	274
Trouble Ticket 04.01.....	274
Trouble Ticket 05.01.....	275
Trouble Ticket 06.01.....	276
Trouble Ticket 07.01.....	277
Trouble Ticket 08.01.....	278
Trouble Ticket 08.02.....	279
Trouble Ticket 08.03.....	280
Trouble Ticket 08.04.....	281
Trouble Ticket 08.05.....	281
Trouble Ticket 08.06.....	281
Trouble Ticket 08.07.....	282
Trouble Ticket 09.01 - Routing and Multicast.....	283
LAB 14: TROUBLE TICKETS	284
TASK 1: PREPARE THE TROUBLESHOOTING SETUP	285
Introduction	285
TROUBLE TICKET SETUP 01.....	287
Steps.....	287
Trouble Ticket 01.01.....	288
Trouble Ticket 01.02.....	289
TROUBLE TICKET SETUP 02.....	290
Steps.....	290
Trouble Ticket 02.01.....	290
TROUBLE TICKET SETUP 03.....	291
Steps.....	291
Trouble Ticket 03.01.....	291
TROUBLE TICKET SETUP 04.....	292
Steps.....	292
Trouble Ticket 04.01.....	292
TROUBLE TICKET SETUP 05.....	293
Steps.....	293
Trouble Ticket 05.01.....	293
TROUBLE TICKET SETUP 06.....	294
Steps.....	294

<i>Trouble Ticket 06.01</i>	294
TRouble TICKET SETUP 07	295
Steps	295
<i>Trouble Ticket 07.01</i>	295
TRouble TICKET SETUP 08	296
Steps	296
<i>Trouble Ticket 08.01</i>	296
<i>Trouble Ticket 08.02</i>	296
<i>Trouble Ticket 08.03</i>	296
<i>Trouble Ticket 08.04</i>	297
<i>Trouble Ticket 08.05</i>	297
<i>Trouble Ticket 08.06</i>	297
<i>Trouble Ticket 08.07</i>	298
TRouble TICKET SETUP 09	299
Steps	299
<i>Trouble Ticket 09.01</i>	299

Aruba Training-Confidential

Lab 0: Testing Lab Connectivity (optional)

Overview

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

- **PC-1:** This client is used for device management, traffic analysis, connectivity testing, accessing the Web-UI of NetEdit and your switches and accessing the CLI over SSH of switches.
- **PC-3:** This client is used connectivity testing.
- **PC-4:** This client is used connectivity testing.
- **6300-A switch:** This is Sw1, primarily used for Layer3 functions.
- **6300-B switch:** This is Sw2, primarily used for Layer2 functions.
- **8325-A switch:** This is Agg1, acting as the first aggregation switch.
- **8325-B switch:** This is Agg2, acting as the second aggregation switch.
- **OOBM switch:** You have NO access to this switch.
- **Server Switch.** You have NO access to this switch.
- **NetEdit:** You will access this device over a HTTPS session via PC-1
- **Windows Server:** You have NO access to this server; it is used as DNS and has a DHCP server function.
- **ClearPass server:** You have read-only access to this server, it will be used as a AAA server.

Objectives

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

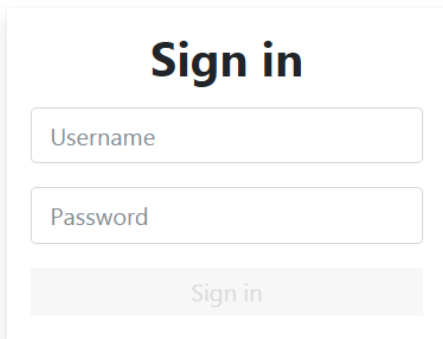
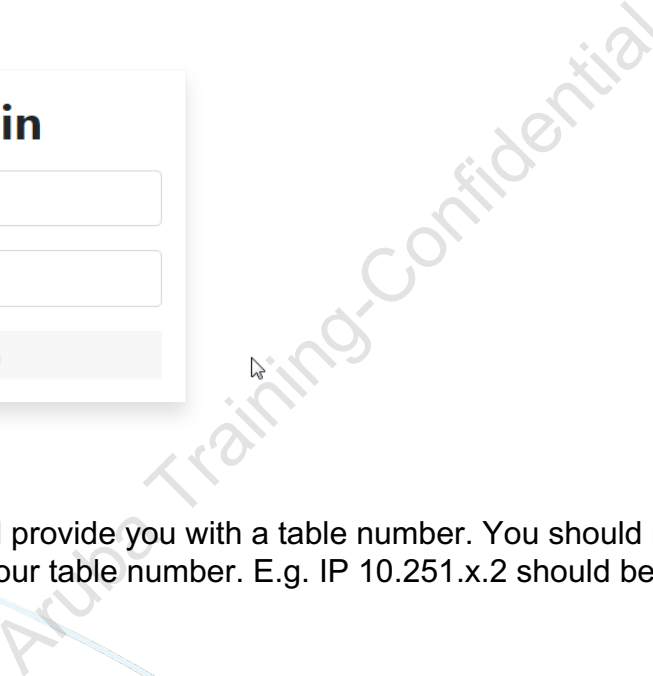
Task 1: Aruba Training Lab Access

Objectives

In this task you will check that you have connectivity to the remote lab and can successfully login. This will ensure that you have access to your remote lab equipment during this training.

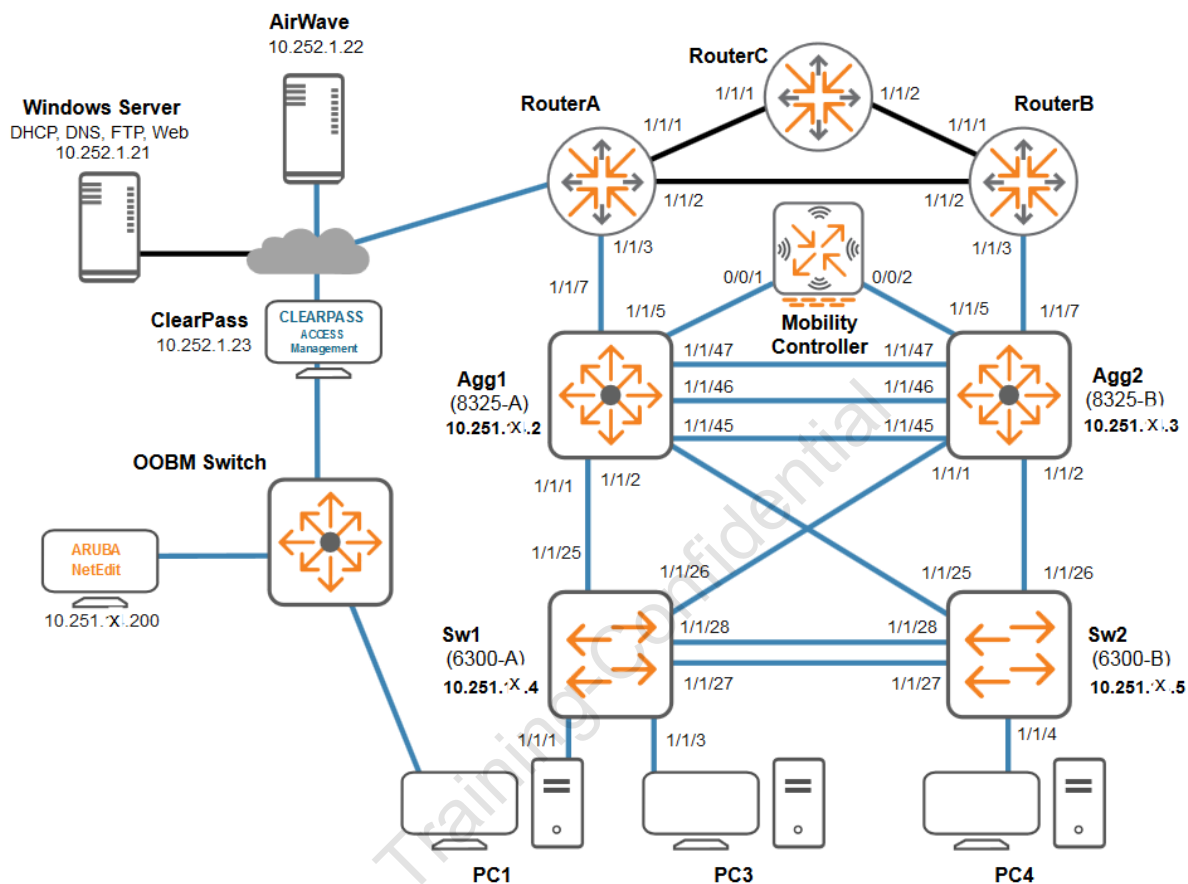
Steps

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



A screenshot of a web form titled "Sign in". It contains two input fields: "Username" and "Password". Below these fields is a "Sign in" button. A mouse cursor is pointing at the button. The form is set against a light gray background.

3. Your instructor will provide you with a table number. You should replace 'x' in the lab instructions with your table number. E.g. IP 10.251.x.2 should be updated based on your table number.



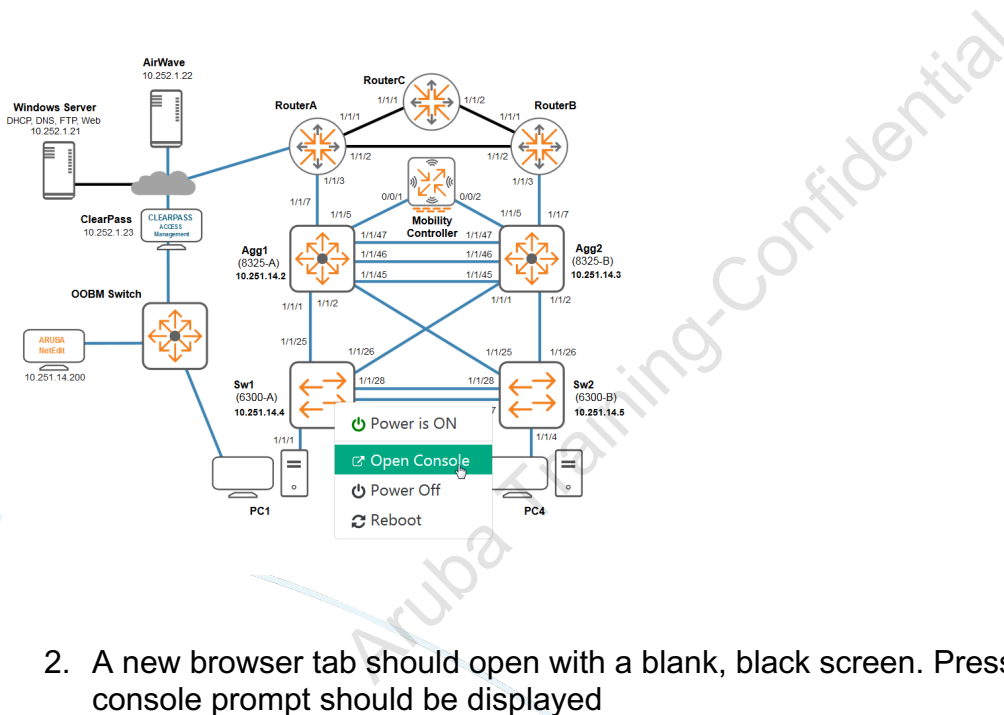
Task 2: Testing Connectivity

Objectives

To test connectivity and authentication credentials for each of the devices. Working from the Aruba Training Lab diagram, you will connect to and log into the Access switches and your client PCs.

6300-A and 6300-B

1. To connect to the console of the 6300-A switch, right-click on the icon in the lab diagram and select **"Open Console."**



2. A new browser tab should open with a blank, black screen. Press 'Enter' a few times. The console prompt should be displayed

```
(C) Copyright 2017-2020 Hewlett Packard Enterprise Development LP

RESTRICTED RIGHTS LEGEND
Confidential computer software. Valid license from Hewlett Packard Enterprise
Development LP required for possession, use or copying. Consistent with FAR
12.211 and 12.212, Commercial Computer Software, Computer Software
Documentation, and Technical Data for Commercial Items are licensed to the
U.S. Government under vendor's standard commercial license.

We'd like to keep you up to date about:
* Software feature updates
* New product announcements
* Special events
Please register your products now at: https://asp.arubanetworks.com

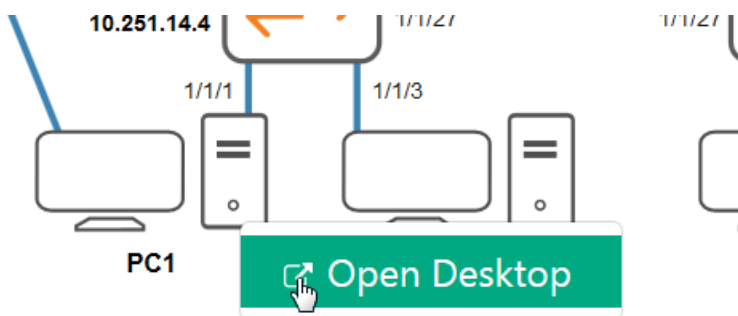
6300 login: 
```

3. Repeat the steps 1 to 2 on Sw2 (6300-B), Agg1 (8325A) and Agg2(8325B).

Aruba Training-Confidential

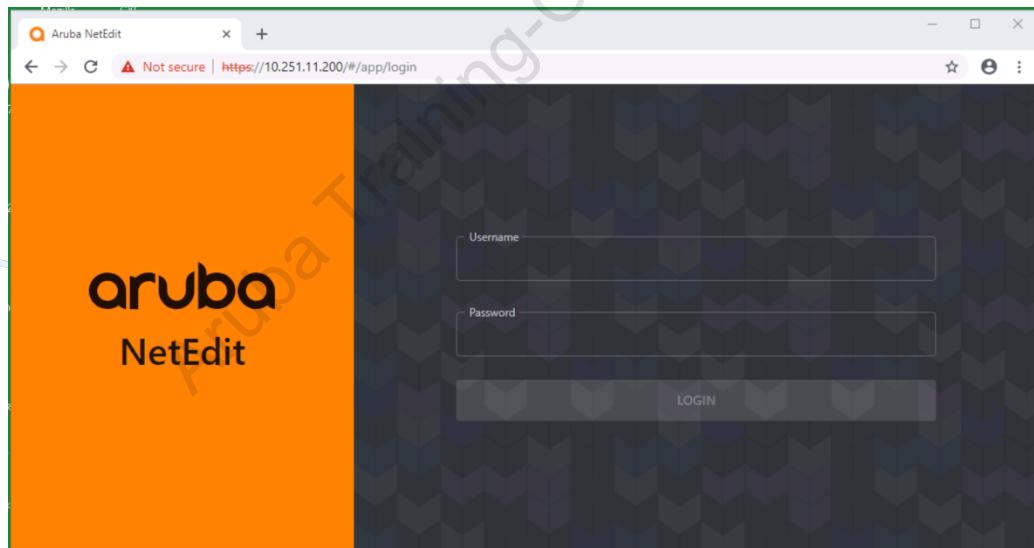
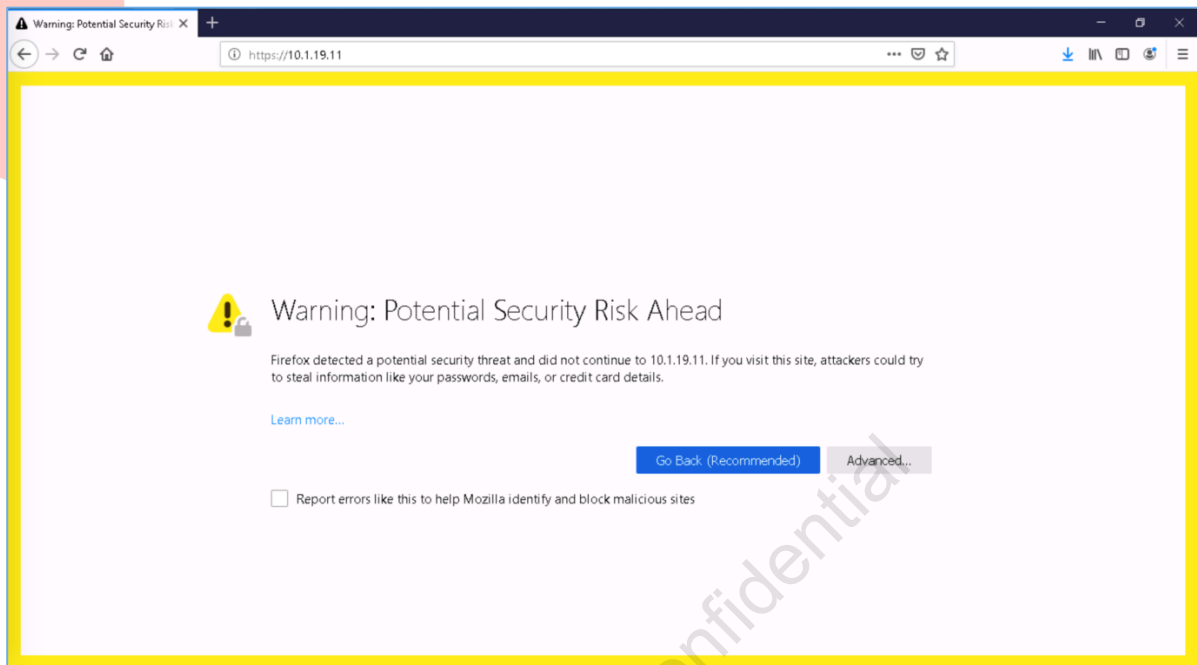
PC-1, PC-3 and PC-4

4. To access the desktop PC-1, right-click on the icon in the lab diagram and select “**Open Desktop**.”



NetEdit (via PC-1)

1. On PC1, open a browser and type the NetEdit IP address in the URL field (**10.251.x.200**) then hit **[Enter]**. You will be presented a security certificate warning. Each table has its own NetEdit host.
2. Accept the warning. You will see the NetEdit login page.



CPPM Host (via PC-1)

3. On PC1, open a browser and type the ClearPass Policy Manager IP address in the URL field (**10.252.1.23**) then hit **[Enter]**. There is 1 common ClearPass server for the tables in the class.
4. You will be presented a security certificate warning.

5. Accept the warning. You will see the login page right after.
6. You have read-only access with the credentials readonly / readonly .

You have completed this Lab!

Aruba Training-Confidential

Lab 01: Initial Setup

In this lab you will perform the initial switch setup.

- Reset the switches to factory default
- Apply the basic initial configuration

Aruba Training-Confidential

Task 1: Factory Reset the Switches

Objectives

- Erase the configuration and any previous settings from the switches
- Verify and correct the firmware version on the switches if necessary

Steps

1. Open a console connection to Agg1 (8325A). Login with username 'admin'

Q: In which scenario would the switch prompt you to change the admin password?

A: When it has just completed a zeroize procedure. This means the switch is completely factory default.

2. When you are prompted to change the admin password, set it to 'aruba123'.

NOTE: If you were not prompted to change the admin password, perform a zeroize of the switch settings. The switch will reboot and then you will be prompted to set the admin password.

3. Since the training lab is shared with other switching courses, the firmware may be different for these other courses. Verify the current firmware version.
4. In case the active version is not 10.05.0021:
 - a. Verify the current images in the flash

NOTE: If you don't have 10.05.0021 on the switch, notify your instructor.

- b. Set the correct image (primary or secondary based on the previous output) as the default boot to make 10.05.0021 the default boot image.
- c. Verify the correct version has been set with the `show image` command and reboot the system.

NOTE: If you would perform the 'erase all zeroize' procedure at a later point, remember to set the default boot again. This setting is cleared by the zeroize procedure.

5. Repeat the previous 4 steps for Agg2(8325A)

6. Repeat the previous 4 steps for Sw1 (6300A)

7. Repeat the previous 4 steps for Sw2 (6300B)

The expected version of the switches for the ASTS training labs is 10.05.0021

Aruba Training-Confidential

Task 2: Load the Initial Configuration

Objectives

- Load the initial configuration to the router virtual devices
- Load an example initial configuration for the switches

Steps

Load initial configuration to the router devices

1. Use the training lab interface to access the console of RouterA, login using admin/aruba123
2. Review the list of checkpoints

```
RouterA# show checkpoint list
```

3. Copy the checkpoint 'asts-base-l3' to the running configuration, save the configuration and reboot.

```
copy checkpoint asts-base-l3 run
write mem
boot system
```

4. Repeat steps 1-3 for RouterB and RouterC.

OOBM Network Access

5. Open a console connection to Agg1 and check the IP address of the mgmt interface.

Q: What is the OOBM IP address?

A: By default DHCP is enabled on the OOBM port, this could be used for Zero-Touch Provisioning. Since there is no DHCP server on the Out-Of-Band Management network in the lab, the switch will not have a management IP address.

6. On Agg1, configure a static OOBM IP 10.251.x.2/24 on the 'mgmt' interface, set 10.251.x.254 as the mgmt default gateway.

NOTE: Replace x with your assigned table number.

7. Make a checkpoint with name 'oobm' (all lowercase).

IMPORTANT: Some scripts that are required later in the training labs depend on this checkpoint. Make sure you have created this checkpoint **oobm**.

8. Review the checkpoint list and verify that checkpoint **oobm** is listed.
9. Open console connections to Agg2, Sw1 and Sw2 and configure the static management IP and GW.

Device	OOBM IP
8325A (Agg1)	10.251.x.2/24
8325B (Agg2)	10.251.x.3/24
6300A (Sw1)	10.251.x.4/24
6300B (Sw2)	10.251.x.5/24

10. Open a connection to PC1, open a command prompt (cmd.exe) and use ipconfig to check the IP address of the 'OOBM NIC' (should be 10.251.X.91). Whenever you need to make a TFTP copy from a switch, you can use the mgmt interface and VRF of the switch and use this 10.251.x.91 IP address of the PC1 as the destination.

```
ipconfig
```

Push Base configuration to your devices

Configuration files have been prepared for your table and a basic SSH script can be used to push these configurations to your lab devices.

Before you push the commands, review the current interface status on Agg1.

11. On Agg1, review the current interface status

Q: What is the status of interface 1/1/1? What is the cause of this status?

A: The status is down because of 'Group speed mismatch'

Q: What does this mean?

A: On the 8325, port groups can be configured to operate at 10g or 25g speed. By default, they are set to operate at 25Gbps speed. When a different speed transceiver is connected, the 'Group speed mismatch' is reported. This will be corrected by using the 'system interface-group 1 speed 10g' command.

12. On PC1, in the ASTS folder on the desktop, run the '**ASTS-Lab01-task2.cmd**'.

NOTE: The script will make an SSH connection to the OOBM IP of the switches (using admin/aruba123 credentials) and enter the prepared commands from the ASTS\config folder on your desktop.

NOTE: The script simply deploys commands using SSH to the devices. Feel free to manually open SSH connections to the 4 devices and copy/paste the commands from the config files yourself if you prefer to push the configuration manually.

IMPORTANT: Since this is the first time Putty is making a connection, you will see messages about the public key. You can **ignore** these messages; Putty will automatically continue.

13. Close the script once it has completed.
14. On PC1 **Desktop > ASTS > configs** folder, review the contents of the 'ASTS_cfg_lab01-done-agg1.txt'. You can also use "show checkpoint asts-lab01-done" command on the switch. These are the commands that were pushed to the Agg1.
15. Repeat this review for Agg2, Sw1 and Sw2.

Task 3: Verify Device Access

Objectives

- Verify access between the lab devices

Device IP Address Table

Device	In-Band IP Address	OOBM IP Address
Agg1 (8325)	10.x.1.2/24	10.251.x.2
Agg2 (8325)	10.x.1.3/24	10.251.x.3
Sw1 (6300)	10.x.1.4/24	10.251.x.4
Sw2 (6300)	10.x.1.5/24	10.251.x.5
RouterA	10.255.101.11	
ClearPass	10.252.1.23	
NetEdit		10.251.x.200
PC1 (MGMT)		10.251.x.91

Steps

NOTE: The remote lab Webgate Guacamole interface is very convenient for accessing the console of the devices. However, for copy/paste, resize or screen buffer operations, an SSH connection is more flexible.

It is **highly** recommended to use PC1 and open SSH connections to the OOBM IP of the devices instead of the 'Console' connections in the remote lab interface.

1. On PC1, use Putty (on the desktop) or MTPutty (Multi-Tab Putty) to open an SSH connection to Agg1 OOBM IP address. Putty connections have been setup for you.

Verify in-band connectivity

2. On Agg1, ping to Agg2, Sw1 and Sw2 using the default routing table. Use the in-band IP address listed in the table above.
3. On Agg1 ping to the RouterA using the default routing table
4. On Agg1 ping to the ClearPass host (10.252.1.23) using VLAN1 as source IP address.

OOBM Connection tests

5. Using Agg1, ping the NetEdit host on the OOBM mgmt network.

NOTE: In case any ping tests would fail, verify that:

- The configuration was loaded correctly
- The ports are up
- The correct VRF was used for the ping command

IMPORTANT: Inform your instructor if some devices cannot be reached

Configuration file formats

6. On PC1, verify the TFTP server (3CDaemon on the desktop) is running. The TFTP root folder has been set to the folder Desktop > TFTP Root.
7. On Agg1, explore the copy command to upload the current configuration file with TFTP to the management PC using the OOBM network. Use filename 'asts-lab01-agg1.txt'.

Q: Which file format options are available when an AOS-CX file is copied?

A: cli and JSON based.

8. Upload the running configuration as CLI file to PC1, use filename asts-lab01-agg1.txt
9. Upload the running configuration as JSON file to PC1, use filename asts-lab01-json.txt
10. On the PC1, open the saved files (Desktop > TFTP Root) using Notepad++ and review the JSON file contents. This shows the hierarchical structure of the AOS-CX configuration file.

NOTE: When NetEdit is used to push configuration changes, NetEdit generates a new configuration based on JSON and uploads this to the target devices.

Table configuration backup and restore on Management PC

While the TFTP method works, the same JSON file can be downloaded using the RESTAPI.

NOTE: You will explore more about the REST API in Lab03.

A basic backup/restore script is available to backup and restore the running configuration of your 4 devices to local JSON files.

11. On PC1, Open Desktop > ASTS
12. Run '**ASTS-config-running-backup.cmd**', enter '**lab01-yourname**' (avoid spaces in the backup name) and press Enter.
13. Open Desktop > ASTS > backups and review the 4 backup files.

Example Restore

14. On Agg1, set the hostname to 'test'
15. To compare your current configuration with a saved checkpoint, use the checkpoint diff command.
16. On PC1, run the '**ASTS-config-running-restore.cmd**' and enter the name of your backup.
17. Verify on Agg1 that the hostname was reverted by the backup.
18. You can use this backup method at any time during the labs if you want to save your configurations.

Table checkpoint configuration (Optional)

19. You may also save 'local' configuration checkpoints on the devices at any point during the labs using the '**ASTS-config-checkpoint-save-from-running.cmd**' script.

NOTE: It is only possible to save a new checkpoint if the configuration is different from the other checkpoints. This is why an sample change is generated first.

20. On Agg1, make an sample configuration change (example: hostname Agg1-mylab01)
21. On PC1, run '**ASTS-config-checkpoint-save-from-running.cmd**'. Enter 'lab01-<your name>'.
22. On Agg1, review the checkpoint list.
23. Revert the checkpoints: On PC1, run '**ASTS-config-checkpoint-restore.cmd**' and enter 'asts-lab01-done'. This checkpoint was created earlier in this lab.
24. On Agg1, verify that the hostname has changed.
25. Review the last 10 events in the event log, you should see an event that the switch configuration asts-lab01-done was copied to the running-config.

Course labs 'known good' checkpoints

The training labs assume that you have completed the previous lab activities in the order of the class. In case you are unable to complete some of the labs, the actual commands for that activity can be found per switch on PC1 > Desktop > ASTS > configs.

You must also run a script to load all these configurations in the correct order and save them as checkpoints. This will allow you to compare your configuration to a known good configuration.

26. Deploy the 'known good' checkpoints on your devices. The script will also upload a JSON and CLI version of each lab to the PC1 TFTP folder. Make sure your TFTP server is running on PC1.
27. On PC1, open > Desktop > ASTS and run **ASTS-Lab01-config-deploy-all.cmd**

This will take about 6-7 minutes to load and save all the labs and the checkpoints. At the end, the asts-lab01-done will be loaded.

28. After the load completes, you can verify the saved checkpoints on each switch.
29. On PC1, open the Desktop > TFTPRoot folder. This will now contain a completed configuration of each lab activity. This can be used in case you want to perform a complete restore or compare your lab configuration with a known-good configuration.
30. Reboot all 4 switches to complete the process.

You have completed this Lab!

Aruba Training-Confidential

Optional Lab 02: Troubleshooting Tools

AOS-CX Troubleshooting and Analysis

NOTE: This lab activity simply demonstrates the available troubleshooting tools. These tools will be used again in the remainder of the labs. This is why the Lab activity is marked as optional.

In this lab activity several troubleshooting tools will be introduced. Some of these should be used in later lab activities to analyze or troubleshoot other technologies.

Requirements

This lab requires the completion of Lab 01.

Objectives

- Review Show commands
- Copy support logs
- Explore Debug options
- Diagnostics tools
- Traffic mirroring and packet captures

Task 1: Introduction to show commands and support files

Objectives

- Show commands
- Copy support log files
- Accounting logs
- Event logs
- Terminal monitor

Steps

Show commands and support files review

1. On Agg1, review the physical interface state.

Q: What is the speed of the link to Sw1 (1/1/1) and Agg2 (1/1/46)?

A: The link to Sw1 is 10Gbps, the link to Agg2 is operating at 25Gbps.

2. Review the available interface groups and the configured speed. Note that all the ports in a group are configured together.
3. Review the link-status

Q: How many link transitions did interface 1/1/1 have?

A: This depends on your lab environment.

4. Review the interface transceiver details. This is useful to detect the type, product number and serial number of the transceiver.

Interface statistics

5. Review the interface statistics.

Q: Does the output include all interfaces? Is this useful?

A: Interfaces that have 0 value statistics can clutter the output.

6. Review the interface statistics of the interfaces that have non-zero statistics
7. Clear the statistics and review the statistics again. It may take a few seconds before new statistics are available.
8. Open an extra SSH session to Agg1, review the statistics again.

Q: Have these statistics been reset as well?

A: No, the clear statistics command only clears the statistics for the current session.

9. Leave the second connection in place. On the original connection, clear the statistics at the device 'global' level.
10. Verify on the second connection that the statistics have now been reset at the device global level. You may close the second session now.

Q: What is the impact of clearing the statistics at the global level?

A: NAE and external SNMP monitoring systems will also notice the reset of the port statistics when they are cleared at the global level.

Support Logs

11. Copy support files for all features to PC1 (10.251.x.91 in VRF mgmt) with filename 'agg1-support-all.tar.gz' using TFTP.

NOTE: It may take a minute to complete the collection.

12. On PC1, unzip the file in the TFTP folder (7Zip is installed).
13. Untar the tar file (7Zip is installed).
14. Take a moment to review the files.
15. Open the ovssdbdump.html file. This is an HTML dump of the system configuration and state database. All settings and protocol states of the switch are stored in this database.
16. Search for the word '**lldp_neighbor**'. This will show the current LLDP neighbors. Close the HTML file.
17. Take a moment to review the folders in the 'Feature' subfolder.

Accounting logs

You will review the accounting logs in a moment. First make a sample configuration change, you should find this change back in the log.

18. On Agg1 define LAG 256 with LACP using interfaces 1/1/46 and 1/1/47. Enable the LAG256 and ports 1/1/46 and 1/1/47. Do not configure Agg2 at this point, this will cause an LACP blocked state that will be used in the next section.
19. Review the accounting log for the last 10 entries. The last entries should show the LAG commands.

Q: Which commands are stored in the accounting log?

A: Both configuration and show commands, as well as HTTP calls made using the REST API.

Event log

In this section you will review the status of the LAG and attempt to find the issue.

20. On Agg1, review the status of the LACP interfaces.

NOTE: This 'show lacp interfaces' command also applies to non-LACP enabled (static) aggregations!

21. Review the status of the LAG 256 interface

Q: What is the current status of the physical and the logical (LAG) interfaces?

A: The physical interfaces are lacp-block status, the logical interface is 'Disabled by LACP or LAG' status.

Q: Why is the LACP blocked status shown?

A: Let's review the event log!

22. Review the event log in reverse order. Note that each event is logged by a daemon (process) on the switch, and the daemon name is included after the hostname.

Q: What is the daemon name that logs 'LAG' messages?

A: This is logged by the lacpd (LACP daemon)

Q: Why are the ports in a blocked state?

A: The partner timed out; this means that there were no LACP messages received from the peer device. This is 'expected' since, you only configured the Agg1 side of the LAG, not Agg2.

23. Now filter the events based on the last 20 entries of the 'lacpd' daemon. Use 'show event - d ?' to see a list of the available processes.

Terminal Monitor

In this section, the terminal logging for events will be reviewed.

NOTE: Extended Terminal monitor functionality has been available since release 10.05, Terminal Monitor is supported in a SSH session only

By default, there is no live logging of events in the active session.

You will now explore some of the 'terminal monitor' options to get live logging of events.

24. Enable the terminal-monitor without any options

25. Enable and disable interface LAG256 again

Q: Was a message displayed in your session about this?

A: No.

26. Review the terminal-monitor severity options

Q: What is the default severity logging level?

A: The default level is error and above.

27. Adjust the terminal-monitor to include INFO messages and above

28. Disable and enable the interface LAG 256. The Link status events should now be shown.

NOTE: It may take a few seconds for the events to be displayed on your screen.

29. Reconfigure the terminal-monitor with a filter to see only messages that contain the string **'lacpd'** in them.

30. On Agg2, complete the LACP LAG 256 configuration on Agg2 and make sure the LAG is fully up with both interfaces 1/1/46 and 1/1/47.

31. On Agg1, you should now see in the terminal monitor events that the LAG is now online.

32. Disable the terminal monitor using the 'no terminal-monitor' command.

Task 2: Diagnostic tools

Objectives

- Enable diagnostics commands
- Explore the available diagnostic tools

Steps

Explore dropped frames - Jumbo frames

In this section Sw1 will be configured with jumbo frames, while the Agg1 will use the standard MTU. You will explore how you can detect this configuration error.

1. Use PC1 to open an SSH session to Sw1.
2. Configure interfaces 1/1/25 and 1/1/26 on Sw1 with an MTU of 9198.
3. Configure SVI 1 on Sw1 with an IP MTU of 9198.
4. On Sw1, ping to 10.x.1.2 (the IP of Agg1) to confirm a normal sized packet works fine.
5. On Sw1, ping to 10.x.1.2 with a datagram-size of 2000 and the 'do-not-fragment' option enabled. This ping should fail.
6. On Agg1, attempt to detect the drop reason, review the interface 1/1/1 statistics.

Q: Any dropped frames?

A: No, these are not counted in these statistics.

7. Review the available 'diag?' command options.
8. Enable diagnostics mode.
9. Review the available 'diag?' command options again.
10. Explore the 'diag interface 1/1/1 statistics' command.

Q: What categories have non-zero counters?

A: Several statistics may have non-zero counters. In this example, the 'frame size exceeded drops' will be increasing.

11. Correct the jumbo frames on Agg1 and Agg2, set all ports to MTU 9198.
12. Reset the Sw1 SVI1 MTU to 1500.

Explore dropped frames - VLAN mismatch

You will now configure a VLAN mismatch on the LAG256 between Agg1 and Agg2. VLAN 11 will exist on LAG256 on Agg2, but it will not be defined yet on Agg1.

1. Use PC1 to open an SSH session to Agg2.
2. Define VLAN11, SVI 11 on Agg2 and assign IP address 10.x.11.3/24
3. Configure LAG256 as VLAN trunk and allow all VLANs.
4. Start a ping on Agg2 to 10.x.11.2 with 1000 repetitions. This will force Agg2 to send ARP broadcast requests in VLAN 11, these will be transmitted over the LAG256 to Agg1.
5. Use PC1 to open an SSH session to Agg1.
6. Review the interface statistics for 1/1/46, 1/1/47 and LAG256.

Q: What do you notice about the statistics?

A: The member port statistics are automatically added to the LAG statistics.

Q: Are there any dropped frames?

A: No, these statistics only show hardware dropped frames.

7. Review the extended statistics for ports 1/1/46 and 1/1/47. Repeat the command 2 or 3 times to see which values are changing.

Q: Are there any dropped frames?

A: No, but the 'filtered' frames value seems to increase.

8. For 1/1/46 and 1/1/47, use the 'diag interface <interface> statistics' command a few times. Note the 'RX Statistic' category.

Q: Which value is increasing?

A: The 'Invalid VLAN drops' is increasing for this port.

9. Correct the configuration: On Agg1, define VLAN11, SVI 11 and assign IP address 10.x.11.2/24, configure LAG256 as VLAN trunk and allow all VLANs.

Task 3: Debugging options

Objectives

- Enable debugging for daemons
- Configure debug destinations
- Capture debug

Steps

Review default logging and logging destination

In this section you will explore the debug options, review the debug buffer and generate some debug events. Familiarize yourself with the process, so you can use this in the next chapters.

1. Review default debug status on Sw1.
2. Review default debug destinations on Sw1.

NOTE: The debug 'buffer' is by default enabled and does not show in the debug destination list.

3. Clear the debug buffer on Sw1.
4. Review the debug buffer on Sw1.
5. Clear the debug buffer and review it again to confirm it is empty.
6. Enable debugging for 'ndm unresolvedunicast' and 'ndm statemachine' on Sw1. NDM is the neighbor lookup daemon, this includes ARP.
7. Clear the ARP table on Sw1.
8. Review the contents of the ARP table on Sw1.

Q: Do you see any entries? Why?

A: Sw1 has 10.x.1.2 configured at the default gateway. This 'next hop' will be automatically resolved by the device.

Review normal ARP operation

In this section you will use the debug functions to review the normal operation of an ARP lookup. This will be done by performing a ping to 10.x.1.3. The ping will trigger an ARP request.

9. Enable terminal-monitor on Sw1 with severity debug and filter on 10.x.1.3 events.
10. On Sw1, ping 10.x.1.3, then wait a few moments.

NOTE: It can take a few moments for the terminal-monitor events to appear in your session.

11. In the debug output, you should find a message from the NDM_STATEMACHINE, 'Updated MAC from old 00:00:00:00:00:00 to a new MAC:'

Review missing ARP reply operation

Now review the operation when there is no response from the host. You will attempt to ping IP 10.x.1.30, an IP address that is not online in the lab network.

12. Update the terminal-monitor filter with severity debug on 10.x.1.30 events.
13. On Sw1, ping 10.x.1.30 a single time (repetition 1)
14. Wait a few moments and review the generated output. You should find a message from the NDM_STATEMACHINE about 'Max probe count reached' and 'Neighbor is INCOMPLETE'. This indicates that no response was received after the neighbor discovery messages were sent out.
15. Disable all debugging on Sw1 with the 'no debug all' command.
16. Disable the terminal-monitor on Sw1.
17. Review the last 40 lines of the debug buffer for the 'ndm' module on Sw1. This output will also show the NDM entries.

Task 4: Traffic mirroring and packet captures

Objectives

- Data plane - Local mirror
- Data plane - Remote mirror over GRE tunnel
- Data plane - Mirror to CPU
- Control plane - TCPDUMP analysis

Steps

Data plane mirror - Local mirror

In this section, a local mirror will be defined. Due to the remote lab environment and the use of VMs, the mirror packets cannot be observed on the MGMT PC. So, this mirror cannot be verified and only the configuration will be reviewed.

1. Open an SSH connection to Sw1
2. Review existing mirrors
3. Configure local mirror session 1
4. Add port 1/1/25 as the source interface and monitor both receive and transmit traffic
5. Add port 1/1/1 as the destination interface
6. Enable the mirror
7. Review the mirror
8. Review the mirror details

Q: How many output packets were handled by this mirror?

A: Depends on your setup

9. Disable the mirror
10. Remove the destination interface from this mirror

TIP: Use 'show run current' to see the configuration of the current context.

Dataplane mirror - Remote mirror over GRE tunnel

NOTE: The GRE destination IP must be reachable using the in-band network. It is not possible to configure the GRE tunnel over the OOBM network.

NOTE: The destination interface of the GRE tunnel should not be monitored by any mirror session.

11. On Sw1, verify IP connectivity to the destination (PC1 on port 1/1/1) with a ping to 10.x.1.91.
12. Configure the mirror 1 with a destination tunnel to PC1, use the switch VLAN1 source IP address. There is no need to use the DSCP and VRF options.

NOTE: The command supports a 'vrf' option, this applies to all VRFs, except the 'mgmt' VRF.

13. Enable the mirror.
14. Start Wireshark on PC1.
15. On the Capture menu, select Options.
16. Select the Lab NIC connected to 6300 port 1 and select Start.
17. On Sw1, ping to Agg1 VLAN 1 IP address (10.x.1.2)
18. On the PC1, stop the Wireshark trace after about 30 seconds.

Q: Are there any ICMP Destination Unreachable packets? Why?

A: Yes, since the PC does not have a real tunnel endpoint, the PC complains that the GRE packet cannot be handled. This is expected, so these ICMP messages should be ignored / filtered in the output

19. Apply a display filter in Wireshark. Ensure that any packet with the source IP of the PC is not shown, and that only GRE packets are shown.

`!ip.src == 10.x.1.91 and gre`

NOTE: Make sure to press ENTER after entering the display filter

!ip.src == 10.12.1.91 and gre						
No.	Time	Source	Destination	Protocol	Length	Info
46	23.998990	88:3a:30:a4:1f:4f	Spanning-tree-(for-...	STP	161	MST. Root = 32768/0/88:3a:30:a4:1f:40 Cost = 0 Port = 0x801a
47	24.550433	10.12.1.4	10.12.1.2	ICMP	184	Echo (ping) request id=0x2e70, seq=1/256, ttl=64 (reply in 48)
48	24.550435	10.12.1.2	10.12.1.4	ICMP	184	Echo (ping) reply id=0x2e70, seq=1/256, ttl=64 (request in 47)
50	25.555544	10.12.1.4	10.12.1.2	ICMP	184	Echo (ping) request id=0x2e70, seq=2/512, ttl=64 (reply in 51)
51	25.555545	10.12.1.2	10.12.1.4	ICMP	184	Echo (ping) reply id=0x2e70, seq=2/512, ttl=64 (request in 50)
53	25.900666	88:3a:30:a4:1f:4f	Spanning-tree-(for-...	STP	161	MST. Root = 32768/0/88:3a:30:a4:1f:40 Cost = 0 Port = 0x801a
> Frame 47: 184 bytes on wire (1472 bits), 184 bytes captured (1472 bits) on interface 0 > Ethernet II, Src: 88:3a:30:a5:52:80 (88:3a:30:a5:52:80), Dst: Vmware_b1:14:aa (00:50:56:b1:14:aa) > Internet Protocol Version 4, Src: 10.12.1.4, Dst: 10.12.1.91 > Generic Routing Encapsulation (Transparent Ethernet bridging) > Ethernet II, Src: 88:3a:30:a5:52:80 (88:3a:30:a5:52:80), Dst: b8:d4:e7:41:38:00 (b8:d4:e7:41:38:00) > Internet Protocol Version 4, Src: 10.12.1.4, Dst: 10.12.1.2 > Internet Control Message Protocol						

20. Take a moment to parse any of the packets.

Q: What is the outer source IP?

Q: What is the GRE protocol type?

21. Attempt to find an LLDP frame from the Agg1 switch.

TIP: You can just add '**and lldp**' to the display filter.

22. Attempt to find the ICMP ping between Sw1 and Agg1

TIP: If you applied the previous filter, just replace 'lldp' with 'icmp' and press ENTER.

NOTE: The egress interface of the GRE tunnel cannot be shared for multiple mirror sessions.

ip.src == 10.12.1.91 and gre and lldp						
No.	Time	Source	Destination	Protocol	Length	Info
8	3.492947	b8:d4:e7:41:38:fd	LLDP_Multicast	LLDP	153	TTL = 120 System Name = T12-Agg1 System Description = Aruba J1635A GL.10.05.0001
21	8.493918	b8:d4:e7:41:38:fd	LLDP_Multicast	LLDP	153	TTL = 120 System Name = T12-Agg1 System Description = Aruba J1635A GL.10.05.0001
30	12.760239	88:3a:30:a5:52:90	LLDP_Multicast	LLDP	165	TTL = 120 System Name = T12-Sw1 System Description = Aruba J1668A FL.10.05.0001

> Frame 8: 153 bytes on wire (1224 bits), 153 bytes captured (1224 bits) on interface 0
 > Ethernet II, Src: 88:3a:30:a5:52:80 (88:3a:30:a5:52:80), Dst: Vmware_b1:14:aa (00:50:56:b1:14:aa)
 > Internet Protocol Version 4, Src: 10.12.1.4, Dst: 10.12.1.91
 > Generic Routing Encapsulation (Transparent Ethernet bridging)
 > Ethernet II, Src: b8:d4:e7:41:38:fd (b8:d4:e7:41:38:fd), Dst: LLDP_Multicast (01:80:c2:00:00:0e)
 > Link Layer Discovery Protocol

23. On the Sw1, disable the mirror 1 session

24. Remove the tunnel destination

Dataplane mirror - TSHARK - Mirror to cpu - Live

In this section, the Tshark utility will be used. Tshark can be used on the AOS-CX switches to capture and view **dataplane** packets that are mirrored to the CPU.

This may include control plane traffic if that happens to pass through the physical interface that is being mirrored.

The mirror to the CPU will continuously send packets to the CPU for as long as the mirror session is enabled. Make sure to disable the mirror session when the analysis has completed.

The packets that arrive at the CPU can be shown live on the cli or they can be stored in a PCAP file and transferred at a later time.

25. Configure mirror session 1 with the destination CPU

26. Enable the mirror session
27. Launch the diagnostic utility “tshark” without any options on Sw1. This enables the live view.
28. Open an SSH session to Agg1.
29. From the Agg1, ping the PC1 (10.x.1.91). This traffic passes through the Sw1 port 1/1/25 that is being mirrored.
30. In the Sw1 session, packets should be shown in the session
31. Use CTRL-C to stop the live view
32. Verify that the test ping packets were received.

TSHARK - Mirror to cpu - PCAP

The AOS-CX system can also save these data plane packets that are mirrored to the CPU in a PCAP file, so the PCAP file can be transferred to an external system for analysis.

33. Clear any existing PCAP file on Sw1. If this is the first time, the system will report that the capture file does not exist.
34. Start the tshark utility with the file option on Sw1 and leave it running.
35. On Agg1, ping to PC1. This traffic should go through Sw1 via port 1/1/25.
36. On Sw1, stop the tshark utility using CTRL-C
37. Copy the tshark file using the copy command from Sw1 to PC1. Either the in-band IP (10.x.1.91) or the OOBM IP and VRF can be used for the copy operation to PC1.
38. On PC1, open the PCAP file and verify the Agg1 ping is included.

Apply a display filter ... <Ctrl-/>					
No.	Time	Source	Destination	Protocol	Length Info
1	0.000000000	88:3a:30:a4:1f:4f	Spanning-tree-(for-	STP	119 MST. Root = 32768/0/88:3a:30:a4:1f:40 Cost = 0 Port = 0x801a
2	1.005241180	b8:d4:e7:41:38:fd	LLDP_Multicast	LLDP	111 TTL = 120 System Name = T12-Agg1 System Description = Aruba JL635A GL.10.05.0001
3	2.000140290	88:3a:30:a4:1f:4f	Spanning-tree-(for-	STP	119 MST. Root = 32768/0/88:3a:30:a4:1f:40 Cost = 0 Port = 0x801a
4	3.935310393	10.12.1.2	10.12.1.91	ICMP	142 Echo (ping) request id=0x01ce, seq=1/256, ttl=64 (reply in 5)
5	3.935837594	10.12.1.91	10.12.1.2	ICMP	142 Echo (ping) reply id=0x01ce, seq=1/256, ttl=128 (request in 4)
6	3.999543978	88:3a:30:a4:1f:4f	Spanning-tree-(for-	STP	119 MST. Root = 32768/0/88:3a:30:a4:1f:40 Cost = 0 Port = 0x801a
7	4.969154935	10.12.1.2	10.12.1.91	ICMP	142 Echo (ping) request id=0x01ce, seq=2/512, ttl=64 (reply in 8)
8	4.969832497	10.12.1.91	10.12.1.2	ICMP	142 Echo (ping) reply id=0x01ce, seq=2/512, ttl=128 (request in 7)

> Frame 4: 142 bytes on wire (1136 bits), 142 bytes captured (1136 bits) on interface 0
 > Ethernet II, Src: b8:d4:e7:41:38:00 (b8:d4:e7:41:38:00), Dst: Vmware_b1:14:aa (00:50:56:b1:14:aa)
 > Internet Protocol Version 4, Src: 10.12.1.2, Dst: 10.12.1.91
 > Internet Control Message Protocol

39. When the test is successful, on Sw1, disable the mirror 1 session
40. Remove the mirror source from the mirror session 1

NOTE: If there is any issue with the PCAP file, it can be removed manually with the command `diag utilities tshark delete-file`

Class-based mirroring

The Classifier Policy Engine on AOS-CX includes support to mirror traffic as a policy action. In this section a classifier will be defined to match on ICMP traffic that passing through the dataplane of the switch. Traffic matching these classes will be mirrored to mirror session 1, that is the CPU in this example.

41. On Sw1, define a new class named 'icmp' that defines ICMP to any destination
42. Define a new policy named 'mirror'. Add the 'icmp' class and assign action to mirror to the previously defined mirror session.
43. Enable the mirror session
44. Apply the policy 'mirror' to the interface 1/1/25 to both in and out directions.
45. Open a second SSH session to switch Sw1.
46. In the second SSH session, start a live Tshark session.
47. On Agg1, ping 1 time to PC1 (ping 10.x.1.91 repetitions 1).
48. This ICMP echo request and reply should be shown in the live Tshark session.
49. On PC1, open a ssh session to Agg1 (10.x.1.2). You do not need to login. This traffic should not be shown in the live Tshark session since it does not match the classifier.
50. On PC1, close the new SSH session.
51. On Sw1, disable the mirror session.
52. On Sw1, remove the policy 'mirror'

TCPDUMP

TCPDUMP operates on the software of the AOS-CX device. Any control or management plane protocol can be monitored and analyzed using TCPDUMP, including packets that would use the OOBM port.

Example control plane packet trace on switch.

NOTE: Additional output options can be enabled, such as:

- verbosity level4
 - print link_level_header
 - print data
-

53. Configure a test RADIUS server on Sw1 and define a RADIUS tracking account.
54. Open a second SSH session to switch Sw1
55. Run tcpdump in one session. Port 1812 is used for RADIUS communication.

56. In the other SSH session, enable RADIUS tracking.
57. Use the disable/enable feature of the tracking command to trigger a new RADIUS request if you want to repeat this test
58. Review TCPDUMP output

NOTE: TCPDUMP will stop automatically after 20 packets by default.

59. Disable the RADIUS tracking.
60. Optionally you can save a new checkpoint on Agg1, Agg2, Sw1 and Sw2.

You have completed this Lab!

Lab 03: Monitoring and Automation Tools

AOS-CX Monitoring and Automation

In this lab activity several monitoring and automation tools will be introduced. Some of these may be used in later lab activities to monitor or automate other technologies.

Requirements

This lab requires completion of Lab 02.

If you did not complete the optional Lab02, load the checkpoint '**asts-lab02-done**' on Agg1, Agg2, Sw1 and Sw2.

Objectives

- NetEdit monitoring and logging
- Network Analytics Engine
- Using the AOS-CX REST-API
- SNMP

Scenario

NetEdit

The customer wants to discover all devices in NetEdit so they can be monitored and managed. They heard that it is also possible to monitor SNMP devices, such as the Aruba Gateway. They heard that NetEdit can automatically discover other devices once a seed device has been added, but that did not seem to work for them when they tried it in their lab. You must troubleshoot this issue and ensure that additional switches are automatically discovered in NetEdit.

RESTAPI

The customer is interested in the RESTAPI of the switches. They want you to demonstrate how to see the available RESTAPI functions and demonstrate how the LLDP neighbors of all the switch ports can be collected. They also want to understand how the output of a RESTAPI call can be restricted to just the statistics instead of all of the details, such as the configuration. They also noticed that after a few RESTAPI logins, the system does not accept new logins. They want you to explore this issue and explain it.

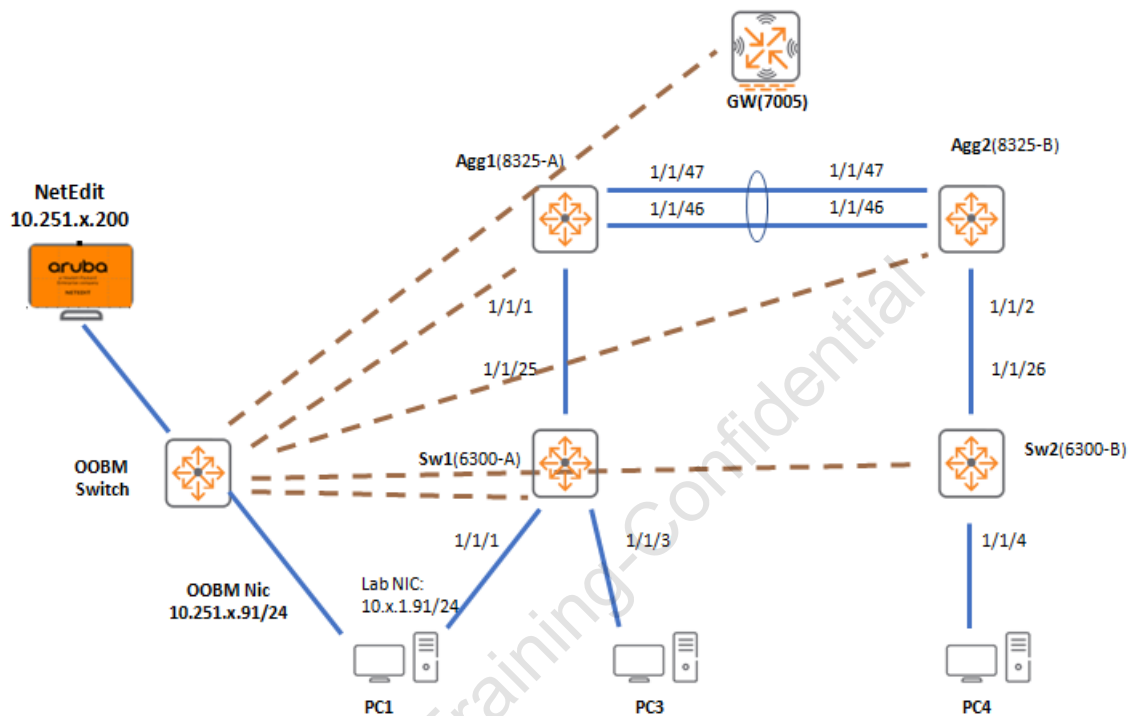
NAE

The customer has heard about the configuration features of the Network Analytics Engine. They would like to see a demonstration of the automatic configuration of the interface description based on the LLDP peer of that port. They expect to get some information about the logical hierarchy of an NAE script. They would also like to understand how to troubleshoot any errors or changes made by an NAE script, so they expect to get some debug output from the running NAE LLDP script.

SNMP

The customer wants to use an existing SNMP monitoring solution. SNMP should be operational in the default VRF. They also want to ensure that the default SNMP community 'public' cannot be used anymore. The existing SNMP solution will connect using community string 'aruba123'. The customer also wants to configure Airwave SNMP monitoring using SNMPv3 (credentials amp/aruba123). Note: SNMP configuration on the switches is sufficient, you do not need to discover the devices in Airwave. In the past, they had issues with some SNMP monitoring solutions and found that it is hard to discover which systems were sending bad SNMP community strings. They want you to demonstrate debugging of an SNMP request with an invalid community string.

Lab Diagram



Task 1: NetEdit

Objectives

- Discover Devices in NetEdit
- Review the NetEdit log files
- Discovery subnets
- Using NetEdit topologies

Steps

1. Using PC1, verify RESTAPI is enabled and configured with read-write access on the Agg1 and Agg2 switches

NOTE: With release 10.05, this has become the default for all AOS-CX platforms. In previous releases the REST API had to be enabled for read-write manually on the 8300 and 8400 platforms.

2. Using PC1, open a session to NetEdit (10.251.x.200), login with admin/no password

NOTE: When prompted to change your password, change it to 'aruba123'

NOTE: If the login fails, login with admin/aruba123

3. Define a new credential set 'asts'
 - REST admin/aruba123
 - SSH admin/aruba123
 - SNMP SNMPv2 public
4. Define a new managed subnet for 10.0.0.0/8, assign the 'asts' credential set
5. Discover the Agg1 switch in your table (Use the OOBM IP 10.251.x.2) and the Aruba Gateway (10.251.x.6) by entering them into the seed list.

TIP: When the managed subnet has been defined, the discovery window does not need to be completed with the discovery subnets anymore, just enter the seed addresses.

6. Verify in the device list that the Agg1 and Aruba Gateway have been discovered.

Discovery issues

Q: Are the other switches added automatically?

A: The auto-discover feature is enabled, but the other switches are not automatically discovered.

7. Investigate why they are not discovered. Check the 'Logs' of NetEdit.

Q: What is the error message you see for the device discovery?

A: NetEdit shows that the device 10.x.1.y is unreachable.

Q: How did NetEdit find this 10.x.1.y address?

A: NetEdit uses LLDP peer information. Since the peer switches report their in-band management IP address, that will be the IP address that NetEdit attempt to reach. Since NetEdit is on the OOBM network, this will not work.

8. On Agg1, review the LLDP neighbor details for port 1/1/46. Note the current Neighbor Management-address.
9. On Agg2, adjust this by manually assigning the OOBM IP address as the IPv4 management IP in LLDP.
10. On Agg1, apply the LLDP OOBM management IP configuration.
11. On Agg1, confirm the LLDP neighbor details have been updated with the correct management IP address.
12. In NetEdit, confirm that the Agg2 has now been automatically discovered. It can take some time to complete the discovery.
13. Correct the LLDP Management IP on Sw1 and Sw2.
14. Verify all 4 switches are now listed in the device list. This may take a few minutes.

Diagnostics page

15. Attempt to discover IP 10.251.x.254 by adding it as a seed device. This is the default gateway on the OOBM network and this device cannot be accessed.
16. Check the NetEdit Diagnostics page to see the status.
17. Once you have seen the device with the discovery error on the Diagnostics page, you may delete all the unreachable devices.

NetEdit Insights

18. In NetEdit, navigate to Devices using the left menu. Select Agg1 > Edit Configuration.
19. Select interface 1/1/1 in the configuration. Notice the Insights window on the right-hand side that shows the neighbor and the neighbor port configuration.

NetEdit Debugging

20. Under Logs > Actions > select 'Enable Debug Logging'
21. On Agg1, enable terminal-monitor with severity debug by CLI if using ssh to connect the device.
22. Use NetEdit to define a plan and change the Agg1 interface 1/1/1 description from 'sw1' to 'sw1test'. Do not validate or deploy the plan yet.
23. On Agg1, enable debugging for the 'rest' module
24. In NetEdit, use the Validate configuration

Q: Where any REST call made to the switch?

A: Yes, NetEdit will perform a dry-run of the configuration on the switch itself. This ensures that only a switch-validated configuration will be sent to the switch when the deployment is triggered.

The hpe-restd debug logs should reveal:

- a new session is started by user admin
- the privilege level that is assigned to this login, this controls the allowed actions of the account
- POST, GET and DELETE actions are listed, this is NetEdit that submits the 'dryrun' configuration that will be tested by the switch. After the validation has completed, the 'dryrun' configuration is deleted.

25. Use NetEdit to deploy the configuration. The hpe-restd debug logs should reveal that a new switch configuration has been written to the running configuration.

26. On the Agg1, review the debug status using 'show debug'

Q: What do you notice?

A: Debugging has been disabled

Q: What does this mean?

A: Debugging settings are stored as part of the configuration file. When the new NetEdit configuration was pushed, it did not include the debug configuration, since the debugging was enabled after the Plan was created.

IMPORTANT: Make sure to disable any debugging after you have completed your troubleshooting, since the debugging would be saved in the configuration, so it would also survive a device reboot. This also applies to saving checkpoints.

Collect NetEdit logs

27. In NetEdit, navigate to Logs > under the Action menu, download the support logs with the 'Export Support Bundle' action.
28. On the PC1, use 7zip to extract the support logs.
29. Take a few moments to check the access.log and the netedit.log files.

This concludes the NetEdit task. Close the browser session.

Aruba Training-Confidential

Task 2: RESTAPI

Objectives

- Using the built-in Swagger interface
- Example using CURL to access the REST-API
- REST Sessions and clearing REST sessions on AOS-CX

Steps

Using the built-in Swagger interface

1. On PC1, open an HTTPS connection to the Sw1 (10.251.x.4)
2. Login with the admin credentials
3. At the right-top, use the 'gear' icon to access the 'System' menu
4. Select 'API', a new window will open with the Swagger interface. It will take a few moments before the interface completes loading.

Wildcard

5. In the Swagger interface, select LLDP_Neighbor
6. Select the 'GET /system/interfaces/{pid}/lldp_neighbors' link
7. Enter '1/1/25' for the **pid**, that is the instance of the interface.
8. Scroll down and click on 'Submit'.

Q: How many LLDP Neighbors are shown in the response? Is this correct?

A: Just 1, yes, there is only 1 LLDP neighbor on this port

9. For many 'instances', the REST API supports using a wildcard '*'. Change the **pid** to * and submit the request again.

NOTE: Be careful with this wildcard character. Requesting all statistics for all interfaces may be resource intensive for the switch. Attempt to limit the use of this wildcard when possible.

10. Verify that LLDP neighbors for all dataplane interfaces are shown now

Escape Characters

11. In the Swagger interface, select 'Interface' > Get '/system/interfaces/.
12. Scroll down and submit the request.

Q: What is the request URL? What is the current browser URL?

A: The actual RESTAPI is different from the current browser URL, since the current browser URL shows the Swagger application, not the actual RESTAPI.

13. Review the response body.

Q: What do you notice about the interface name?

A: The interface name 1/1/1 contains '/', the forward slash is encoded since it would conflict with the / in the URL. This is seen as the %2F character in the URL.

Filtering output

14. Under 'Interfaces' , select GET /system/interfaces/{id}.
15. Enter '1/1/1' for the id > Submit

Q: What type of information can you see in the output?

A: Statistics, status and configuration are shown in the response.

16. Scroll back to the request, enter 'statistics' for the selector and submit again
17. The response should now only show the all the statistics for this interface, but the configuration and status is now filtered.
18. Scroll back to the request, in the attributes list, select 'lldp_statistics' and submit again.
19. Notice how the response now only includes the statistics for this interface. This demonstrates the flexibility of filtering using the RESP API.

Configuration Depth

20. Navigate to 'MAC' > GET /system/vlans/{pid}/macs
21. Enter * for the pid (represents the VLAN ID here) and Submit

Q: What does the %3A in the MAC address mean?

A: Just like the / character, the : character is escaped in the REST URLs

22. You can see the details of each MAC by making another query for each response URL. However, the REST API can make this easier for you by automatically querying URLs in the

response. If that response would contain a URL, it could do it again, etc. You can control the depth of this detail using the 'depth' option. In many cases, a depth of 1 would be sufficient to get 1 additional level of detail.

23. Set the depth option to 1 and Submit again.

Q: Do you see more information now?

A: Yes, for each MAC Address entry, the port information is included. Each entry also contains a 'mac_addr' field that shows the classic MAC address format.

REST Sessions and clearing REST sessions on AOS-CX

24. Open an SSH session to Sw1.

25. Open the online help. This shows an example to access the REST API using the Command-line URL (CURL) utility.

https://techhub.hpe.com/eginfolib/Aruba/OS-CX_10.04/5200-6725/index.html#GUID-68022E2E-62AD-4905-95A8-B88334E3659A.html

26. Use PC1 command prompt to send a CURL command, save the session cookie in a file 'sw1.txt'.

NOTE: The ASTS lab folder on the desktop contains ready to use batch files for the commands below.

```
curl -X POST --insecure --header "Content-Type: application/x-www-form-urlencoded" --header "Accept: application/json" -d "username=admin&password=aruba123" "https://10.251.x.4/rest/v1/login" -c sw1.txt
```

27. Repeat this command a few times until this message is shown

Login failed: session limit reached.

28. Open a private browser window and attempt to open a new HTTPS session to the switch.

Q: What is the error?

A: Login failed: User session limit reached

29. Use the SSH connection to clear all HTTPS sessions

This demonstrates how the active sessions can be cleared.

IMPORTANT: Any REST or Web application should be configured with a proper 'logout' to clean up the session on the switch.

30. Repeat the CURL login (just once now) and save the cookie in the sw1.txt file.

31. Query the VLAN list using CURL, by using the saved cookie file.

```
curl -X GET --header "Accept: text/plain" "https://10.251.x.4/rest/v1/system/vlans" -b  
sw1.txt -k
```

32. Attempt to query using depth=2

```
curl -X GET --header "Accept: text/plain" "https://10.251.x.4/rest/v1/system/vlans?depth=2" -  
b sw1.txt -k
```

Aruba Training-Confidential

Task 3: NAE

Objectives

- NAE Script Hierarchy
- NAE Script Action
- NAE status integration with NetEdit

Steps

Introduction

In this task, you will activate an NAE script. The example script will automatically set the correct interface description based on the received LLDP information on a port. This script will monitor the link state changes and make configuration changes (description) when needed.

NAE Script Hierarchy

An NAE script contains several main sections: the Manifest, Parameter definition, Monitor, Rules, Conditions and Actions.

1. On PC1, from the ASTS lab files folder, open the ASTS - Lab - 03 - LLDP_Renaming_2.0.py NAE script using Notepad++.
2. Review the Manifest

Manifest is an object and it is defined as a dictionary. A dictionary is like an array, where you can have a key, and a value assigned to that key, so basically a key-value pair. In the Manifest dictionary you see four key-value pairs: Name, Description, Version and Author. You can assign your own values to the keys and these values will be used in the NAE to display description, author and version information.

Parameter definition

3. Review the Parameter Definitions in the script.

Parameter Definitions is the main dictionary, and within that dictionary there is another dictionary (interface_id) with multiple objects. The Name, Description, Type are mandatory objects whereas the Default key is not required.

Monitor, Rules, conditions and actions

4. Review the Monitors in the script.

The monitor, rules, conditions and actions are typically combined in a Python class. A class is a collection of Python definitions (similar to a function) that can be assigned to an object.

Q: What is the URL1 that will be monitored by this NAE script?

A: The URI1 is '/rest/v1/system/interfaces/*?attributes=link_state'.

Q: What does this mean?

A: The * means that 'all' interfaces will be monitored. The 'Attributes' is used to ensure that only the 'link_state' is queried instead of all configuration and statistics. This is an example of a selective query.

NAE Script Debugging

5. Review the script for debug commands (self.logger.debug). The NAE script developer can include debug statements in the script. These debug statements can be seen in the switch debug logs when NAE debugging is enabled.
6. Open an SSH connection to Sw1
7. Activate nae script debugging.
8. Use the web interface to upload the script to Sw1. Notice that the Manifest data is shown when uploading the script.
9. Define a new NAE agent named 'lldp' based on the LLDP_Renaming script. Notice the Parameters and their default values are shown as you have seen in the script.
10. Stop and start the script.
11. Review the debug log using 'show debug buffer module nae'. The NAE debug logging is logged by the 'hpe-policyd' daemon.
12. Use 'show interface brief' to verify that the port descriptions have been updated based on the LLDP peer name.
13. Disable all debugging.

NAE Script deployment

The NAE scripts and agents are also available via the REST API. An example CURL script is provided to deploy the same LLDP Script and Agent to Agg1, Agg2 and Sw2.

14. On PC1, open the ASTS folder on the desktop. Run the '**ASTS-Lab03-nae-lldp-add.cmd**' script.
15. Open an SSH session to Agg1, use 'show interface brief' to verify that the interface descriptions have been updated.

tip: The NAE scripts are saved in the configuration, so the 'show running-config' command will now include the Base64 encoded version of the NAE script. To see the classic running configuration, you can use the 'show run | exclude nae' command.

Clearing NAE data

An erase zeroize action will clear all configuration and it will also clear the NAE time series database. It is possible to clear the time series database without performing the zeroize action using the command 'clear nae-data'. Keep in mind that any historical information of NAE will be removed with this action.

16. Clear the nae data on Sw1.

NOTE: This command is hidden and must be entered without command completion.

17. Use the web interface on Sw1 > Analytics. There should be no more historical data on the dashboard now.

Aruba Training-Confidential

Task 4: SNMP Configuration

Objectives

- Configure SNMP Support for Airwave
- SNMP VRF Support
- SNMPv3 and SNMPv2
- Troubleshooting SNMP
- Diagnostics SNMP

Steps

SNMP VRF Support

1. On PC1, open a connection to Agg1
2. Verify the default enabled VRFs for SNMP.
3. Enable snmp-server on Agg1 for mgmt VRF and the default VRF.
4. Verify the active VRFs for SNMP again, both default and mgmt should be active.

NOTE: In 10.4, the SNMP process is bound to a single VRF. Configuring an additional VRF will disable SNMP on the previously configured VRF. In 10.5, the SNMP server can be bound to multiple VRFs.

SNMPv2 communities

5. Review the SNMP communities
6. Attempt to remove the public community

NOTE: It is not possible to just delete the public community. It will be automatically removed when a new community is defined.

7. Configure a new community 'aruba123'
8. Review the SNMP communities

Q: What does this mean?

A: AOS-CX will automatically remove the 'public' community when a user-defined community is added to the system.

Configure SNMPv3 Support for Airwave

9. Define an SNMPv3 user account on Agg1: amp. Use SHA with key 'aruba123' and AES with key 'aruba123' for the authentication and privacy.
10. Define an SNMPv3 context 'aruba' and link it to the default VRF and community 'aruba123'
11. Link the SNMPv3 account to the SNMPv3 context.

Troubleshooting SNMP packets

12. Use tcpdump to review the SNMP packets and community strings. Make sure to run TCPDUMP in the correct VRF context.
13. On PC1, run '**ASTS-Lab03-SNMP-GET-badcommunity.cmd**'. This will send an SNMP v2 request using 'badcommunity' as the community name to Agg1.
14. Stop the TCPDUMP.

SNMP Events

15. Review the last 200 events in the event log, filter on the 'snmpd_wrapper' daemon.

Q: In which the namespaces is the SNMP Agent running?

A: The SNMP Agent is running in the 'swns' and the 'VRF_1' namespace.

Q: What would these namespaces indicate?

A: Namespaces are a technology to isolate processes and contexts. In AOS-CX, the VRF 'default' is represented as the 'swns' namespace, the VRF 'mgmt' is represented by the 'VRF_1'.

Diagnostics SNMP

16. It is also possible to perform local SNMP queries on the switch.
17. Send a local SNMP GET request to query the device uptime.
18. Optional: Save a new checkpoint on Agg1, Agg2, Sw1 and Sw2

You have completed this Lab!

Lab 04: VSX and Layer2 Technologies

In this lab activity VSX is configured and several Layer2 technologies are reviewed

Requirements

This lab requires completion of Lab 03.

Objectives

- Configure VSX
- Apply VSX best practices
- Configuring STP with VSX
- Configuring Loop protection

Aruba Training-Confidential

Scenario

VSX

1. The customer wants you to configure Agg1 and Agg2 into a VSX pair with Agg1 as the primary.
2. They want to use LAG256 (ports 1/1/46 and 1/1/47) for the ISL.
3. The keepalive should be port 1/1/45, assigned to the dedicated VRF named 'KA'.
4. The keepalive IP subnet is 192.168.0.0/31.
5. They have heard that AOS-CX has 2 types of native VLANs (tagged and untagged), they have noticed this on the VSX ISL and would like to understand what this means.
6. For the VSX configuration, they expect to have configuration synchronization for these features:
 - dhcp-relay
 - lldp mclag-interfaces
 - qos-global
 - stp-global
 - vsx-global
7. They expect an active gateway configuration for SVI1 and SVI12 with
 - virtual MAC 02:00:00:00:00:01
8. They do not want to manually copy the configuration of the SVI Active Gateway to the secondary node (Agg2).

VSX LAG - Aruba Gateway

The customer expects an LACP VSX LAG called 'lag5' to connect to the Aruba Gateway using port 1/1/5 on both Aggregation switches. It should allow VLANs 1,11,12.

VSX LAG - Sw2

1. The customer expects an LACP VSX LAG called 'lag2' to connect to Sw2 configured using ports 1/1/2 on both Aggregation switches. It should allow VLANs 1,11,12.
2. They want to understand the behavior of LACP on the AOS-CX platform. They want you to demonstrate what happens when Sw2 is connected with a static LAG to the VSX LAG. After troubleshooting, Sw2 should connect using an LACP LAG to the VSX pair.
3. They also want to see how they can see the LACP statistics.

VSX - DHCP Relay

1. The DHCP server that supports VLAN12 is 10.252.1.21.
2. On Sw2 port 1/1/4 they want to have a PC connected to VLAN12.
3. The customer wants to understand how the DHCP relay function and fail-over works in a VSX pair.
4. They would like to see the relay debug output from both Aggregation switches.
5. They also want to understand what happens with DHCP relay when e.g. SVI12 is shutdown on Agg1.
6. The customer expects you to explain this so you should understand how DHCP Relay works.
7. They also want to understand if there is any impact on the user traffic when an SVI is shutdown on a VSX member. This can be tested on PC4 with a ping to 10.251.1.23 and 10.251.1.24 when SVI12 is shutdown on Agg1.

VSX - Virtual Active Gateway

The customer has heard that it is possible to use a single IP for both Aggregation switches as the interface and Active Gateway IP.

They want to see this demonstrated on SVI11, that includes the DHCP relay configuration.

Client in VLAN11

1. The customer has a PC3 connected to port 1/1/3 on Sw1. They want this PC in VLAN11 connect via the link Sw1-Sw2 to the VSX pair. The port 1/1/27 should be used for the Sw1-Sw2 link and it should become a VLAN trunk with VLAN 1 and VLAN 11, VLAN 1 is the native VLAN.
2. They notice that the PC3 does not get an IP address in VLAN 11 and they want you to troubleshoot this.
3. The customer informed you that the uplink ports 1/1/25 and 1/1/26 on Sw1 will be used for L3 IP only, not for a Layer2 LAG, so they can be shut down after the troubleshooting.

VSX Best Practices

1. In order to make it stable, the customer wants to have a static MAC address configured on the VSX system: 02:01:00:00:00:00.
2. On the VSX, they want to ensure that all Layer2 and Layer3 SVI interfaces support jumbo frames.
3. On the Sw2 access switch, they want all Layer2 interfaces to support jumbo frames.

VSX Linkup Delay

1. The customer wants to understand the impact of an ISL failure and recovery. They have noticed that Agg2 interfaces remain disabled for some time when it reboots or when the ISL is restored. You must investigate this behavior and understand which Layer2 or Layer3 interfaces are impacted by this feature.
2. They want to ensure that the LAG to the Aruba Gateway is not affected by this feature.
3. They also want to understand what must be done, so that an SVI is not impacted by this feature, since they may use SVI as an upstream routed connection.

Spanning-Tree

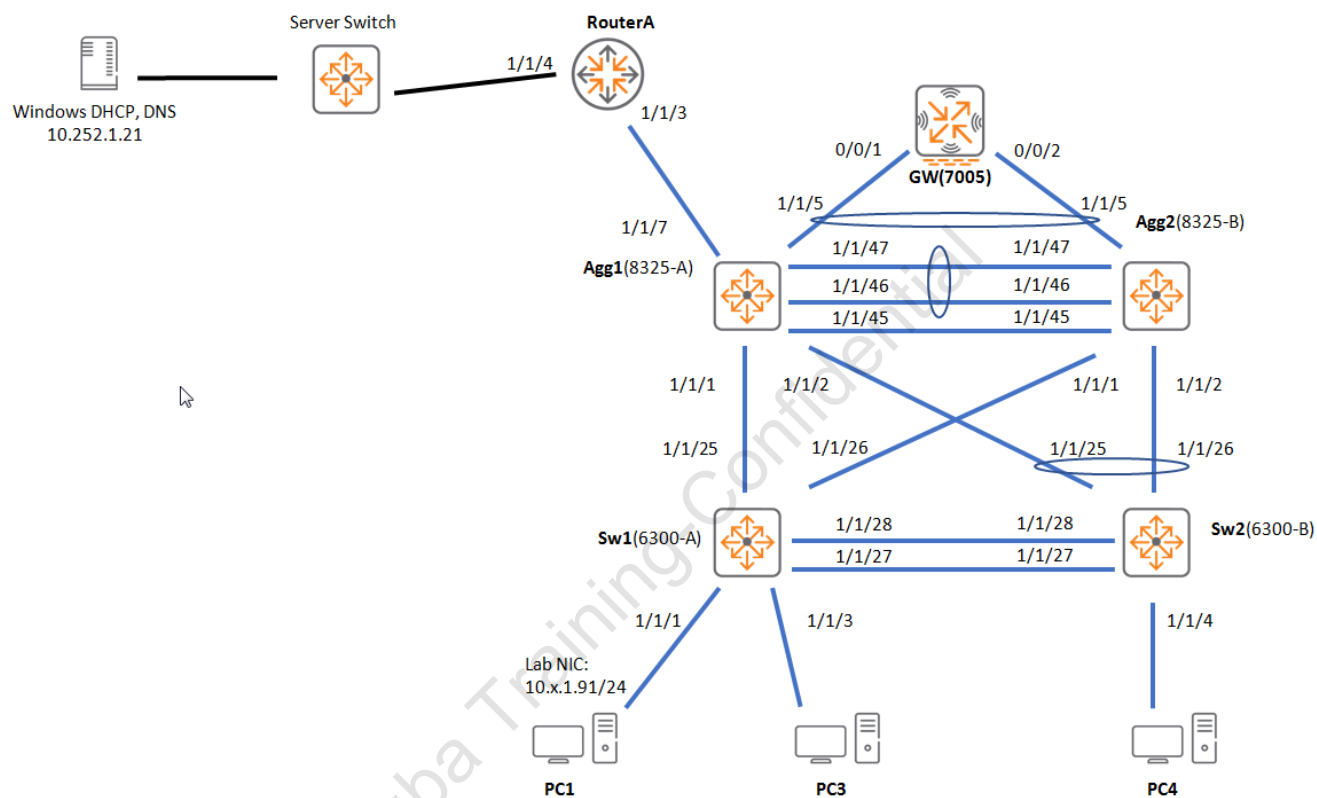
1. The customer wants to have STP on the VSX pair, with the VSX pair acting as the root bridge. Since Agg1 and Agg2 are independent switches, they want to understand if they should be configured with the same or a different STP priority.
2. They want you to explore which system is transmitting the STP BPDUs over the VSX LAGs.
3. On the VSX, they want you to demonstrate the STP root-guard feature on the LAG2 to Sw2, so when Sw2 attempts to become root bridge, the LAG2 will be blocked.
4. On the Sw2, they want you to demonstrate the STP loop-guard feature on the LAG2 to the VSX, so when STP is stopped or shutdown on the VSX, LAG2 will transition to a blocking state.
5. On the Sw2, they want to ensure that the access port 1/1/4 transitions to a forwarding state quickly.
6. They want to understand the real time difference between STP admin-edge and admin-network before an access port transitions to a forwarding state.

7. They have had many issues with TCNs in the past, so they also want you to investigate if and when a TCN is generated on access ports with admin-edge and the default admin-network configuration.
8. The customer wants a demonstration of the BPDU-Guard feature on Sw2. It should be enabled on the ports 1/1/4 and 1/1/27. Port 1/1/27 connects to Sw1, since it generates STP BPDUs, the port 1/1/27 on Sw2 should be disabled by the BPDU-Guard.
9. They want the disabled port to automatically recover after 30 seconds.

Loop Protection

1. The customer has had issues with phones that were blocking the STP BPDUs, so when a loop was created between 2 IP phones, the STP BPDU-Guard could not detect the loop.
2. They want you to demonstrate any other loop prevention protocols available on top of STP.
 - a. To demonstrate this, apply a BPDU-filter on Sw1 ports 1/1/27 and 1/1/28. This will cause a loop that cannot be detected by the BPDU-Guard on Sw2.
 - b. Note: Make sure to shutdown Sw1 port 1/1/1 and 1/1/3, as well as Sw2 port 1/1/4 before you create this loop. Otherwise the loop may also affect the lab backbone infrastructure devices.
3. They want any looped ports to automatically come back online after they have been shut down for 30 seconds.
4. They want to see the counters and debug output of this feature when the loop is detected.

Lab Diagram



Task 1: Configure VSX

Objectives

- Configure VSX on Agg1 and Agg2

Steps

VSX Configuration

1. Configure VSX between the Aggregation 1 and Aggregation 2 switches using the information in the table to configure VSX

Parameter	Value	
Primary	Agg1	
Secondary	Agg2	
ISL	LAG256 (1/1/46 1/1/47) LACP	
Keepalive	1/1/45	
Keepalive VRF	KA	
Keepalive IP Subnet	192.168.0.0/31	

2. Configure Agg1 for VSX
3. Configure Agg2 for VSX
4. Verify that the VSX is operational on Agg1 switch
5. Review the VLAN status of the ISL port (LAG256)

Q: What do you observe? What is the Mode for VLAN1?

A: The mode is native-tagged.

6. Review the running configuration of LAG256

Q: What does this mean?

A: This means that the native VLAN can be either untagged (default) or tagged (current configuration) on a VLAN trunk. This can be used by customers that do not want to have untagged traffic on their VLAN trunks, without sacrificing a dummy untagged VLAN.

7. Configure VSX management synchronization on Agg1 switch for the following features: vsx-global, dhcp-relay, lldp, mclag-interfaces, qos-global, stp-global

VSX SVI Active Gateway - SVI 1

8. On Agg1, configure SVI 1 with Active Gateway, enable the vsx-sync command for the active gateway settings.

NOTE: SVI 1 has interface IP address on Agg1 (10.x.1.2) and Agg2(10.x.1.3).

VLAN ID	Active Gateway MAC	Act GW IP
1	02:00:00:00:00:01	10.x.1.1

VSX SVI Active Gateway - SVI 12

9. On Agg1, define Layer 2 VLAN 12, use the vsx-sync function to sync it to Agg2.

10. On Agg1, configure SVI 12 with Active Gateway, enable the vsx-sync command for the active gateway settings.

VLAN ID	Active Gateway MAC	Act GW IP	Agg1 IP
12	02:00:00:00:00:01	10.x.12.1	10.x.12.2/24

11. Verify the configuration is active on both nodes from Agg1.

Q: Did the configuration of SVI 12 appear on Agg2?

A: No, the VLAN 12 was synced, but the SVI 12 did not appear.

12. Review the VSX status of the configuration sync

13. Review the differences in the running config

Q: What do you observe?

A: The SVI 12 interface is missing on the secondary node, it must be manually defined.

14. Define SVI 12 on Agg2, assign IP 10.x.12.3/24.

15. Repeat the VSX config sync checks on Agg1, they should be ok now.

16. Verify **on both Agg1 and Agg2** that you can reach the DHCP server on 10.252.1.21, using the VLAN 1 source IP address.

17. On Agg1, configure DHCP Relay for SVI 12.

Q: Is DHCP Relay any different on a VSX cluster compared to 2 standalone Layer3 switches?

A: DHCP Relay on a VSX cluster will only be performed by the primary node. The secondary node will take over the role when the primary node is down. This will ensure the DHCP server receives the DHCP information only once.

18. NAE: On Agg1, use 'show interface brief' to verify that the port descriptions have been updated by the LLDP NAE script. Notice the special descriptions for the ISL and the Keep Alive links.

VSX LAG to Aruba Gateway

The Aruba Gateway has been pre-configured with a port-channel with LACP. You should configure the VSX system with a LAG to the Aruba Gateway.

19. On Agg1, configure VSX LAG 5 with LACP to the Aruba Gateway, connected to port1/1/5.

Setting	Value
Port	1/1/5
LAG ID	LAG 5
Native VLAN untagged	1
Allowed VLANs	1,11,12

20. On Agg2, define the LAG5 as multi-chassis. Verify VSX-sync completed the VLAN sync.

21. On Agg2, assign port 1/1/5 to the LAG5 and enabled the port.

22. On Agg1, verify the status of the LAG5 is 'up', both locally and on the VSX peer

VSX LAG to Sw2

23. On Agg1, configure VSX LAG with LACP to Sw2, activate the vsx-sync for the VLAN configuration. Allow VLANs 1,11 and 12.

Setting	Value
Port	1/1/2
LAG ID	LAG 2
Native VLAN untagged	1
Allowed VLANs	1,11,12

24. Verify the configuration on Agg1.

25. Review the 'show interface brief' output.

Q: What is the current LAG2 status?

A: Down, since Sw2 has not been configured yet.

VSX LAG LACP Status

In this section you will explore the impact on the VSX LAG when there are LACP issues. For example, 1 side has LACP configured, while the other side has a static LAG without LACP.

26. On Sw2, configure ports 1/1/25 and 1/1/26 with a LAG uplink to the VSX core. Use LAG ID 255. Define VLANs 11 and 12. Allow VLANs 1,11 and 12 on the LAG. Do not enable LACP.

27. On Agg2, create LAG 2 multi-chassis interfaces and assign port 1/1/2 to the LAG2.

28. On Sw2, review the LAG status. Both Sw2 interfaces will be 'up' (forwarding).

Q: Do you see an 'Actor' system-id?

A: No, the system-id is only shown for an LACP enabled LAG.

29. On Agg1, verify the LAG status, check the VSX peer as well.

Q: What is the current status?

A: blocked. Why?

A: When AOS-CX is configured with an LACP enabled LAG, it will not enable the interface, unless LACP data units are received from the peer device.

NOTE: This is a different behavior compared to the AOS-Switch based switches.

30. Use diag-dump to review if any LACP PDU have been received on the LAG2

Adjust the LAG to LACP

31. On Sw2, enable LACP mode active on the LAG255.

32. On Agg1 and on vsx-peer, verify the LAG interfaces are in the UP and forwarding state.

33. On Agg1, use 'diag-dump lag basic' to verify LACP PDUs are now being received

Client PC Validation

34. On Sw2, assign port 1/1/4 (connected to PC4) to VLAN12

35. On PC4, verify you have an IP address in the 10.x.12.0/24 range.

VSX DHCP Relay Services

36. On Agg1, review DHCP relay statistics using the command 'show dhcp-relay' and repeat for the Agg2 VSX peer.

Q: What do you notice?

A: Only the primary node statistics have increased. This is the expected behavior.

37. On **both Agg1 and Agg2**, enable debugging for dhcprelay, enable terminal monitor for debug.

38. On PC4, release and renew the IP address.

```
ipconfig /release  
ipconfig /renew
```

39. Review the debug output on Agg1 and Agg2.

Q: What do you observe?

A: Only Agg1 DHCP relay is processing the requests.

40. Disable the debugging.

Impact of Primary node SVI Shutdown

In this section you will explore the effect of shutting down an SVI interface on the primary VSX node. This will have impact on the traffic and the DHCP Relay function.

Impact of Primary node SVI Shutdown - DHCP Relay

DHCP Relay is active on the primary VSX node, and the secondary will take control when the primary node is unavailable.

However, keep in mind that this may affect DHCP operation when the SVI is disabled on the primary node.

41. On PC4, perform an ipconfig /release in an administrator command prompt.

42. On Agg1, shutdown the VLAN 12 SVI.

43. On PC4, perform an ipconfig /renew.

Q: What is the impact of shutting down a Layer3 VLAN interface on the primary node?

A: This impacts the DHCP relay function.

Q: Did you observe any DHCP relay messages in the debug?

A: No, since only the primary VSX node is supposed to process the DHCP relay packets.

44. On Agg1, enable the VLAN 12 SVI again, renew the IP on PC4.

45. Disable all debugging **on both Agg1 and Agg2**.

Impact of Primary node SVI Shutdown - Traffic

The Active Gateway MAC address is programmed into the ASIC for the VLAN. When the Layer3 VLAN interface is disabled, the Layer2 VLAN is still up. However, the node will not pass traffic for the Active Gateway MAC over the ISL to the node where the Layer3 interface is up.

This section starts with both Agg1 and Agg2 VLAN 12 SVI interfaces in the UP state.

46. On PC4, ping to 10.252.1.23 and 10.252.1.24. These are 2 IP addresses of the ClearPass server. Due to a difference of 1 between these 2 IP addresses, the LAG hashing on Sw1 will send 1 ping over the link to Agg1, and the other ping over the link to Agg2. (you don't know exactly which ping goes to which Agg, but you do know that they go to different Aggregation switches).

47. Verify both pings were successful.

48. On Agg1, shutdown the Layer3 VLAN12 SVI.

49. On PC4, attempt both pings again.

Q: What is the impact of shutting down a Layer3 VLAN interface on the primary node?

A: This impacts the traffic forwarding for traffic that requires the default gateway.

50. On Agg1, enable the VLAN12 SVI again.

51. On PC4, verify the ping succeeds again.

Virtual Active Gateway vs Active Gateway

AOS-CX also supports the virtual Active Gateway function. This should only be used when it is not required for each VSX member to have a unique identity on the network. You should not use this feature for example when the switches need to run OSPF for example on that same Layer3 interface, since that requires each switch to have its own IP address.

The typical scenario for this single IP Active Gateway is subnets where only endpoints are connected.

VLAN ID	Active Gateway MAC	Act GW IP	Agg1 IP	Agg2 IP
11	02:00:00:00:00:01	10.x.11.1	10.x.11.1/24	10.x.11.1/24

NOTE: The table shows the same IP for the Active Gateway IP and the Agg1/Agg2 IP. This is not a typo, this is the point of 'virtual' Active Gateway (to use the same IP as interface and gateway IP).

52. On Agg1, configure SVI 11 to use 10.x.11.1/24 as both the Interface IP and the Active Gateway IP.

53. On SVI 11, enable vsx-sync for the active gateway function.

54. On SVI 11, configure the DHCP Relay

55. On Agg2, define the SVI 11 with IP 10.x.11.1/24.

Q: What do you observe?

A: Agg2 shows an error that active gateway is using the current subnet. (means the currently active IP cannot be changed since Active Gateway is using it). This is because the interface was previously configured with the 10.x.11.3/24 IP address.

56. Attempt to correct this on Agg2:

- remove the active gateway IP
- attempt to remove the IP address 10.x.11.3/24

Q: What do you observe?

A: Even when the Active Gateway configuration is removed, the VSX sync adds it almost immediately, so it is not possible to remove the IP address.

57. Resolve the issue by:

- On Agg1: disable the vsx-sync on the SVI 11
- On Agg2: remove the active gateway configuration and the current IP
- On Agg2: assign the 10.x.11.1/24 IP
- On Agg1: enable the VSX-sync for SVI 11

58. Verify that VSX-sync has pushed the rest of the active gateway and DHCP relay configuration.

You have now completed the 'Virtual' Active Gateway configuration.

Verify the configuration and troubleshoot the Layer2 connection

Sw1 will have a 'routing' role in most lab activities, while Sw2 will have the 'Layer2 Access Switch' role. However, since PC3 endpoint is connected to port 1/1/3 on Sw1, you will now connect PC3 over the port 1/1/27 to Sw2 so it 'appears' to be connected on Sw2 port 1/1/27.

59. On Sw1, define VLAN11. Configure port 1/1/27 (to Sw2) as VLAN trunk, allow VLAN 11 and VLAN 1 (leave VLAN1 as the default native VLAN). Enable the port. Assign interface 1/1/3 as access port to VLAN 11.

60. Review the VLAN port configuration of port 1/1/27 on Sw1.

61. On Sw2, define VLAN 11, configure port 1/1/27 as VLAN trunk, allow only VLAN 1 and 11 on the VLAN trunk. Enable the port.

62. On PC3, release and renew the IP address to attempt to confirm the operation.

Q: What do you observe?

A: The PC3 is not getting an IP address.

63. On Sw1, attempt to troubleshoot the issue, start from the physical layer and work your way up:

Is the interface UP?

Is there an LLDP neighbor?

Does the VLAN contain the correct ports?

Is the VLAN tagged or untagged status correct?

Does the switch learn any MAC addresses on the ports? (Active Gateway will send gratuitous ARP packets, so the Active Gateway MAC address should be listed in the MAC tables)

Is there any Layer2 protocol that is impacting the forwarding?

NOTE: You may need to repeat these steps on the Sw2.

64. This should reveal that Spanning-Tree is blocking the port 1/1/27 between Sw1 and Sw2.

NOTE: Even though there is no Spanning-Tree running on the Aggregation switches at this point, you should realize that they will simply forward Spanning-Tree BPDUs as 'data frames'. This is why Sw1 and Sw2 appear to have 2 links from a Spanning-Tree point of view.

65. Since Sw1 will have a routing role, it will be Layer3 connected to the Agg1 and Agg2, so you should shutdown ports 1/1/25 and 1/1/26 for now.

66. On PC3, renew the IP address, this should now be successful.

67. On Sw1, verify the MAC-address table again for VLAN 11, MAC addresses should be learned on both 1/1/3 and 1/1/27.

Task 2: Apply some VSX Best Practices

Objectives

- Apply System MAC
- Apply MTU
- Review the link-up delay timer
- VSX LAG MTU
- Review VSX LAG Hashing algorithm

Steps

1. On Agg1, review the default VSX system MAC.

Q: Why is it recommended to apply a user-defined system MAC?

A: This simplifies the hardware replacement of the VSX primary switch. Without the user-defined system MAC, the VSX cluster would change MAC address to the replaced VSX primary switch, this would affect STP and LACP peer communication.

2. What would be the impact on a production network when you change the system-mac?
3. **On Sw2 and Agg2**, check the LAG status and repeat the command every 5 seconds.
4. On Agg1, apply a user-defined VSX system MAC.

NOTE: Since the ISL is temporarily 'broken', VSX will be out of sync. Once the LAG 256 is re-established, the VSX sync will take a few moments to apply the VSX system MAC to the secondary node. Only then, the VSX LAG to the Sw2 can revert to the UP state again.

5. On Sw2, verify the LACP interface status.

Q: Did the Partner System-ID change?

A: Yes, the VSX system mac is now used for the Layer2 protocols.

Q: Was there any impact?

A: Yes, VSX ISL LAG is reset, this requires the VSX systems to re-synchronize once the LAG256 is up again.

6. On both Agg2 and Sw2, stop the repeat using CTRL-C key combination.

MTU Configuration

7. You have already applied the jumbo MTU to the Agg1 and Agg2 Layer2 interfaces in the previous lab.
8. **On both Agg1 and Agg2**, apply the maximum Layer3 MTU on VLAN1, 11 and 12.
9. Configure all ports **on both Sw1 and Sw2** to support jumbo frames.

Uplink delay-timer

10. On Agg1, review the VSX link uplink-delay timer.

Q: What is the purpose of the VSX uplink-delay timer?

A: Provides time for the table sync between primary and secondary for large tables.

Q: In which scenario would a LAG be excluded from the uplink delay timer process, so it transitions immediately to forwarding?

A: In case a VSX LAG is intended for Layer3 communication with a peer device. This will enable the Layer3 adjacencies (routed uplinks) to be established before the Layer2 intended VSX LAGs (switched downlinks) come online.

TIP: The linkup-delay timer can be adjusted based on the number of entries that need to be synced between the VSX systems.

See this link in the 10.05 configuration guide for example numbers:

<https://www.arubanetworks.com/techdocs/AOS-CX/10.05/HTML/5200-7325/index.html#GUID-A61A012E-C780-4FE8-806B-3E63D40E4660.html>

Validate the linkup delay-timer

Review the normal state first

11. On Agg2, review the VSX MAC address table and check the table for the VSX peer.
12. There is a convenient command that groups these 2 commands:

Disable and re-enable the ISL

13. On Agg2, prepare the status output views. This will allow you to see the transitions of the linkup-delay on Agg2.
14. On Agg1, disable the LAG 256 of the ISL.

15. On Agg1, enable the LAG 256 again.

16. On Agg2, review the status changes.

Q: Did the switch have to wait the complete Linkup-delay time?

A: No, after about 40 seconds, it will be reported as completed.

Q: Why is this different from the configured timer?

A: When the table synchronization is complete, there is no need to wait for the timer, so the ports transition to forwarding.

Review impact on the VSX LAG interfaces

In this section you will repeat the ISL failure and review the impact of the linkup-delay timer on the interfaces.

17. On Agg1, bounce the LAG 256 again (shutdown/no shutdown).

18. On Agg2, use the command 'show interface brief' to see the status of interfaces and the Reason. Look for the 1/1/2 and 1/1/5 interfaces (VSX LAG to Sw2 and to the Aruba Gateway).

NOTE: It may take a few moments before VSX is established and the command can be run. Just repeat the command until it works.

Q: What is the reason port 1/1/2 and 1/1/5 of the VSX peer is down?

A: During the linkup-delay, the reason will be shown as 'Disabled by VSX'

Q: Pay attention to the SVI interfaces, what is their status during the linkup-delay?

A: SVI interfaces that are bound to a VLAN that is enabled on a VSX LAG will also be in the 'Disabled by VSX' state. This means that both the Layer2 LAG and the Layer3 SVI are currently disabled by VSX.

19. After about 40 seconds, the sync should be complete and the ports should transition to forwarding.

Linkup-delay exclusion

Note that any VSX LAG that connects to a routing peer, such as an OSPF peer, can be excluded from the linkup-delay. This ensures that the OSPF adjacency can be formed and routes can be learned, even while the Layer2 VSX LAG interfaces are still waiting to synchronize the Layer2 tables.

You will attempt to configure this for the VSX LAG with the Aruba Gateway.

First you will exclude the LAG5 and verify why the Layer3 communication still does not work.

Next you will correct that configuration and verify.

20. On Agg2, configure the linkup-delay-timer to exclude the LAG 5 to the Aruba Gateway.
Notice that this command is local to the Agg2.

21. Verify the configuration on Agg2.

22. Agg2 is now ready for the next test. On Agg1, bounce the LAG 256.

23. On Agg2, review the interfaces, you may need to repeat the command a few times.

Q: Do you see any difference between the ports 1/1/2 and 1/1/5?

A: 1/1/2 is still disabled by VSX, while 1/1/5 is in the 'up' state. The linkup-delay exclude for the LAG5 has worked fine.

Q: What is the status for the SVI 1?

A: The 'vlan1' interface still shows as Disabled by VSX.

Q: Would Agg2 be able to establish a Layer3 adjacency or peering using this SVI?

A: No.

Q: What could be the reason that the SVI 1 is still kept down during the linkup-delay timer?

A: When the VLAN is enabled on any VSX LAG that is still kept 'Disabled' during the timer, the Layer3 SVI interface for that VLAN will also be kept in the 'Disabled' state. In this lab, the VLAN 1 is still enabled on the VSX LAG 2, this is why the SVI 1 remains down.

NOTE: In a real deployment, make sure that the upstream SVI VLAN ID is not enabled on any VSX LAG that connects to the Access Layer switches. This is a strong argument against the use of the 'vlan trunk allowed all' command on the VSX system.

Adjust the configuration

24. On Agg1, remove VLAN 1 from the LAG2 allowed list. This will be synced to Agg2.

25. On Agg1, bounce the LAG 256 to test.

26. On Agg2, review the status.

Q: Did the status for SVI 1 (vlan1) change?

A: Yes, it is now enabled as soon as the VSX cluster forms. This is the behavior that would be needed for an 'uplink' Layer3 SVI interface.

Q: Do you see the difference with the SVI 11 and SVI 12?

A: Yes, these are 'downstream' SVI interfaces, they should remain down until the sync completes.

Restore the VLAN 1 configuration

This section showed how to configure an upstream SVI, in this lab, SVI 1 will be used by Sw2 to get management access to the network, so enable the VLAN 1 again on the LAG2.

27. On Agg1, allow VLAN 1 on the LAG2

VSX LAG Hashing algorithm

Review the VSX LAG hashing algorithms

28. On Agg1, review the default algorithm on the LAG2

Q : What is the default hashing method?

A: I3-src-dst

29. Review the algorithm options: enter the interface LAG2 context and use the 'hash' command to explore the options.

Q: What are the available hashing options?

A: I2-src-dst, I3-src-dst and I4-src-dst

Task 3: Configuring STP with VSX

Objectives

- Configure and review STP with VSX
- Review STP on the Access switches

Steps

VSX operates with independent control planes, but for some protocols, such as STP, the 2 nodes will be syncing their operational state so they can appear as 1 device for the Layer2 peers.

Configure STP

1. Review the STP status **on both Agg1 and Agg2.**

Q: What is the default status of STP on the switches?

A: On the 83xx series, STP is disabled by default.

2. Configure STP priority 1 on Agg1
3. Configure STP priority 1 on Agg2.
4. On both Agg1 and Agg2, enable STP.

NOTE: Best practice is to configure priority 0 on the root. In this lab, priority 1 is used, so another switch may attempt to become root.

5. On Agg1 and Agg2, use 'show spanning-tree' to review and compare the status.

Q: What do you observe?

A: Both Agg1 and Agg2 spanning tree processes state that they are the root (This bridge is the root). This is the expected behavior in a VSX topology.

VSX STP BPDUs Transmissions

In a VSX system, the primary VSX node is responsible to transmit the STP BPDUs over the VSX LAG interfaces. For any interfaces that are only connected on the secondary node, the secondary node will transmit the BPDUs.

6. On Agg1, use the 'show span' command and check the BPDU-Tx statistics for LAG2.
7. On Agg1, use the 'show span vsx' command to see the BPDU-Tx statistics for LAG2 on Agg2.

Q: What do you notice?

A: Only the Agg1 transmits BPDUs on the VSX LAG interface. This demonstrates that only the primary VSX node transmits STP BPDU frames on the VSX LAG interfaces.

STP features - Root-guard

8. Configure Agg1 with STP root-guard on the VSX LAG interfaces towards the Access switch Sw2 (LAG2)

Q: What is the purpose of root-guard in an STP environment?

A: It ensures that STP operates normal with the peer device, as long as the peer device does transmit a superior BPDU. This effectively protects the current root bridge.

Validate the root-guard using Sw2

9. Configure Sw2 with STP priority 0
10. On Agg1, verify the local STP status and the VSX peer STP status.

Q: Is the Agg1 still root bridge?

A: Yes, the root-guard feature should have prevented Sw2 from becoming root bridge, even when Agg1 has a priority 1, that is higher than the Sw2.

Q: What is the State of lag2 interface?

A: LAG 2 is reported as 'Root-Inc' (Inconsistent)

11. On Agg1, check the inconsistent ports, this will reveal the reason.

Q: What is the reason?

A: Root Guard.

12. Review the last 20 events on Agg1 and filter on the hpe-mstpd daemon.

Restore the configuration

13. On Sw2, set the default STP priority.
14. On Sw2, verify you can now ping 10.x.1.2 and 10.x.1.3 again.

VSX configuration consistency check

Since both Agg1 and Agg2 have an independent configuration, it is important to keep their STP configuration state in sync. In this lab, this is done automatically by VSX sync.

In case you would not have the VSX sync active, you can easily compare some key configuration elements using the configuration consistency command.

15. On Agg1, review the current VSX configuration consistency and repeat the command.

16. On Agg2, change the STP mode rpvst.

17. On Agg1, review the updated VSX configuration consistency check

18. Wait a few seconds, the VSX config sync (you enabled STP configuration sync) should automatically correct this for you.

NOTE: If the mode is not synchronized automatically, you can disable and enable the LAG 256 to trigger a full configuration sync.

Configure and review STP on the Access switch Sw2.

In this section, STP features on the access switch will be explored.

19. On Sw2, verify the status of STP in the running configuration

Q: Is STP by default enabled?

A: On the 6300 switches, STP is by default enabled. Note that the existence of the command in the configuration does not mean that it is not default.

Q: What is the root port for STP?

A: The root port is the LAG to the VSX system.

STP Loop-guard

STP Loop-guard ensures that a port on which BPDUs are expected will transition to blocking when no BPDUs are received. A normal STP port would transition to forwarding when no BPDUs are received. This feature can be used in large environment to handle the scenario where the STP process on the Agg/Core switches would fail and would cause loops in the network.

You can easily test this by blocking STP BPDUs on the Agg switches to Sw1. Sw1 will no longer receive BPDUs and should then block the port that has the loop-guard enabled.

NOTE: Do not confuse this feature with Loop Protect . Loop Protect is a feature that is independent of STP and that can be used to detect loops between ports.

20. On Sw2, configure STP Loop-guard on the uplink LAG 255 to the VSX

Validating STP features - Loop-guard

You can now test a 'failure' of STP on the Agg switches by applying a BPDU filter on their ports to the Sw2.

21. On Sw2, enabled terminal-monitor with the INFO severity level.
22. On Agg1, enter the LAG2 context and apply the BPDU-filter.
23. On Sw2, review the events that were reported.
24. On Sw2, review the spanning-tree state of the port LAG255.
25. On Sw2, review the STP inconsistent ports
26. After you have confirmed the feature, restore the configuration by removing the BPDU-guard on the Agg1 LAG2 port.
27. On Sw2, disable terminal monitor and verify that port lag255 is now forwarding again.

Configure the access ports

STP has 2 port types: admin-edge and admin-network. The default spanning-tree port-mode in AOS-CX is admin-network. In the next section you will see how the port will transition when it comes up.

28. Open a second SSH connection to Sw2, run the 'show span | i 1/1/4' command to see only port 1/1/4 state and repeat the command
29. In the first SSH connection, bounce the port 1/1/4 (shutdown > no shutdown).
30. Follow the state in the second SSH connection

Q: How long does it take for an admin-network port to transition to forwarding on an AOS-CX switch?

A: When AOS-CX does not detect a peer STP device (no received BPDUs), it will transition to forwarding after about 4-5 seconds.

NOTE: On AOS-Switch based devices, this feature was known as 'auto-edge'. In AOS-CX, this functionality is part of the admin-network port role.

Topology Change Notifications

TCN packets are required to notify the switches in an STP domain about a topology change. The switches can then clear their Layer2 tables and ensure that MAC addresses are learned on the correct (possibly new) ports.

While this is good for backbone topology changes, it cause unnecessary load on the network if a TCN would be transmitted for every access port that would go up or down.

In this section you will explore the TCN impact.

31. On Agg1, enable debugging of MSTP events. Enable terminal-monitor for debugging.
- Verify TCN generation for access ports

32. On Sw2, bounce the port 1/1/4.

33. On Agg1, review the messages.

Q: Was there any Topology Change received by the Agg1?

A: No, since Sw2 has detected that the port 1/1/4 is an access port (based on the auto-edge principle), it will not generate a TCN for this port.

Verify TCN generation for STP connected ports

34. On Sw2, bounce the port 1/1/27. This is the port that connects to the Sw1.

35. On Agg1, review the messages.

Q: Was there any Topology Change received by Agg1?

A: Yes, since Sw2 received STP BPDUs on the port 1/1/27, it has generated a TCN packet to the rest of the network.

36. On Agg1, review the number of received TCN packets (TCN-Rx). This is a good place to start the troubleshooting TCN issues. If ports are receiving continuous TCN updates, you need to check the STP switch that is connected to these ports.

37. This demonstrates how AOS-CX handles both scenario's well with the default STP configuration state.

38. On Sw2, close the second SSH connection.

39. On Agg1, disable the debugging.

BPDU-Guard configuration

BPDU Guard is typically an access switch feature to ensure that no other STP devices can be connected to the network.

When a BPDU guard enabled port receives an STP BPDU, it will be disabled.

40. On Sw2, enabled terminal-monitor for INFO severity.

41. On Sw2, set the global spanning-tree BPDU-Guard timeout to 30 seconds.

42. On Sw2, configure the access port 1/1/4, port 1/1/27 and 1/1/28 with BPDU-guard. Port 1/1/4 connects to the PC, while port 1/1/27 and 1/1/28 connects to Sw1 (that still has STP enabled on these ports). Port 1/1/28 is currently down and will only be used in the next section.

43. On Sw2, bounce the port 1/1/27 and watch the events.

Q: What happened with the port 1/1/27?

A: Sw2 received a BPDU from Sw1, so the BPDU Guard has disabled the port. 30 seconds later the port is enabled again and the port will be disabled again.

44. On Sw2, review the spanning tree port states

45. You may leave the BPDU-Guard enabled on Sw2 for now.

Aruba Training-Confidential

Task 4: Configuring Loop-Protection in the Access Layer

Objectives

- Configure a network loop that Spanning Tree does not detect.
- Configure Loop protection on the Access layer switches to detect the loop.

Steps

Loop Setup

Sw1 will be used to test the loop protect feature. By configuring the BPDU-filter on Sw1, Sw1 will be blocking the STP BPDUs, so STP BPDU-guard on Sw2 will not be able to detect the loop.

1. On Sw1, shutdown ports 1/1/1 and 1/1/3.

IMPORTANT: To limit the loop impact Sw1 must shutdown port 1/1/1 (PC1) and 1/1/3 (PC3) to prevent the looped traffic to enter the VM infrastructure.

2. On Sw1, verify that the ports 1/1/1 and 1/1/3 are down.

Configure the network loop without Loop Protect

First you will make a network loop that STP cannot detect without the Loop Protect feature. In the next section you will use Loop Protect to detect the loop.

3. On Sw1, configure BPDU-filter on ports 1/1/27 and 1/1/28. This ensures that STP on Sw2 will not be able to detect and block the loop, and it simulates an edge device that 'consumes' STP BPDU frames.
4. On Sw1 only, enable these 2 ports. Port 1/1/27 was already enabled, this applies to port 1/1/28. The 'looped' device is now ready.
5. On Sw2, review the current (default) loop protect interface configuration and status. Loop protect is not enabled by default.

Enable the loop

6. On Sw2, enable port 1/1/28, wait about 10 seconds.
7. On Sw2, verify the port status of the ports 1/1/27 and 1/1/28, both should be 'up'
8. Clear the global statistics on Sw2 and review the non-zero port statistics. You should now see many received and transmitted packets on the interfaces due to the loop.

Loop Protection configuration

9. On Sw1, configure the global loop-protect re-enable timer to 30 seconds.
10. On Sw1, enable loop-protect on 1/1/27 and 1/1/28.

11. Observe how the ports are disabled by loop-protect.
12. Review the last 20 events on Sw1, filter for the hpe-lpd daemon.
13. Use 'show interface brief', check the 'Reason' description.

Restore the configuration

14. On Sw2, disable port 1/1/28
15. On Sw1, after about 30 seconds, verify the port states of ports 1/1/27 and 1/1/28. They are no longer disabled by loop protect.
16. On Sw1, enable port 1/1/1 (PC1) and 1/1/3 (PC3).
17. Optional: On Agg1, Agg2, Sw1 and Sw2, you can make a new checkpoint named 'asts-lab04-<your-name>'.

You have completed this Lab!

Lab 05: Review of Layer3 Basics

This lab activity reviews OSPF configuration and troubleshooting, OSPF in a VSX environment and Policy Based routing.

Requirements

This lab requires completion of Lab 04.

Objectives

- Debugging OSPF features
- Configuring OSPF for VSX with Active Forwarding and max-age on startup
- Policy Based Routing

Scenario

OSPF

The customer wants you to configure OSPF area 0 in the Main location between Agg1 and Agg2. You should enable SVI 1 and the Loopback interfaces.

1. SVI 11 and SVI12 must be advertised into OSPF but Agg1 and Agg2 should not form an OSPF Adjacency on these interfaces.
2. On Agg1, a tagged SVI41 (1/1/1) must be created and connect to Sw1 (1/1/25).
3. On Agg2, a tagged SVI42 (1/1/1) must be created and connect to Sw1 (1/1/26).
4. Both links must be enabled in OSPF area 0. Sw1 must be configured for OSPF area 0 as well. When the interfaces 1/1/25 and 1/1/26 are enabled, you may notice that the OSPF adjacency does not form on the 1/1/26. If you see this behavior, you should troubleshoot the issue and correct the configuration.
5. The customer wants to see the debug log of the normal OSPF adjacency setup between Sw1 and Agg1. They also want to see the debug log or show command output to detect various OSPF adjacency errors on the link between Sw1 and Agg1 including:
 - a. different OSPF hello timers
 - b. dead-timers
 - c. invalid authentication or authentication type
 - d. conflicting router-id between 2 peers
 - e. control plane trace of the OSPF traffic on Sw1.
6. The customer wants to configure the SVI41 and SVI42 links so that the OSPF adjacency comes up fast.

VSX Active Forwarding

The customer has heard that Active Gateway can be used for active/active forwarding of traffic from clients or static routes. However, when an OSPF peer is connected to a VSX LAG, the peer will use the OSPF next-hop IP, which will be unique for each VSX node. They are aware traffic may arrive on the wrong VSX node.

1. Configure a tagged SVI40 on the LAG5 connected to the Aruba Gateway and enable it for OSPF area 0. The Aruba Gateway has been pre-configured with OSPF.
2. Configure VSX so that traffic received from the Aruba Gateway OSPF peer can be handled by either VSX node.

OSPF Key-chain and tuning (Optional)

The customer is interested in the following:

1. On the link between the Sw1 and Agg1, the customer wants to enable OSPF authentication using a key-chain and key 1 aruba123. They also want a demonstration of a live migration to key 2 ArubaRocks without any impact on the OSPF adjacency.
2. The customer wants to ensure that when a VSX node is rebooted, the node will make itself the least preferred path to the other OSPF peers for 1 minute. They want to see the output of the actual change on the OSPF route costs before and after this minute.

Policy Based Routing (PBR)

The base configuration for the customer's PBR is:


- Note: a template configuration is available for these steps if you do not want to configure this manually. It is on PC1, Desktop > ASTS > configs > ASTS_cfg_lab05-steps-v2.txt (Task6 section).
- The customer has a branch office with 2 links to the Main location. In the lab, the branch office will be simulated using a VRF on Sw2, the Sw1 operates as the routing switch for the Main location.
- On the link 1/1/27 between Sw1 and Sw2, configure SVI46 and SVI47, the customer wants L3 counters on both SVIs.
- On Sw2, the customer wants you to:
 - a. configure VRF 'branch' with SVI46 and SVI47
 - b. enable OSPF area 0 in the VRF 'branch'.
 - c. assign SVI46 OSPF cost 10
 - d. assign SVI47 OSPF cost 100.
 - e. configure a DHCP server in the VRF 'branch' for subnet 10.x.15.0/24
 - f. assign port 1/1/4 (PC4) to VLAN 15
 - g. verify PC4 is connected to the network.

PBR Path selection

To verify PBR path selection:

1. Use PC4 to verify that a trace to 10.x.12.1 takes the shortest path
2. Use Sw2 to verify the selected SVI by checking the L3 counters.
3. The customer wants to verify that OSPF failover works properly before starting the PBR configuration.
 - a. They notice that the failover of the traffic is very slow after a shutdown of the SVI46.
 - b. Investigate this and adjust the configuration so that PC4 traffic to 10.x.12.1 recovers within 3 seconds after a shutdown of SVI46 on Sw1.
4. After link failover validation, the customer wants to see PBR active for the branch SVI15.
 - a. Any traffic from the branch to 10.x.12.0/24 subnet should take the SVI47 path.
 - b. Any other traffic from the branch should take the default best path on SVI46.
 - c. Failover should still work fine for all of the traffic.
5. The customer wants to see this demonstrated based on the L3 statistics on the SVIs.

Lab Diagram



Task 1: Configure VSX with OSPF

Objectives

- Configure OSPF between the VSX nodes

Steps

OSPF Configuration on Agg1 and Agg2

Lab Diagram - Layer3 Topology



1. On Agg1, configure loopback interface 0.

Device	IP
Agg1	10.x.0.2/32
Agg2	10.x.0.3/32

2. On Agg1, configure OSPF area 0. Use 10.x.0.2 as the router-id.
3. On Agg2, configure OSPF area 0. Use 10.x.0.3 as the router-id.
4. On Agg1, configure and enable the loopback 0 and interface VLAN 1 for OSPF.
5. On Agg2, configure and enable the loopback 0 and interface VLAN 1 for OSPF.
6. On both Agg1 and Agg2, enable the SVI 11 and SVI 12 for OSPF and configure them as passive interfaces.
7. Wait about 30 seconds for the adjacencies to establish, then verify the OSPF LSDB

Q: How many LSA entries do you expect?

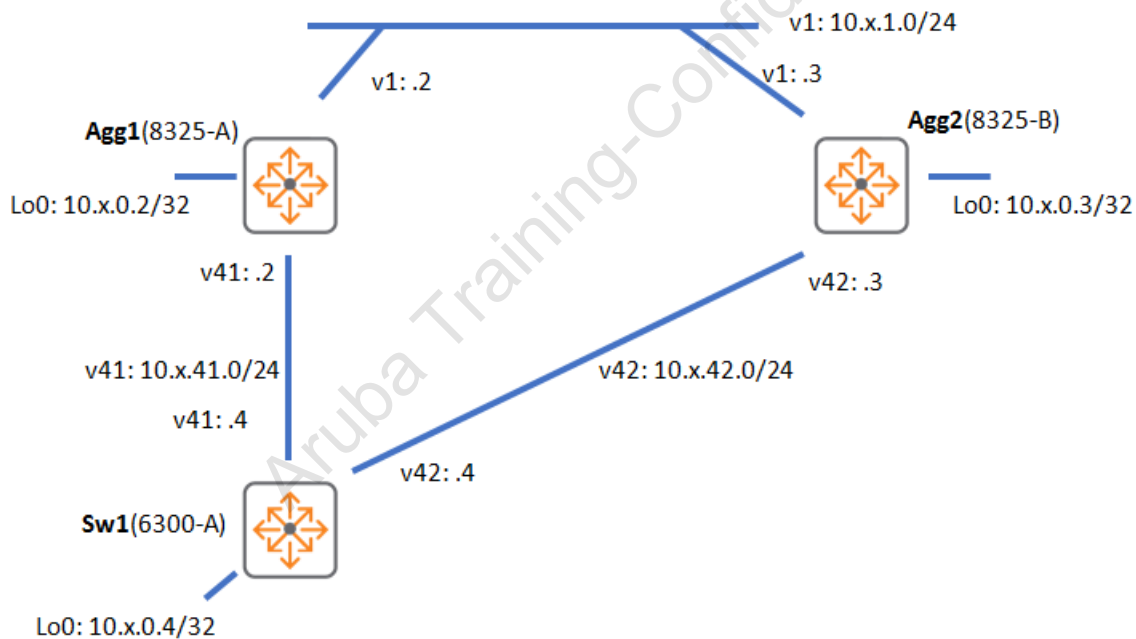
A: There would be 3 LSA entries based on the Agg1 and Agg2 configuration. You should see 2 Router LSAs and 1 Network LSA.

Task 2: Debugging OSPF

Objectives

- Adjacency debugging
- LSDB debugging

Lab Diagram - Layer3 Topology



Steps

Adjacency Debugging

In this section you will establish the Layer3 connections between the VSX and Sw1. During the process you will introduce some configuration errors, so you will need to use the debug and network traces to discover the configuration error.

Prepare Sw1 to VSX links

Agg1

1. On Agg1, define VLAN 41, make 1/1/1 VLAN trunk and allow VLAN 41.
2. On Agg1, assign SVI 41 IP 10.x.41.2/24, enable it for OSPF area 0.

Agg2

3. On Agg2, define VLAN 42, make 1/1/1 VLAN trunk and allow VLAN 42.

NOTE: Remember to convert port 1/1/1 to a switched port. By default ports on the 83xx platform are routed ports.

4. On Agg2, assign SVI 42 IP 10.x.42.3/24, enable it for OSPF area 0.

Sw1

5. On Sw1, define VLAN 41, make 1/1/25 VLAN trunk and allow VLAN 41.
6. On Sw1, assign SVI 41 IP 10.x.41.4/24.
7. On Sw1, define VLAN 42, make 1/1/26 VLAN trunk and allow VLAN 42.
8. On Sw1, assign SVI 42 IP 10.x.42.4/24.
9. On Sw1, define loopback 0 with IP 10.x.0.4.
10. On Sw1, configure OSPF with router-id 10.x.0.4 and define area 0.
11. On Sw1, enable interfaces SVI41, SVI42 and the Loopback0 for OSPF.
12. On Sw1, enable interfaces 1/1/25 and 1/1/26.
13. On Sw1, disable SVI 1.
14. On Sw1, verify the OSPF adjacencies to Agg1 and Agg2.

Q: Do you have adjacencies with both VSX nodes?

A: No, only 1 adjacency came up.

Link Troubleshooting

15. Attempt to troubleshoot this issue. Start from the physical layer and then move up to Layer2 and Layer3.

Are the interfaces 'up'?

Are there any LLDP neighbors?

Are the ports in the VLAN correctly assigned?

Q: What do you observe?

A: VLAN 42 is down, there are no forwarding ports in the VLAN.

Q: You know Layer1 (physical layer) is working, and frames can be exchanged (LLDP). What could be the reason for a port being blocked?

A: Review the spanning tree state of the ports 1/1/25 and 1/1/26.

Q: How can you prevent STP from blocking ports that will only perform routing functions?

A: Configure STP BPDU-filter on these ports.

16. On Sw1, review the active VLANs on the ports. Since no VLANs are active on both ports, there is no Layer2 loop.

17. On Sw1, configure BPDU filter on ports 1/1/25 and 1/1/26.

18. On Sw1, verify you have now OSPF adjacencies with both Agg1 and Agg2.

Review the normal OSPF adjacency setup

In this section you will review the normal setup of an OSPF adjacency. In the next sections you will explore some issues, so you can consider this to be the 'known good' output.

19. On Sw1, enable OSPFv2 debugging.

20. On Sw1, enable terminal-monitor for DEBUG severity.

21. On Sw1, disable the port 1/1/25. Wait a few moments, then review the debug messages.

22. On Sw1, enable the port 1/1/25. Wait a few moments, then review the debug messages.

Q: What are the OSPF states during the adjacency setup?

A: Down > Init > Two-way > Exstart > Exchange > Full

23. Disable the terminal monitor.

Invalid Timers

24. On Agg1, configure VLAN 41 IP interface with an OSPF hello timer of 5 seconds.

25. On Agg1, configure an OSPF dead timer of 15 seconds.

26. On Agg1, verify that the OSPF adjacency with Sw1 was lost on SVI 41.

OSPFv2 Debugging

27. On Sw1, clear the debug buffer, next review the debug buffer in reverse order.

Q: What do you observe?

A: You should see an OSPF ERR message about the neighbor hello interval.

28. On Sw1, adjust the SVI 41 interface OSPF hello interval to the peer interval.

Q: Did the adjacency come online now?

A: No, so something else must be wrong.

29. On Sw1, SVI 41 interface OSPF dead interval to the peer interval.

Q: Did the adjacency come online now?

A: Yes, with the same timers, the adjacency will come online.

30. On Sw1, review the OSPF interface statistics of interface SVI41

Q: What do you observe?

A: The statistics show counters for the various OSPF packets, including the dropped hello packets due to Hello and Dead mismatch.

Invalid authentication

31. On Agg1, configure a plain-text password authentication (key aruba123) on the SVI 41.

32. On Sw1, review the OSPF interface statistics

33. On Sw1, clear and review the debug buffer in reverse order.

Q: What do you observe?

A: The switch reports that there is an unexpected authentication type (1 > id for cleartext/simple authentication) , while it expects no authentication (0 > id for no authentication).

34. On Sw1, run a TCP dump of protocol 89 (the OSPF IP protocol number) for 4 packets.

NOTE: It may take a few moments before the TCPDUMP output is processed, be patient.

Q: Did you see the received OSPF packet? What information did TCPDUMP provide?

A: TCPDUMP shows the OSPF timers and authentication information. For a clear-text password, the password is shown as well.

35. On Agg1, reset the SVI 41 authentication to none.

Router-id Conflict on Peer

36. On Agg1, change the router-id to 10.x.0.4, that is the same router-id as Sw1.
37. Confirm to reset the OSPF process.
38. On Sw1, review the OSPF adjacencies.
39. Review the debug buffer. You should see an OSPF ERR message about the neighbor router ID.
40. Review the OSPF interface statistics, notice the 'Duplicate router ID' counter.
41. On Agg1, correct the router-id again to 10.x.0.2.

OSPF Point to Point networks

When only 2 OSPF routers are expected on a link, it is possible to skip the DR/BDR election process on an OSPF link by configuring the link as Point to Point (P2P).

42. On Sw1, configure SVI 41 and SVI 42 with the OSPF point to point network type.
43. On Agg1, configure SVI 41 as OSPF point to point.
44. On Agg2, configure SVI 42 as OSPF point to point.

Test the point to point

45. On Agg1, disable port 1/1/1 (connected to Sw1, transporting SVI 41)
46. On Sw1, enable OSPF debugging and terminal-monitor for INFO
47. On Agg1, enable the port 1/1/1.
48. On Sw1, review the debug output.

Q: What do you observe?

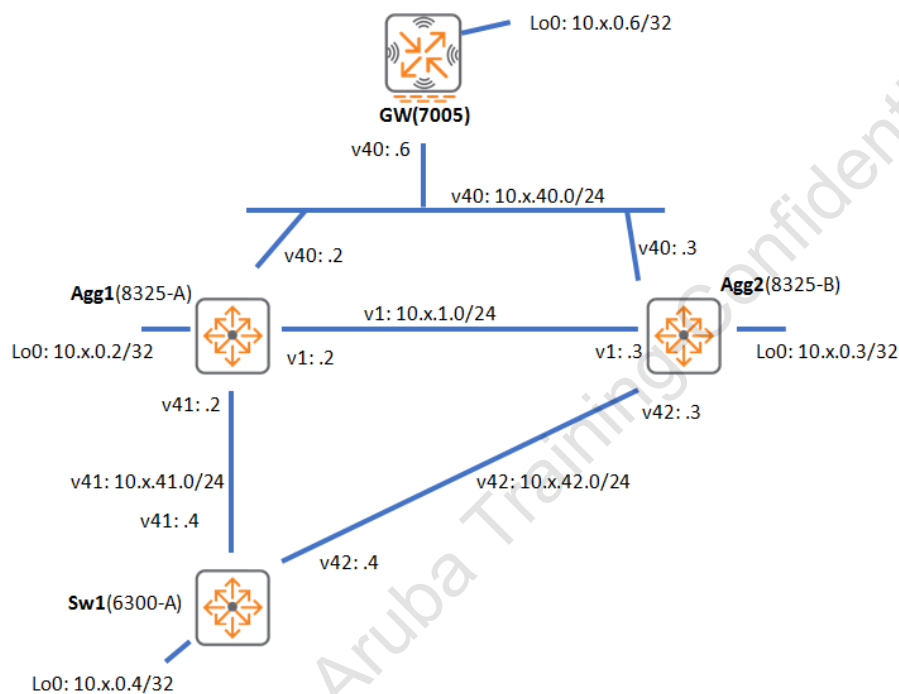
A: Immediately after the Two-way phase, the adjacency exchange is started. This speeds up the initial time required to establish an adjacency with the peer router.

Task 3: VSX Active Forwarding

Objectives

- Understand the need for Active Forwarding
- Configure Active Forwarding

Lab Diagram - Layer3



Steps

An OSPF neighbor of a VSX pair can be connected in various ways:

- Routed port to each VSX node.
- VLAN interfaces (SVI), 1 on each link to each VSX node.
- 1 VLAN interface on a Layer2 LAG to the combined VSX system.

Using routed ports is the most convenient method, but in that case only 1 VRF can be active on the port.

When multiple VRFs are required between the routers, VLAN Interface need to be used to transport multiple SVI contexts over the same physical port.

Since every VRF would require a unique VLAN and subnet for each port, VSX also supports connecting OSPF peers over a Layer2 LAG with multiple VLAN interfaces. This reduces the number of VLAN interfaces that are required.

The administrator should note that instead of point to point connections, the transit network becomes an OSPF broadcast network with 3 routers on the subnet: VSX primary, VSX secondary and the peer router.

Since the peer router will forward traffic for e.g. the VSX primary on the LAG, the LAG hashing may actually deliver the traffic on the VSX secondary.

On the VSX secondary, the traffic would then need to be forwarded to the VSX primary over the ISL.

Since the two VSX systems typically should have the same routing table information, you can configure the VSX systems to route traffic locally in case they would receive traffic that is destined to the VSX peer system.

This is known as 'Active Forwarding'. With 'Active Forwarding', the VSX nodes will learn the SVI MAC address from each other, and when traffic arrives on the SVI with destination MAC address of the peer, it will also be locally routed.

Configuration

In this task, the Aruba Gateway will have the role of the external OSPF router that is connected using a VSX LAG.

You will configure VLAN 40 as a routed link between the Aruba GW and VSX nodes for active forwarding.

1. On Agg1, review the VSX LAG 5. This connects to the Aruba Gateway.
2. On Agg1 and Agg2, define VLAN 40 and an SVI for VLAN 40, enable OSPF in area 0. Add VLAN 40 to interface LAG 5 allowed VLAN list on Agg1.

Device	Value
Agg1 SVI 40	10.x.40.2/24
Agg2 SVI 40	10.x.40.3/24
Aruba Gateway SVI 40 (pre-configured)	10.x.40.6/24

3. On Agg1, review the OSPF neighbors on the SVI 40 interface. Both Agg2 and the Aruba Gateway (router id 10.x.0.6) should be in the list.
4. Use PC1 to open an SSH connection to the Aruba Gateway (10.251.x.6) with credentials admin/aruba123. Review the OSPF neighbors, both Agg1 and Agg2 should be listed.

Active Forwarding - Understand the need for Active Forwarding (optional)

This section shows how some traffic that may be destined to VSX node1 may be sent to the other VSX node due to the LAG hash.

In the next section, Active Forwarding will be enabled to handle this.

Validation:

- Sw1 will be the test device that sends traffic. You will add an additional IP (10.x.0.7) to the Loopback 0 interface in order to test the LAG with a different src/dst IP hash.
 - Sw1 will ping 10.x.40.6 (the Aruba Gateway IP on VLAN40) using source IP 10.x.0.4 and 10.x.0.7 (an even and odd IP address to force a hashing difference). The Aruba Gateway will apply distribute the test traffic over the LAG to the VSX.
 - On Agg1 and Agg2, you will enable a traffic monitor for ICMP traffic on the port to Sw1 and the LAG256 - VSX ISL.
 - This will reveal when a ping is forwarded over the ISL.
5. On Sw1, configure Loopback0 with secondary IP 10.x.0.7/32.
 6. On Sw1, send a single ping to 10.x.40.6 using source IP 10.x.0.4 and 10.x.0.7, both pings should succeed.
 7. On the Aruba Gateway, review the ARP table so you know the MAC of Agg1 and Agg2 on VLAN40.
 8. On Agg1, define a class ip 'icmp' that matches any ICMP traffic.
 9. Define a policy 'mirror' that applies action 'mirror 1' to the class 'icmp'
 10. Apply the policy 'mirror' to LAG5 (Aruba GW) and LAG256 (VSX ISL).
 11. Define mirror 1 with destination CPU, enable the mirror.
 12. On Agg2, repeat the mirror definition and verify the other settings have been synced.
 13. Start 'diag utilities tshark' on **both Agg1 and Agg2**.

Testing

14. On Sw1, send a single ping to 10.x.40.6 using source IP 10.x.41.4.
15. Check the Agg1 and Agg2 traces. If the packet was only visible on 1 switch, it means that that switch performed the routing, so there was no sub-optimal path (and you should run the other ping). In case you saw the ICMP packet on both switches, it means the traffic is redirected to the other VSX node, so you don't need to run the next ping test.

16. This step only applies if the previous ICMP test only showed on 1 switch: On Sw1, send a single ping to 10.x.40.6 using source IP 10.x.42.4.
17. Take note of the source/destination IP combination ping that showed up on both switches.
18. In the TSHARK, you should also check the destination MAC of the ICMP packet. This will not match 1 of the 2 switches, so that switch will forward the traffic over the ISL. The VSX Active Forwarding will resolve this suboptimal forwarding.

Active Forwarding

19. **On both Agg1 and Agg2**, review the current VSX active-forwarding state, it is disabled by default.
20. On Agg1, disable IP ICMP redirects and enable 'vsx active-forwarding' on the VLAN 40 SVI.
21. Review the VSX Active-forwarding state on Agg1.

Q: Did Agg1 learn the SVI 40 MAC address of Agg2?

A: No information from Agg2 is available currently.

22. Review the VSX Active-forwarding state on Agg2.

Q: Did this configuration get synced automatically?

A: No, the Active Forwarding must be enabled locally on the Agg2.

23. Next enable it on Agg2 on SVI VLAN 40.

24. Review the VSX Active-forwarding state on **both Agg1 and Agg2**.

Q: Did the switches learn the MAC and IP of their neighbor VSX system?

A: Yes, the MAC of the peer was learned and is now programmed in the ASIC for local routing.

Validation

25. In the initial test section, you should have found a ping combination that resulted in the ICMP packet being visible on both aggregation switches.
26. On Agg1 and Agg2, enable the 'diag util tshark' again to see ICMP packets.
27. On Sw1, repeat the ping that was previously visible on both aggregation switches.

Q: Was the ICMP packet received on both switches?

A: No, thanks to the Active Forwarding, the traffic is now routed locally by each Aggregation switch for this SVI.

28. In the Tshark output, verify the destination MAC address of the ICMP packet. This should still be the other switch MAC address, but now it is handled by the local switch.

Cleanup steps

29. On both **Agg1** and **Agg2**, remove the policy 'mirror' from the LAG5 and LAG256 interfaces.
30. On Sw1, remove the secondary IP 10.x.0.7 from the Loopback0 interface.

Aruba Training-Confidential

Optional Task 4: OSPF Key-chain and Max-metric On Startup

Objectives

- Configure OSPF authentication using key chains.
- Configure and understand the max-metric on startup for OSPF.

Steps

OSPF interface authentication can be configured using a clear-text or MD5 hash key.

While this works, these methods are very difficult when an environment wants to change the keys while the network is in production, since it is very difficult to change the keys on both sides at the exact same moment.

Key chains can be used to handle this scenario.

Both sides can be configured with a series (chain) of keys. As long as 1 key matches, the authentication is considered successful.

For a migration, a customer just needs to add a key on both sides. Once this is done, the old key can be removed on both sides. Keep in mind that the key id has to match on both sides.

You will now configure a key-chain on VLAN 1 between Agg1 and Agg2.

Configuration

1. On Agg1, define a key-chain named 'asts' with key id 1 using key 'aruba123'.

NOTE: Do **not** add the 'send-lifetime' command, the key will use the default lifetime.

2. Enter the SVI 41 context on Agg1, link the 'asts' keychain to VLAN 41 SVI using the command 'ip ospf keychain asts'
3. Explore the 'ip ospf authentication' command options, however, do **not** enable keychain authentication yet.

Q: What are the authentication options?

A:	null	default> no authentication
	simple-text	clear-text password
	message-digest	single MD5 key with a single key-id
	keychain	multiple MD5 keys, each with their own key-id and optional lifetime

4. On Sw1, create the same keychain and bind it to SVI 41. Do not enable the authentication yet.

Activation

5. **On both Agg1 and Sw1**, under the SVI41 context, apply the 'ip ospf authentication keychain' command at the same time (within a few seconds).
6. On Agg1, verify that the adjacency is formed correctly with the new key without resetting the adjacency.

Migration - Attempt 1

Now attempt to perform an online migration to a new key in the keychain.

7. On Sw1, add a key with id 2 to the key-chain using key 'arubaRocks'
8. On Agg1, wait about 15 seconds, check the OSPF neighbors.

Q: What do you observe?

A: The adjacency with Sw1 was lost.

9. On Agg1, review the ip ospf statistics on interface SVI 41.

Q: What do you observe?

A: The authentication failures counter is increasing.

10. On Agg1, enable ospf all debugging.
11. On Agg1, enable terminal-monitor with severity DEBUG and filter on 'hpe-routing'

Q: What do you notice on Agg1?

A: The debug log shows that the packet cannot be authenticated.

12. On Sw1, review the keychain 'asts' status.

Q: What is the active Send key?

A: Sw1 is using the new key 2 as the active send key, but this key does not exist yet on Agg1.

13. Disable the OSPF debug

TCPDUMP on Agg1

14. On Agg1, enable diagnostics mode
15. Use TCPDUMP of the OSPF protocol (protocol 89).
16. Stop the TCPDUMP after you see an OSPF packet with authentication type MD5. Look for packets from source Sw1 (10.x.41.4).

NOTE: It may take a few moments before the TCPDUMP shows output

packets.

Q: What is the key-id on the inbound OSPF packet?

A: Key-id 2

Q: What is your conclusion?

A: The previous steps revealed that the new key is used immediately as the outbound key.

Migration - Think about the migration steps

17. On Sw1, review the keychain.

Q: How could you use the Send and Receive Key validity to handle the key migration?

A: You can add a key that is a valid 'Receive' key, but that cannot be used for sending yet. Once both sides have the key in the valid Receive state, the send validity can be adjusted.

Migration - Attempt 2

18. Rollback: On Sw1, first remove the key 2 from the keychain.

19. On Sw1, verify the OSPF adjacency with Agg1 is established again.

20. On Agg1, add key id 2, but do not add the key-string yet.

21. Configure the send-lifetime with some start date in the future and duration infinite.

22. Configure the key-string 'arubaRocks'

23. Review the keychain status.

Q: What is the active Send key id?

A: Key 1 is still the active Send key.

Q: What are the active Receive key ids?

A: Both key 1 and 2.

24. On Sw1, verify the OSPF adjacency with Agg1 is still established.

25. On Sw1, add the key2 'arubaRocks' to the keychain.

26. On Sw1, review the keychain 'asts'.

Q: What is the active send key?

A: The active send key is now key id 2

NOTE: You are now half-way in the migration. Sw1 is using key 2 for transmit, and this is accepted by Agg1. Agg1 is still using key 1 for transmit, this key is still accepted by Sw1.

27. On Agg1, change the key 2 send-lifetime so it can be used now.

28. On Agg1, review the keychain status.

Q: What is the current Send Key ID?

A: Due to the start-time change, the current send Key ID is 2.

29. Complete the migration: **On both Agg1 and Sw1**, remove the key id 1 from the keychain and verify that the adjacency is still active.

Max-metric on Startup

When an OSPF router comes online, it may need some time to establish links. During this period, it may not be ideal to use this router as an OSPF transit router, since it may not know all routing paths yet.

With the max metric on startup, an administrator controls OSPF to announce the maximum metric for its links, so it will be the least preferred path for a period of time.

After the max age on startup timer expires, the normal cost will be applied and normal routing can resume. This allows OSPF to fully converge before any transit traffic must be processed.

30. On Sw1, check the OSPF routing table, check the 10.x.0.6 network (loopback of the Aruba Gateway), this is network where Agg2 is a 'transit' router.

Q: What is the cost for the 10.x.0.6/32 network?

A: Currently, the cost is 201.

Q: How many paths are there for this destination?

A: There are 2 paths: via 10.x.41.2 and 10.x.42.3.

31. On Agg2, disable OSPF.

32. Configure the max-metric option to 60 seconds on Agg2. The default timeout is 600 seconds.

33. On Sw1, run the commands 'show ip ospf routes' and 'show ip ospf neighbors'

34. On Sw1, repeat the show ip route 10.x.0.6

35. On Agg2, enable OSPF again.

36. On Agg2, use the 'show ip ospf routes' command to see the local version of the OSPF routing table.

Q: What is the cost for the OSPF routes?

A: The routes are listed with a cost of 65535, this is the maximum metric for the link.

37. On Sw1, the repeat command should still be running.

Q: How many paths are available to reach the 10.x.0.6/32 network?

A: Only the path via Agg1 is currently shown, since the path via Agg2 has a much higher cost.

38. After about 1 minute, there should be 2 paths to the 10. Tx.0.6/32 network.

39. On Agg2, review the 'show ip ospf route' command.

Q: What are the costs of the routes?

A: Once the timer expires, the normal metric is applied to the interfaces again.

40. On Sw1, stop the repeat.

Task 5: Policy Based Routing (PBR)

In this task you will configure Policy Based Routing between a simulated Branch location and the main network.

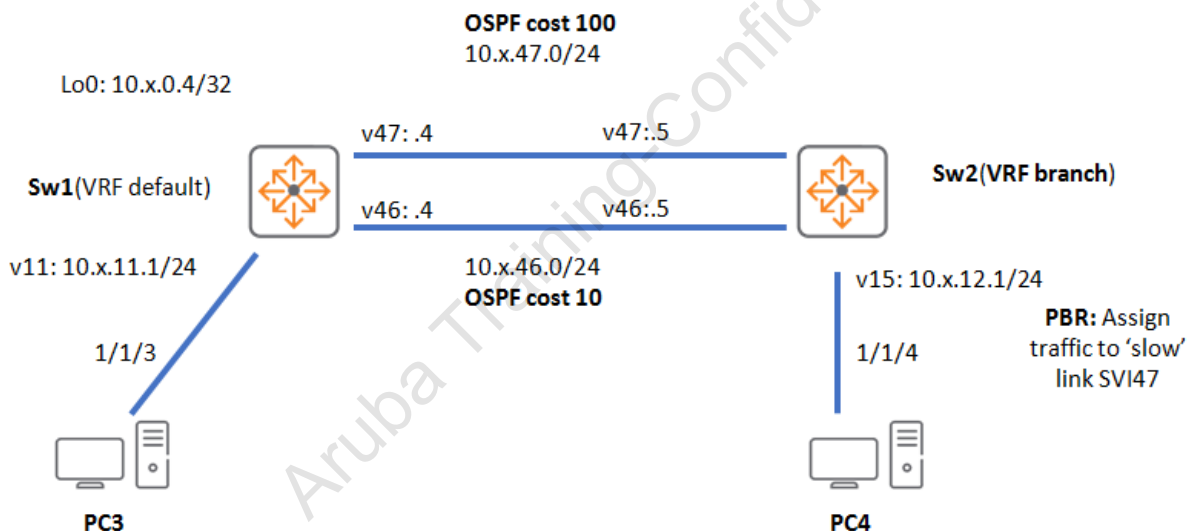
The Branch location is simulated by using a VRF on the Sw2, it will connect using 2 SVI interfaces to the Sw1 (acting as the main network router).

The 2 links between Sw2 and Sw1 are used as 'WAN' links, with different OSPF costs applied to them.

Objectives

- Configure and monitor PBR

Lab Diagram



Steps

Prepare the setup

You will need additional interfaces between Sw1 and Sw2. You can perform this configuration manually (Option1) or use the template files to copy/paste the configuration (Option2).

Option 1: Manual Configuration

1. On Sw1, define VLAN 46 and VLAN47, enable them on the port 1/1/27.

- On Sw1, define SVI 46 and SVI 47, assign IP addresses to the new interfaces, enable L3 counters, OSPF and make them OSPF point to point connections. Assign SVI 46 OSPF cost 10, SVI 47 cost 100.

Interface	IP
SVI 46	10.x.46.4/24
SVI 47	10.x.47.4/24

- On Sw2, define VLAN 46 and VLAN47, enable them on the port 1/1/27.
- On Sw2, define VRF 'branch', enable OSPF process 1 for VRF 'branch' and define area 0.
- On Sw2, define SVI 46 and 47, attach them to VRF 'branch'.
- On Sw2, assign IP addresses to the new interfaces.

Interface	IP
SVI 46	10.x.46.5/24
SVI 47	10.x.47.5/24

- On Sw2, verify that both interfaces SVI 46 and SVI 47 have active OSPF adjacencies.
- On Sw2, you will prepare VLAN 15 as the 'branch-PC' VLAN. PC4 will be moved to this VLAN. Sw2 will act as the DHCP server in the branch using this range: 10.x.15.200 – 10.x.15.210 and 10.x.15.1 as a default-router

- On PC4, renew the IP address, verify it has received an address in the 10.x.15.0/24 range. This concludes Option 1.

Option2: Saved Script configuration

- On PC1, open Desktop > ASTS > Configs, open ASTS_cfg_lab05-steps-v2.txt. Scroll down to the Task6 section and apply the sw1 and sw2 configuration.

- On PC4, renew the IP address, verify it has received an address in the 10.x.15.0/24 range.

This concludes Option 2.

OSPF best path verification

- On Sw2, review the current configuration of SVI 46 and SVI 47. Both SVI interfaces are assigned to the VRF 'branch', OSPF is enabled and a cost has been set:

SVI 46 OSPF cost 10

SVI 47 OSPF cost 100

13. On Sw2, review the OSPF and IP routing table for the VRF branch.

Q: What is the preferred path to 10.x.12.0/24?

A: Based on the configured OSPF interface cost, SVI 46 (next hop 10.x.46.4) will be preferred

Test assigned link and failover without PBR applied

You will now use Sw1 to review the statistics of the inbound traffic in the SVI 46 and SVI 47 interfaces. Since currently, there is no PBR applied, it is expected to receive all traffic on SVI 46 (lowest cost). Only in case this link would fail, the traffic would be sent over SVI 47 (higher cost).

14. On Sw1, clear the interface statistics of the current session.

15. On Sw1, show the SVI 46 and SVI 47 statistics, next repeat the last 2 commands every 10 seconds. This will show you the 'latest' 10 seconds statistics.

16. On PC4, ping 10.x.12.1, review the repeated statistics on Sw1. The traffic should have used SVI 46.

Test link failover

17. On Sw2, enable terminal-monitor for severity INFO and disable SVI 46.

18. On Sw2, after a few seconds, the terminal monitor will log a message AdjChg: Full> Down.

19. On PC4, attempt to ping 10.x.12.1.

Q: What do you observe?

A: Although Sw2 knows that SVI 46 is down, and although OSPF should have recalculated the path, the ping still fails.

Q: What could the reason be?

A: For routing, it is important to consider both directions of the path. While the path from Sw2 to Sw1 may be aware of the change, the path from Sw1 to Sw2 is not aware of the change yet. Since there is no physical port change (SVI 46 was shutdown, not port 1/1/27), Sw1 needs to rely on the OSPF hello packets and dead timers.

Q: What is a typical protocol to help for these scenario's?

A: BFD (Bidirectional Forward Detection) is typically used when OSPF peers are not directly connected.

Enhance the failover detection for the SVI using BFD

20. On Sw2, first enable the SVI 46 again.

21. On both Sw1 and Sw2, enabled BFD globally and on the SVI 46 and SVI 47 interfaces.

22. On Sw1, verify the state of the BFD sessions

23. On Sw1, review the details of BFD session ID 1

Q: What is the default Min Tx, Min Rx Interval and the default detect multiplier?

A: The default Min Tx and Min Rx Interval are 3000ms and the default detect multiplier is 5.

Q: How long does it take using the default values to detect that the peer device is unreachable?

A: Using the default timers, it takes 15 seconds.

24. **On both Sw1 and Sw2**, adjust the global BFD timers to transmit and accept receive every 1000ms and set the detect multiplier to 3.

25. On Sw1, start the statistics repeat again

NOTE: Due to the BFD traffic, additional packets (about 22-24 packet per interval) will be observed in the statistics.

Test link failover

26. On PC4, start a continuous ping to 10.x.12.1.

27. On Sw2, disable SVI 46.

28. On PC4, review the ping output.

Q: What do you observe?

A: The ping continues after 1 lost ping due to the enhanced BFD failure detection.

29. On Sw1, review the statistics. The packets will now use SVI 47.

30. On Sw2 enable SVI 46 again.

You have now verified the default traffic flow and the failover. In the next section PBR will be configured.

Policy Based Routing Configuration

In this section you should configure PBR so that traffic to a voice subnet (simulated by the 10.x.12.0/24 subnet) from the branch (SVI 15) will be sent over the SVI 47 path. Any other traffic will still use the SVI 46 path. You will use a ping to 10.x.11.1 to test the 'other' traffic.

31. On Sw2, define a new IP class named 'voice' for any traffic matching destination IP 10.x.12.0/24. Include the 'count' option.

32. Define a PBR action list named 'nexthop47' with the SVI 47 link next hop ip (10.x.47.4)

33. Define a new policy 'branch' that applies the PBR action 'nexthop47' to the class 'voice'.
34. Apply the new policy to the VLAN 15 interface.

Verification

35. On Sw2, review the PBR configuration
36. On Sw2, review the policy hitcounts

37. On Sw1, the SVI 46 and SVI 47 statistics are still repeated, review the statistics.

Q: What do you observe?

A: For the SVI 47, the Rx statistics have increased to about 33-35 packets.
For the SVI 46, the Tx statistics have increased to about 10-12 packets.

Q: What could be the reason for the observed asynchronous routing?

A: PBR has only been configured on the branch. The return traffic is still using the normal routing table, so the traffic destined to the branch will use SVI 46. This lab activity only demonstrates the PBR in 1 direction, in a real deployment, it will typically be configured on both ends of the link.

Verify the 'other' traffic

38. Verify the 'other' traffic behavior. On PC4, stop the ping to 10.x.12.1 and start a new ping to 10.x.11.1.
39. On Sw1, review the statistics output. This traffic should only use SVI 46.

Cleanup

40. On Sw2, assign the PC4 port 1/1/4 to VLAN 12.

This concludes the PBR activity.

Required: Post-lab checkpoint

41. On Agg1, Agg2, Sw1 and Sw2, make a new checkpoint named 'asts-lab05-<your-name>'.

IMPORTANT: Previous checkpoints at the end of a lab activity were marked as Optional, this checkpoint is **required** since you will need it at the start of Lab7!

At the start of Lab 07, you will need to revert to this checkpoint!

You have completed this Lab!

Aruba Training-Confidential

Lab 06: BGP

In this lab you will configure iBGP and eBGP peering and apply route control for the received and advertised routes.

Requirements

This lab requires completion of Lab 05.

The routerA, routerB and routerC are configured as eBGP peers.
Agg1, Agg2 and Sw1 belong to the same AS 64500. Sw2 will be configured in a separate AS, so its routes will need to transit through the AS 64500.

- iBGP peering using Loopback interfaces
- use BGP peer groups
- configure a BGP route reflector
- Perform and understand BGP route control and route manipulation

Scenario

Load checkpoint asts-bgp on routerA, routerB and routerC. Then, save and reboot the 3 routers.

In this scenario there is a customer (AS64500), three ISPs (AS64511, AS64512, AS64513) and a partner (AS64600) that accesses the Internet through the customer's network. The lab diagram below provides IP address details.

iBGP

The customer plans to setup BGP peering with some external organizations. They want to run a Proof of Concept on several of the BGP features. To do this:

1. Configure iBGP peering between Agg1 and Agg2 for BGP AS64500 using the Loopback interfaces.
 - a. Configure Agg1 without the update-source IP address and use BGP debugging to discover the selected source IP address.
 - b. Once you have seen the BGP debugging, correct the configuration.
2. Configure Sw1 as an iBGP peer in AS64500. Do not use a route-reflector.

eBGP

The customer wants to see how to debug an eBGP peer.

1. On Agg1, configure routerA as an eBGP peer using invalid AS number 64520.
 - a. Use BGP debugging to detect the error.
 - b. Once you have seen the BGP debugging, correct the configuration.
2. Configure Agg2 as an eBGP peer connected to routerB.
 - a. Review why Agg1 does not have a route to 198.51.100.0/24 via Agg2
 - b. Correct the configuration for all of the other routers in the customer AS
3. On Sw2, load the configuration from asts_cfg_lab06-task2-Sw2.txt. This will create a 'partner' BGP peer in a VRF 'bgpext'.
4. On Sw1, define SVI48 on port 1/1/27 and establish eBGP peering with Sw2-BGPext (AS64600).
5. The customer wants Agg1 and Agg2 to advertise a default route into OSPF, even when these routers do not have a default route in their own routing table.

BGP Advertisements

To demonstrate BGP Advertisements:

1. On Agg1 add a secondary address 100.100.1.2/26 to SVI12.
2. On Agg2, add a secondary address 100.100.1.3/26 to SVI12.
3. On Agg1, advertise the 100.100.1.0/24 network using a network command.
4. On Agg2, redistribute the directly connected routes matching 100.100.1.0/26 into BGP.
 - a. Make sure it appears (origin) as if a network command was used.
 - b. Use an aggregate address to advertise the 100.100.1.0/24 network and make sure that there are no subnets of the aggregate address advertised.

BGP Peer Groups

On Agg1, reconfigure the iBGP peers Agg2 and Sw1 using a BGP Peer group. Migrate the existing settings into the peer group.

Controlling Transit Traffic

The customer is concerned that it could advertise not only its own networks, but also the networks it has received from other providers, making it a transit provider. Demonstrate to the customer how to stop this:

1. Ensure that Agg1 and Agg2 advertises only their own networks to the routerA and routerB systems.
2. Once this has been implemented, the partner (AS64600) complains that their networks are not reachable anymore.
 - a. Configure Agg1 and Agg2 so that they also advertise the partner (AS64600) networks on top of their own networks.
 - b. Ensure no other networks are advertised.

Outbound Traffic Route Control

Now demonstrate to the customer how to control the outbound traffic to some destinations that are hosted on routerC. They want you to demonstrate how the traffic path can be controlled using BGP.

1. The prefixes that need to be controlled are:
 - a. 198.19.100.0/24
 - b. 198.19.101.0/24
 - c. 198.19.102.0/24.
2. Note that the routerC has done its own path manipulation for these routes, so the path to:
 - a. 198.19.100.0/24 has an equal AS path length over routerA and routerB
 - b. 198.19.101.0/24 is by default preferred via routerA
 - c. 198.19.102.0/24 is by default preferred via routerB
3. Demonstrate that the traffic destined to any ISP (routerA/B/C) leaves the AS via Agg1. Such that:
 - a. Any external route learned by Agg1 is more preferred than a route learned by Agg2.
 - i. Configure this only on Agg1 and ensure the routing tables of Agg2 and Sw1 prefer this path as well.
4. Demonstrate how to apply selective route control, so traffic destined to:
 - a. 198.19.100.0/24 should have equal preference over Agg1 and Agg2
 - b. 198.19.101.0/24 should be sent via Agg1
 - c. 198.19.102.0/24 should be sent via Agg2
 - i. Any other routes should still take the default preferred path via Agg1.
 - ii. This may be configured on both Agg1, Agg2 or just on one side.
 - iii. If either Agg1 or Agg2 fails or their uplink fails, traffic should be able to use the other Aggregation switch.

5. After this configuration, Sw1 still has only one route in the active routing table for the 198.19.100.0/24 network, even when the AS Path and Local Preference are the same.
 - a. Demonstrate how to ensure that Sw1 can use ECMP for the route to 198.19.100.0/24
 - b. BGP routes to Agg1 and Agg2 should both be active in the routing table.
6. Your demo is going well, next show how to make a small exception of the route control on Agg2.
 - a. When Agg2 needs to send traffic to 198.19.101.0/24, it should go out via the link to routerB. The rest of the AS (Agg1 and Sw1) should still use the preferred path via Agg1.
 - b. This configuration change must be done only on Agg2.
7. Demonstrate how to avoid tearing down the BGP session when a new inbound policy must be processed. Ensure that the Agg1 and Agg2 can process inbound route changes:
 - a. Without resetting the BGP session
 - b. Without requesting a full route refresh from their peers.

Inbound Traffic route control

The customer wants to make sure that the outbound advertisement for the 100.100.1.0/24 route via Agg2 to routerB has a longer AS Path length than the Agg1 advertisement. They want to repeat their AS number 3 additional times in the AS Path attribute. The result should be that routerA, routerB and routerC prefer the path via Agg1 to the 100.100.1.0/24 network due to the shorter AS Path.

The customer wants to control the inbound traffic from the partner network that is running on Sw2 VRF 'bgpext'.

1. Make an IP link between Sw2 and Agg2. SVI12 will be used for that link.
2. Use ASTS_cfg_lab06-task6-Sw2.txt for the Sw2 configuration.
3. Configure Agg2 to peer with Sw2 on the SVI12 interface.

Assuming the Sw2 partner does not perform any BGP route control, the customer wants traffic from the partner Sw2 to 100.100.1.0/24 to enter the AS via Agg2. When that link is not available, the link to Sw1 should be used. This must be configured on Agg2 and/or Sw1, but it should not be configured on the Sw2 partner.

The partner network 'bgpext' on Sw2 is currently advertising these networks to the customer:

- 198.18.14.0/24
- 198.18.15.0/24

The partner wants to control when the 198.18.15.0/24 network is advertised to the ISPs on routerA and routerB. They want to ensure that the 198.18.15.0/24 is no longer advertised to these ISPs. If they do want to advertise it to the ISPs, they should be able to change it themselves, without changing the customer BGP configuration.

AS Number filtering

The customer wants to ensure that the ISPs do not see the partner 'bgpext' AS number in its outbound advertisements to the ISPs. Make sure the AS number of the partner is no longer present in the advertisements to routerA and routerB.

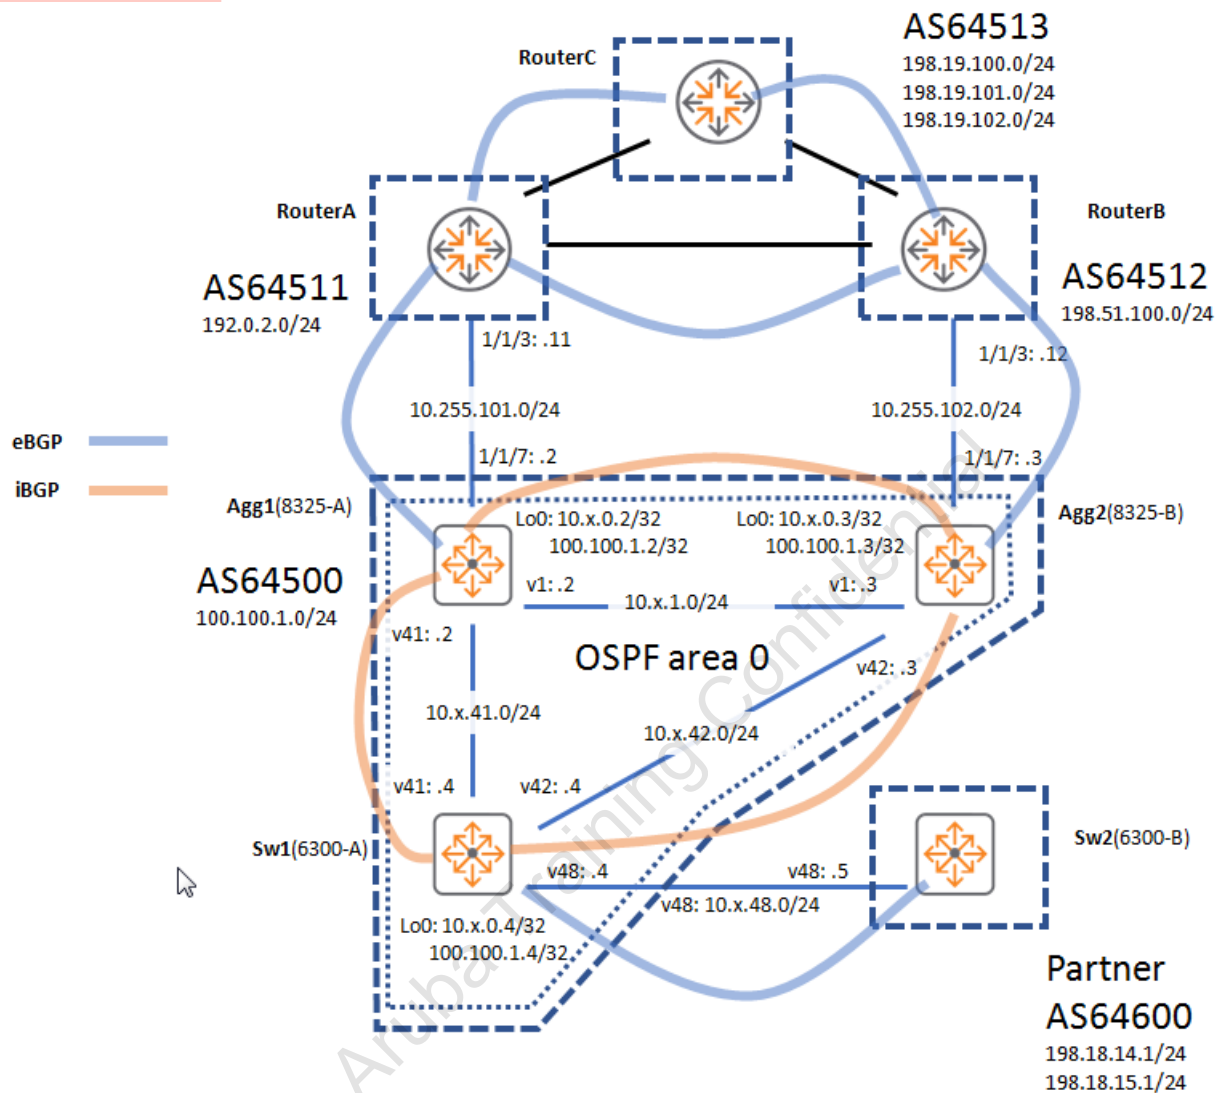
iBGP peering

The customer is worried that the iBGP full mesh requirement will be a challenge to maintain when more iBGP routers are added. Demonstrate an alternative solution to the full mesh.

1. Remove the iBGP peering between Sw1 and Agg2.

Configure BGP so that the Sw1 can still receive the Agg2 routes and vice-versa.

Lab Diagram



Task 1: Prepare the setup

Objectives

- Prepare the router devices for the BGP lab activity

Steps

Router Configurations

1. On routerA, load the checkpoint '**asts-bgp**' to the running-configuration, save and reboot
2. On routerB, load the checkpoint '**asts-bgp**' to the running-configuration, save and reboot
3. On routerC, load the checkpoint '**asts-bgp**' to the running-configuration, save and reboot

Task 2: Configure iBGP and eBGP peering

Objectives

- Configure iBGP peering using the loopback interface
- Configure eBGP peering
- Advertise default route into OSPF
- Advertise internal route to eBGP peer
- Configure BFD for eBGP peers

Steps

In case you feel familiar with the basic iBGP and eBGP peering, you can skip this task and use the commands from the lab06-steps.txt that you can find in Desktop> ASTS > configs

Troubleshooting iBGP peering

Incorrect iBGP session IP address

1. On Agg1, shutdown the uplink port 1/1/7 to routerA.
2. On Agg1, configure BGP AS 64500. Configure iBGP peering to Agg2 loopback IP address, do not configure a source interface yet. Enable the peer for IPv4 address-family.
3. On Agg2, configure BGP AS 64500. Add Agg1 Loopback IP as iBGP peer, enable the IPv4 address family.
4. On Agg1, review the current BGP peers.

Q: What do you observe?

A: There are currently no peers active.

5. On Agg1, debug bgp to see what is happening.
6. On Agg1, review the debug buffer. This may take a few moments based on the default BGP timers.

Q: What do you observe?

A: Either Agg1 or Agg2 will be the first to attempt to send an OPEN message to the peer, so you need to parse the log carefully. You should see that an 'OPEN' message is sent to a neighbor.

Q: What is the 'selected local address' for this OPEN message?

A: The Agg1 has selected an outbound interface IP, that could be 10.x.1.2 for example.

Q: What is the IP address that Agg2 expects?

A: Agg2 expects to receive a BGP message from 10.x.0.2.

7. On Agg1, configure the correct source interface.
8. On Agg2, configure the correct source interface.
9. On Agg1, review the BGP IPv4 unicast summary. The peer should now be established.
10. On Agg1, review the debug log for the 'bgp' module.

Q: What is the configured local address for the OPEN message?

A: Now the loopback 0 will be reported as the configured local address.

Configure iBGP peering between the other switches.

11. On Agg1, configure iBGP peering to Sw1 using the loopback interface.
12. On Agg2, configure iBGP peering to Sw1 using the loopback interface.
13. On Sw1, configure iBGP peering to Agg1 and Agg2 using the loopback interface.
14. **On Agg1, Agg2 and Sw1:** verify the BGP peering is established between with the other 2 switches.

Configure eBGP peering with routerA and routerB

Troubleshoot remote AS number

15. On Agg1, configure eBGP peering with routerA using the physical interface address, configure an invalid AS number (64511 is correct, you will configure 64520).
16. On Agg1, review the BGP peer status
17. Clear the debug buffer and enable port 1/1/7.
18. Review the debug buffer.

NOTE: You may need to wait up to 60 seconds to receive an incoming BGP message from the router.

Q: You should observe an incoming BGP OPEN message. What is the Remote AS number of this message?

A: The AS number is 64511. This does not match the local configuration, so the peering cannot be established.

19. Correct the Agg1 configuration.

20. Review the peering status, it should now be established
21. On Agg2, configure eBGP peering with routerB using the physical interface address.
22. On Agg2, verify the BGP peering is established with the eBGP peer.

Next hop self

23. On Agg1, review the current IPv4 unicast BGP table. Take a closer look at the 192.0.2.0/24 route.

Q: How many routes exist in the BGP IPv4 unicast routing table for this prefix?

A: There are 2 entries listed.

Q: Which route is the 'best' route?

A: The local route with the 10.255.101.11 next-hop IP is the best route.

24. On Agg1, look for the details of this route.

Q: What is the next-hop IP for each route?

A: 10.255.101.110 and 10.255.102.12

Q: Are these next-hop IP addresses reachable for the Agg1?

A: Let's look in the local routing table.

25. On Agg1, check the next-hop IP in the local routing table

Q: Is this IP address reachable?

A: No. When the external subnet is not reachable via the IGP, iBGP should be configured with next-hop self.

26. Review again the BGP details of the 192.0.2.0/24 route on Agg1

Q: Is the route valid?

A: Yes, even though the next-hop IP is not reachable, it is still a valid route in the BGP routing table. The fact that the next-hop IP is not reachable simply excludes the route from the 'best path' selection process.

Solution: configure next-hop self.

27. On Agg2, configure the 'next-hop self' option for the Agg1 peer.

28. On Agg1, review the details of the 192.0.2.0/24 route.

Q: Did the next-hop IP change?

A: No, this change would only be reflected once the session is re-established.

Solution: Reset the BGP peering

To refresh the list of entries from a BGP peer, the BGP peer session can be reset. An example would be 'clear bgp 10.x.0.3'. This will clear the entire BGP session and all address families, this is also known as a 'hard-reset'.

29. On Agg2, enable BGP debug all to see the BGP message.

30. BGP Hard Reset: On Agg1, you can now perform a hard reset of the BGP session with Agg2 (10.x.0.3)

31. On Agg2, review the debug buffer.

Q: What do you observe?

A: With the hard reset, the BGP session is completely cleared and needs to be re-established.

32. On Agg1, review the IPv4 unicast routing table.

Q: Did the next-hop IP change?

A: Yes, thanks to the refresh the latest next-hop IP is shown.

33. Complete the next-hop self configuration on Agg1

34. Complete the next-hop self configuration on Agg2

35. Complete the next-hop self configuration on Sw1

36. **On Agg1 and Agg2**, reset all sessions to make the next-hop self effective.

NOTE: In a lab environment, you can use 'clear bgp *' to reset all BGP sessions, be careful with this command in a production environment.

eBGP between Sw1 and Sw2 (acting as External BGP Partner)

Sw2 will now be configured as a BGP partner network. To isolate this role from the Sw2 acting as 'layer2' switch, this BGP configuration will be assigned to a dedicated VRF. First you will configure a Layer3 link for this communication.

37. On Sw1, add VLAN 48, allow it on the 1/1/27 VLAN trunk.

38. On Sw1, add SVI 48 with IP address 10.x.48.4/24.

Now apply the template configuration to Sw2.

39. On Sw2, apply this configuration. (you can load this from the PC1 > **Desktop > ASTS > configs > asts_cfg_lab06-task2-sw2.txt**

```

config

## VRF 'bgpext' is hosting the eBGP partner peering settings.
vrf bgpext

## Layer3 link to Sw1
vlan 48
  description eBGP-Sw1
interface 1/1/27
  vlan trunk allowed 48
  exit

interface vlan48
  description eBGP-Sw1
  vrf attach bgpext
  ip address 10.x.48.5/24
  exit

## Subnets running at the BGP Partner network
interface loopback 11
  description eBGP-partner-subnet
  vrf attach bgpext
  ip address 198.18.14.1/24
  exit
interface loopback 12
  description eBGP-partner-subnet
  vrf attach bgpext
  ip address 198.18.15.1/24
  exit
ip prefix-list net198-18-15 seq 10 permit 198.18.15.0/24 ge 24 le 24

route-map bgp-import-connected permit seq 10
  match ip address prefix-list net198-18-15
  exit

## BGP peering with Sw1 inside VRF bgpext
router bgp 64600
  vrf bgpext
    bgp router-id 10.x.0.5
    neighbor 10.x.48.4 remote-as 64500
    address-family ipv4 unicast
      neighbor 10.x.48.4 activate
    network 198.18.14.0/24
    redistribute connected route-map bgp-import-connected
  exit
exit

```

Now complete the configuration on Sw1.

40. On Sw1, define the eBGP peer for Sw2 (10.x.48.5) in AS 64600.
41. On Sw1, verify the peer is in 'established' state.
42. On Sw1, verify the BGP IPv4 routing table. You should see entries from the AS 64600.

Advertise default route into OSPF

On the border routers that provide access to the internet, it is possible to inject the default route into OSPF, so that all the internal OSPF systems will use the border routers for the default route.

43. On Sw1, verify that there is no default route in the current state of the lab.
44. **On both Agg1 and Agg2**, ensure that the default route is always advertised into OSPF.
45. On Sw1, verify in the OSPF LSDB that the 2 default routes are in the database and in the IP routing table.

Advertise a route to eBGP peer

Before BGP advertises a route to a BGP peer, BGP will verify if route exists in the local routing table. In this section, you will configure a secondary SVI 12 IP address on Agg1 and Agg2 in the 100.100.1.0/26 range.

Next you will configure BGP on Agg1 and Agg2 to advertise the 100.100.1.0/24 network.

On Agg1, you will use the static route blackhole method.

On Agg2, you will use the BGP summary address advertisement method.

First configure the secondary IP on SVI12. Make sure to set the correct **/26** mask.

46. On Agg1
47. On Agg2

Now let's start with Agg1 advertisement.

48. On Agg1, advertise the 100.100.1.0/**24** network using the network command. Make sure to use the /24 subnet mask.
49. On Agg1, review the BGP IPv4 routing table

Q: Why is the route not shown?

A: BGP requires an exact match in the routing table for the 'network' commands.

50. On Agg1, correct this problem with a blackhole static route and verify the local BGP IPv4 unicast table.

51. Verify the advertisement: On Agg2, review the BGP routing table

52. Open a session to routerA. Verify the BGP routing table.

53. Open a session to routerB. Verify the BGP routing table.

Q: Is routerB directly connected with Agg1?

A: No.

Q: Did Agg2 require any local advertisement configuration for the 100.100.1.0/24 route to be advertised to routerB?

A: No. Agg2 simply advertises the learned routes from its own peers, that includes the prefixes it learned from Agg1, such as the 100.100.1.0/24 prefix.

Q: What would happen if Agg1 would fail?

A: Even when Agg2 has an eBGP connection, the route advertisements would no longer be performed, since Agg2 does not have an advertisement in its own configuration and there are no other routers that are advertising this route.

Now let's configure the advertisement on Agg2 to provide redundancy.

For Agg2, you will use the `bgp aggregate` address method instead of the blackhole route method. The disadvantage of the Agg1 'static route' method is that this advertisement will always be performed, even when the actual route would not be reachable. For example: if the SVI 12 interface would be disabled or go down, the 100.100.1.0/26 subnet would not be reachable, but the advertisement would still be performed.

This is different with an aggregate route, since the aggregate route will only be advertised if a matching prefix exists in the local BGP routing table. When that prefix would disappear, the aggregate route will also be removed.

For a BGP aggregate address to work, a prefix of the aggregate must exist in the BGP table. On Agg2, this will be achieved by redistributing the connected route 100.100.1.0/26 into BGP. To prevent all connected routes to be redistributed, a route-map will be used.

54. On Agg2, define an IP prefix-list 'connected_to_bgp' that matches the 100.100.1.0/26 prefix.

55. On Agg2, define a route-map 'connected_to_bgp' that permits and matches the IP prefix-list 'connected_to_bgp'.

56. On Agg2, redistribute the connected routes using the route-map 'connected_to_bgp' under the BGP IPv4 address-family.

57. On Agg2, review the BGP IPv4 routing table and verify that the /26 prefix has been added to the local BGP table.

Q: What is the Origin code for a redistributed route?

A: Redistributed routes are marked with 'incomplete' origin using the '?' sign.

Q: What is the Origin code of the 100.100.1.0/24 route that is advertised by Agg1?

A: Routes that are advertised using the 'network' command are marked as 'i' for Internal.

Q: What must be done on Agg2 to make the redistributed route appear as an internal route?

A: The redistribution route-map can be used to set the origin value. This will ensure that both advertisements are considered the same.

58. On Agg2, modify the route-map 'connected_to_bgp' to apply origin 'igp'.

59. On Agg2, review the route-map.

60. Since there is no aggregate configured, Agg2 advertises the prefix 100.100.1.0/26 to its peers. Review the advertised routes to routerB

61. On Agg2, now configure the aggregate route under the IPv4 address-family.

62. On Agg2, review the local BGP table and the advertised routes to routerB.

Q: What do you observe?

A: There is now an aggregate /24 prefix, but the /26 prefix is also still advertised to the peers.

63. On Agg2, make sure that only the summary address is advertised, not the subnets.

64. On Agg2, review the local BGP routing table.

Q: What is the status code for the 100.100.1.0/26 prefix?

A: 's' (lowercase). This indicates that the prefix has been suppressed by an aggregate prefix.

65. On Agg1, verify that only the /24 prefix is advertised to the external peers

66. On Agg2, repeat this check

Task 3: Using BGP peer groups

Objectives

- Use BGP peer groups to simplify the BGP peer configuration

Steps

When a BGP router has several BGP peers, many configuration lines may be duplicated for the BGP peers. When several BGP peers need to be configured with the same settings, a BGP peer group can be used to simplify the configuration.

In this task you will configure a BGP peer group on Agg1 and migrate the existing configuration of the Agg2 and Sw1 iBGP peers into the peer group.

Prepare the Peer-group

1. On Agg1, add a new BGP peer-group with name 'internal'.
2. Configure the peer-group with remote-AS 64500.
3. Configure the peer-group with source interface Loopback0.
4. Configure the peer-group with next-hop-self under the IPv4 unicast family.

Migrate the existing peer to the peer-group

5. Remove the existing Agg2 peer (10.x.0.3).
6. Add Agg2 again with the peer-group command.
7. Make sure to activate the neighbor under the IPv4 address family.

IMPORTANT: The activation of the neighbor under the address family is done per neighbor, not at the peer-group level.

8. Repeat these steps for Sw1 peer (10.x.0.4).
9. Verify that the BGP neighbors on Agg1 are back in the established state.
10. Review the current BGP configuration using

Task 4: Controlling transit traffic

Objectives

- Prevent transit AS
- Allow selective transit AS

Steps

Prevent transit AS

In this section you will configure your BGP AS 64500 so that it no longer acts as a transit AS between the external providers.

1. Review the current advertisements: On Agg1, review the routes that are advertised to routerA.

Q: Do you see other routes than the local AS being advertised to the routerA?

A: Yes, all BGP routes, including routes learned via routerB, are advertised to routerA.

2. On Agg2, review the routes that are advertised to routerB.

Q: Do you see other routes than the local AS being advertised to the routerB?

A: Yes, all BGP routes, including routes learned via routerA, are advertised to routerB.

3. On Agg1, define an AS-path list name 'bgp-local' that only contains an empty AS path regular expression.
4. On Agg1, define a route-map named 'bgp-ebgp-out' that permit routes matching this AS-path list.
5. On Agg1, apply the route-map for outbound traffic on the routerA eBGP peer for IPv4 unicast routes.
6. On Agg1, now review the IPv4 routes that are advertised to routerA.

Q: What are routes that are advertised?

A: All routes are still advertised.

Q: Why are all routes still advertised?

A: The policy change will be effective after the BGP session is cleared.

7. On Agg1, clear the BGP session with the routerA.
8. Review the advertised routes to routerA.

Q: What are the routes that are advertised?

A: Only the routes that have an empty AS Path list, that is the 100.100.1.0/24 prefix, should be advertised.

9. Repeat these steps on Agg2.

10. On Agg2, verify that only the 100.100.1.0/24 prefix is advertised to routerB.

Allow selective transit AS

While the previous section ensured that you successfully blocked any transit traffic, your partner network on BGP AS 64600 now complains that they are no longer reachable.

So you should configure your AS to allow transit traffic for routes to your partner AS.

11. On Agg1, make a new AS Path list 'partner' for the routes that originate in your partner AS.

12. On Agg1, update the outbound route-map with a new entry to match on the AS Path list 'partner' and allow it to be advertised to routerA.

13. On Agg1, verify the advertised routes to routerA now include the partner routes.

14. Repeat this configuration on Agg2.

Task 5: Outbound Traffic Route Control

Objectives

- Apply default local preference
- Apply selective local preference
- Apply weight

Steps

Apply default local preference

Local preference allows the BGP administrator to apply a cost to route that will be advertised inside the AS. So when 2 routers advertise the same route into the AS, the route with the highest preference (most preferred) will be selected by the other BGP routers in the AS.

The routerC has been pre-configured to advertise:

- route 198.19.100.0/24 equally to routerA and routerB.
- route 198.19.101.0/24 with a shorter AS path via routerA
- route 198.19.102.0/24 with a shorter AS path via routerB

1. On Agg1, review the BGP IPv4 routing table for the entries

198.19.100.0/24
198.19.101.0/24
198.19.102.0/24

Q: How many entries do you see for each of these routes?

A: On Agg1, for 198.19.100.0/24 and 198.16.102.0/24 there should be 2 entries, 1 via routerA and 1 via Agg2.

For 198.19.101.0/24, only the route via routerA will be present.

2. On Agg2, review the BGP IPv4 routing table for the entries

198.19.100.0/24
198.19.101.0/24
198.19.102.0/24

Q: How many entries do you see for each of these routes?

A: On Agg2, for 198.19.100.0/24 and 198.16.101.0/24 there should be 2 entries, 1 via routerB and 1 via Agg1. For 198.19.102.0/24, only the route via routerB will be present.

3. On Sw1, review the IP routing table for entries

198.19.100.0/24
198.19.101.0/24
198.19.102.0/24

Q: Why do Agg1 and Agg2 have 2 entries in the BGP table for these routes, while Sw1 has only 1 entry?

A: A BGP router will only advertise its best path to its peers. If you would review the best path for 198.19.101.0/24 on both Agg1 and Agg2, you would observe that Agg1 actually points to Agg2 as the next hop (based on the shortest AS Path). Since Sw1 only receives the route to Agg2, only that route is present in the BGP routing table.

Q: What are currently the next-hop paths for the best routes?

A: 198.19.101.0/24 via Agg1, 198.19.102.0/24 via Agg2. The path to 198.19.100.0/24 is advertised by both Agg1 and Agg2, but the BGP route selection process has selected the route to Agg1 as the best path for the routing table.

4. On Agg2, review the IP routing table for entries

198.19.100.0/24
198.19.101.0/24
198.19.102.0/24

Q: What are currently the next-hop paths for the best routes?

A: 198.19.100.0/24 and 198.19.102.0/24 via routerB, 198.19.101.0/24 via Agg1.

5. Now configure the network so that all entries advertised by Agg1 are more preferred than the routes that are advertised by Agg2. You can achieve this by changing the default local preference on Agg1 to '200'.

6. On Sw1, review the BGP IPv4 routing table for entries 198.19.100.0/24 and 198.19.102.0/24.

Q: Do you see the updated local preference?

A: Yes, the route advertised by Agg1 should now have a local preference of '200'.

7. On Sw1, review the IP routing table for entries 198.19.100.0/24 and 198.19.102.0/24.

Q: What is the next-hop for this route?

A: Only the Agg1 route is now in the routing table, since that is the 'best' route.

8. On Agg2, review that it is now using the route to Agg1 for the entry 198.19.102.0/24.

Apply selective local preference

While the default local preference is easy to make 1 router the preferred router, it doesn't help when more granular route control is required.

In this section you will apply granular control of the preference to these routes:

- 198.19.101.0/24 should exit the AS via Agg1 as preferred path with preference 300. The path via Agg2 should get the default preference 100.
- 198.19.102.0/24 should exit the AS via Agg2 as preferred path with preference 300. The path via Agg1 should get the default preference 100.
- 198.19.100.0/24 should exit the AS via both Agg1 and Agg2 with the same preference 300.

9. On Agg1, configure prefix lists to select the required prefixes.
10. On Agg1, configure a route-map that assigns each prefix the required preference.
11. Apply this route-map on the routerA BGP peer in the inbound direction.
12. On Agg2, repeat these steps with the adjusted local preference values.
13. On Sw1, review the BGP IPv4 routing table using 'show bgp ipv4 unicast'. Check the local preference for the 3 requested prefixes.

Q: Are there any other steps required to make this route-map effective?

A: Clear the BGP session to the routerA on Agg1 and to routerB on Agg2.

14. On Sw1, review the BGP IPv4 routing table. Check the local preference for the 3 requested prefixes.
15. On Sw1, review the IP routing table and verify that the correct next hop is selected for each of these routes.

AS multipath-relax

Based on the standard BGP path selection, when 2 paths have the same AS length, local preference and all other values are equal, the lowest next hop ID will be the selection criteria.

In the example of Sw1, the 198.19.100.0/24 prefix has the same AS path length, local preference and metric for the paths via Agg1 and Agg2.

Still, since the actual AS-path of both entries are different (64511 64513 vs 64512 64513), BGP will not apply ECMP load balancing on these routes.

A network administrator can choose to 'relax' this rule, so even when the AS path content is different, when they have the same length, load-balancing can be performed.

This feature is known as AS multipath relax.

16. On Sw1, review the current BGP IPv4 routing table and the IP routing table.

17. On Sw1, configure the AS multipath relax feature, confirm the reset of the sessions.

18. On Sw1, review the current BGP IPv4 routing table.

Q: What do you observe for the status codes of the 198.19.100.0/24 prefix?

A: The second route for 198.19.100.0/24 is now marked with = for multipath.

19. On Sw1, review the IP routing table.

Optional: Test link failure and route failover

20. On Agg1, disable the uplink 1/1/7.

21. On Sw1, verify the next-hop for the 198.19.101.0/24 network is Agg2.(this was Agg1 previously)

22. Review the BGP IPv4 routing table.

Q: What is the local preference listed for the 198.19.101.0/24 entry?

A: The route is now shown with the local preference 100.

23. On Agg1, restore the link 1/1/7.

24. Verify on Sw1 that the route 198.19.101.0/24 has reverted to Agg1 as the next-hop.

NOTE: It may take a few moments for Agg1 to establish the BGP session with routerA.

Using weight locally on a router

While the local preference is a convenient tool to control the traffic path over the entire AS, it may be necessary to overrule the local preference on individual routers, while the other routers in the AS would still follow the local preference rules.

The BGP weight can be used locally on a BGP router to overrule the local preference. The weight value is stored in the BGP tables, but it is not an official attribute that will be exchanged with other BGP routers. As such, the 'weight' is only locally significant on the local router.

Setting the weight on a router will not be visible on any other router in the AS, since it is simply not exchanged as part of a BGP update process.

In this lab, you have configured your AS to send traffic for the 198.19.101.0/24 prefix via the Agg1 by configuring the local preference on Agg1 to '200'. So Sw1 and Agg2 will respect this by default.

In this section, you will overrule this local preference on Agg2 by setting the weight value. This means that Agg2 will forward the traffic directly based on the weight, while the rest of the AS, including Sw1, will still use the local preference for the forwarding.

On Agg2, you will update the route-map that controls the inbound routes from routerB.

25. On Agg2, review the BGP IPv4 routing table and check the current weight value for the 198.19.101.0/24 route.

26. On Agg2, check the IP routing table next hop for the 198.19.101.0/24 route.

Q: What is the current next-hop IP?

A: Currently, the next-hop IPs should point to the Agg1 switch based on the local preference that was set in the previous section.

27. On Agg2, modify the route-map that controls the inbound 198.19.101.0/24 prefix and apply the weight 400.

28. On Agg2, review the BGP IPv4 routing table

Q: Was the weight change applied?

A: No, to make the route-map changes effective, a route refresh or hard reset is required.

Q: What is the advantage of a Route Refresh compared to a hard reset?

A: A Route Refresh does not tear down the BGP session, so it has less impact on the traffic forwarding.

On Agg2, review the debug logging for a Route Refresh.

29. On Agg2, clear the debug buffer.

30. On Agg2, use a route refresh to refresh the routes from the routerB. This should apply the changed weight value.

31. On Agg2, review the debug buffer. The BGP session remained established and the routers simply refresh the routes.

32. On Agg2, review the BGP IPv4 routing table and check the new weight value for the 192.19.101.0/24 route.

33. On Agg2, check the IP routing table next hop for the 198.19.101.0/24 route.

Q: What is the current next-hop IP?

A: The next-hop should have changed to routerB based on the weight configuration.

34. Verify the path selection in the rest of this AS. This should not have been impacted by the weight configuration on Agg2. On Sw1, verify the IP routing table for the 198.19.101.0/24 route.

Q: Did the next hop IP for Sw1 change?

A: No, since the weight is not communicated inside the AS, Sw1 is still using the local preference value and will use Agg1 as the next hop IP.

Soft reconfiguration inbound

A BGP router will process the received routes when they arrive. This means that the inbound route-map will be applied to the inbound routes when the routes arrive from the peer.

As you have previously seen, in case a route-map is modified, these changes do not take effect immediately since they will only be processed the next time the routes are received. This is why the hard reset or the Route Refresh action was required.

It is also possible to maintain all the received routes from a peer in a dedicated table, independent of the main BGP routing table. This configuration allows the BGP process to apply route-map changes based on the local database, without the hard session reset or without a route refresh, since a local 'full' copy of the received routes is available.

This can also make troubleshooting easier. For example: if an inbound route-map would have a mistake and it would remove a route, that route would not be visible in the main BGP routing table. So an administrator may think that the peer did not advertise the route.

When the received routes are kept in a separate table, the administrator could see the exact list of routes that were received.

NOTE: Enabling this feature will consume additional resources on the system, since all the received routes for the configured peers are stored, on top of the main BGP table. In case of large number of routes, the Route Refresh feature may be considered.

You will now configure Agg1 and Agg2 to maintain a separate table for the eBGP peers.

35. On Agg1, review the inbound routes from the routerA

Q: What do you observe?

A: There are no routes listed. This is because the switch does not maintain this separate table by default. So this output is not an error, but by design.

36. Review the command options using the ?

37. To see the currently accepted routes on Agg1 in the main BGP table that use this peer as the next hop, use the 'routes' option.

38. On Agg1, clear the debug buffer and configure the peer 10.255.101.11 with the soft reconfiguration option.

39. On Agg1, review the debug buffer. received routes.

Q: Did this configuration have any impact on the BGP session?

A: Yes, changing the soft-reconfiguration settings for a peer will result in a hard reset to populate the separate table.

40. On Agg1, review the received routes, now the table should have been populated.

Soft reconfiguration inbound impact on route refresh

In the previous section you have seen how a local copy of the received routes can be convenient for troubleshooting. In this section, the local table will be used to apply a configuration change without the hard reset or the route refresh.

In the previous section of this lab, a debug was used to demonstrate that a ROUTE REFRESH message was sent after the command 'clear bgp ipv4 unicast <peer-IP> soft in' was executed.

Now you will configure the separate table and use the same command. This time the changes will be processed based on the local table instead of sending out the ROUTE REFRESH message to the peer.

41. On Agg2, enable the feature for the routerB eBGP peer. Make sure to refresh the routes to populate the table.

42. On Agg2, modify the route-map that controls the inbound 198.19.101.0/24 prefix and apply the weight 500 (in the previous example you applied weight 400).

43. On Agg2, review the BGP IPv4 routing table.

Q: Was the adjusted weight applied automatically?

A: No, the administrator still needs to trigger the re-evaluation process for the routes.

44. On Agg2, clear the debug buffer, then refresh the routes.

45. On Agg2, review the BGP IPv4 routing table.

Q: Was the new weight of 500 applied?

A: Yes, the weight was applied.

46. On Agg2, review the debug buffer.

Q: Was a ROUTE REFRESH message sent to the routerB peer?

A: No, the same command was used, but since the full peer table was now available, the new policy is processed based on this local table.

This concludes the soft reconfiguration inbound feature.

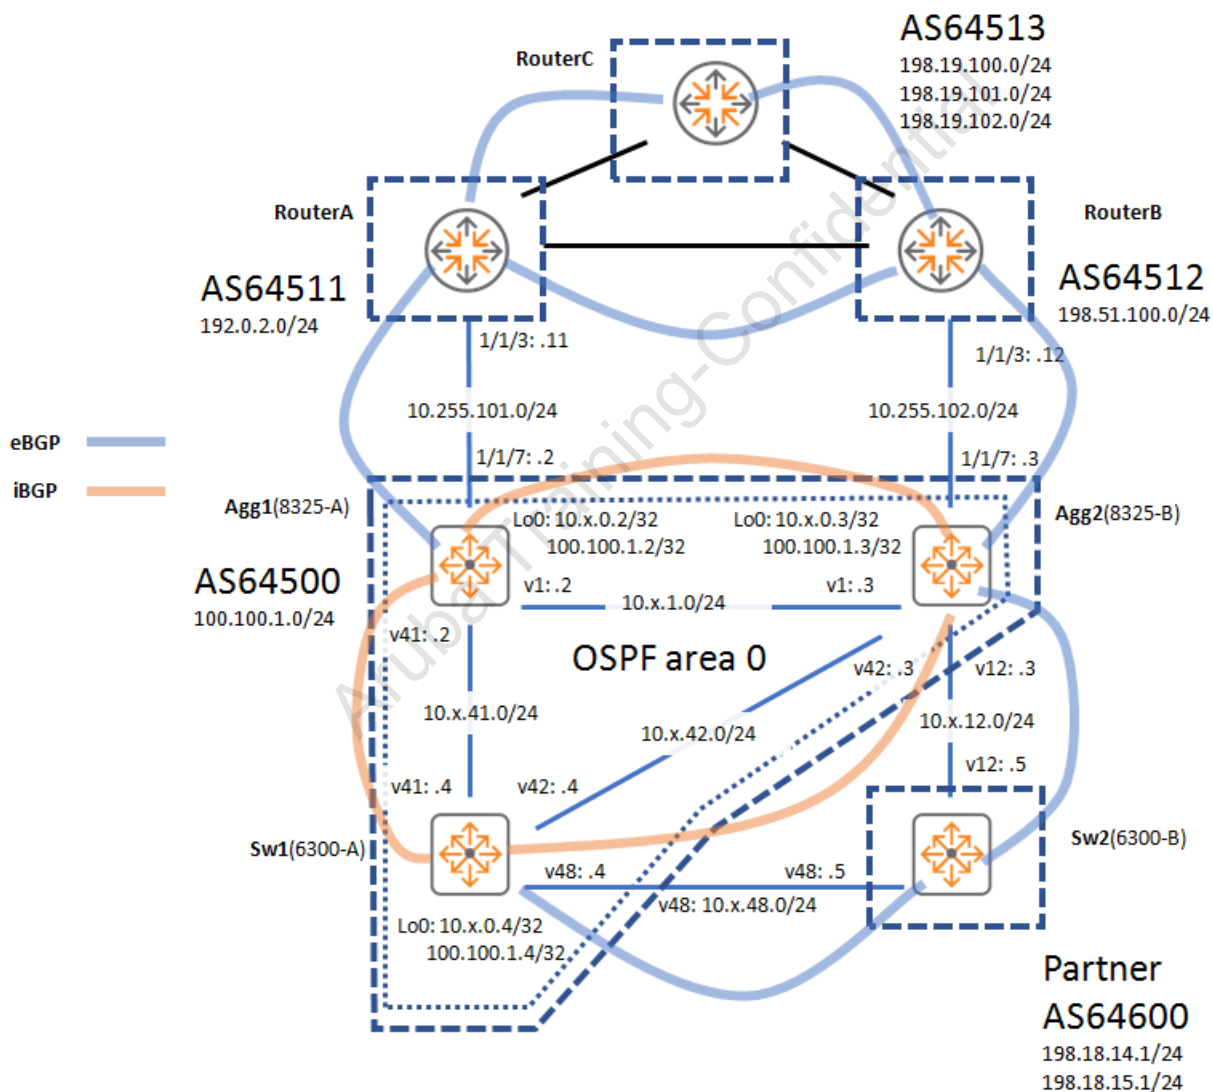
Aruba Training-Confidential

Task 6: Inbound traffic route control

Objectives

- Apply AS-Path length for outbound advertisements
- Use MED for inbound route preference
- Use community attributes to control a routing update

Lab Diagram



Aruba Training-Confidential

Steps

Apply AS-Path length for outbound advertisements

While it is not possible to have perfect inbound traffic control with BGP, some tricks can be applied to make a certain path 'appear' to be better than other paths.

One of the BGP path selection processes (after for example weight and local preference) is the AS Path Length: the shortest AS Path will be preferred.

By tweaking the AS Path length and making it artificially longer on 1 path, the other path could become the more preferred path.

IMPORTANT: Keep in mind that this is not a guarantee, since any upstream BGP router may apply local preference or weight values to your routes, so the AS Path Length would not even be considered.

In this lab, you are announcing your 100.100.1.0/24 network via both Agg1 and Agg2 to the external providers. RouterC will simply pick the shortest AS Path to reach the 100.100.1.0/24 destination.

By artificially increasing the path advertisement on the Agg2 path, you may (not guaranteed!) influence the path selection process for the other BGP peers, since routerC will see that the path via Agg2 will have more 'AS hops' than the path via Agg1.

1. On routerB, review the BGP routing table entry for 100.100.1.0/24.

Q: What is currently the next-hop IP for the 100.100.1.0/24 route?

A: Since no weight or local preference has been configured on the routerB, the AS Path Length is used, so the path to Agg2 will be the shortest path.

2. On Agg2, you have an existing outbound route policy with rules that allow the advertisement of your routes and your partner routes.
3. On Agg2, review the advertised routes to routerB.
4. Change the route-map rule on Agg2 that controls your own routes (empty AS path list) and use the 'set' action to perform an AS Path prepend. Prepend your own AS number 3 times in this entry.
5. On Agg2, review the advertised routes to routerB, they should now have a longer AS Path list.

Q: When is a new policy applied?

A: The next time the routes are exchanged, either based on the peer routerB that would ask for a route refresh, or the Agg2 that could perform a hard reset.

6. On Agg2, perform a hard reset of the routerB session.
7. On Agg2, review the advertised routes to routerB again, they should now have a longer AS Path list.
8. On routerB, review the BGP routing table entry for 100.100.1.0/24 again.

Q: What is the best route?

A: routerB will now see that the path via routerC has a shorter AS Path Length and that will be its next hop IP.

Use MED for inbound route preference

The multi-exit discriminator, or MED, is another way to influence the inbound route control. The MED value can be used to hint your peer AS to use 1 path over another path. Just like the AS Path Length, you should realize that this is just a 'hint', so the partner network can simply overrule your suggestion by configuring a local preference or weight in their setup.

While the AS Path is exchanged over multiple Autonomous Systems, the MED value is only used by your peer AS, and it will not be advertised to other systems beyond the peer AS.

This is why the MED cannot be used to influence the inbound path when you are peering with 2 different AS systems, or 2 providers. In that case the AS Path Prepend can be used to hint a less preferred path.

In the lab setup, you will use the MED value to hint to the Sw2 how you want to receive traffic into your AS.

First you will enable the second link to the Sw2 AS. This is the link between Sw2 and Agg2. This will ensure that Sw2 has 2 paths to choose from to send traffic to your AS.

Remember that the MED acts like a 'cost', so in this case, the lowest MED value will be the preferred path.

Configure your AS to:

- Receive traffic for 100.100.1.0/24 network via Agg2 as metric 10, Sw1 must be backup with metric 100.
- All other routes are still advertised, but without any metric modification.

Prepare the link between Agg2 and Sw2

You will use the existing SVI 12 to setup a Layer3 link to Sw2 VRF 'bgpext'.

9. On PC1, open desktop > ASTS > config > ASTS_cfg_lab06-task6-sw2.txt. Apply this configuration to Sw2. These are the commands from the file:

```
interface vlan12
description BGP-Agg2
```

```
vrf attach bgpext
ip address 10.x.12.5/24
exit

router bgp 64600
vrf bgpext
neighbor 10.x.12.3 remote-as 64500
address-family ipv4 unicast
neighbor 10.x.12.3 activate
```

10. On Agg2, verify you can ping to 10.x.12.5.
11. On Agg2, configure the eBGP peering to Sw2 (10.x.12.5)
12. On Agg2, verify the eBGP peering is established with Sw2.
13. On Sw2, review the current BGP IPv4 routing table, check the metric values for the 100.100.1.0/24 route.

Apply the MED values

14. On Sw1, define a prefix list for the 100.100.1.0/24 prefix.
15. Configure a route-map on Sw1 that applies the requested metric prefix list, make sure all other routes are still allowed.
16. Apply the route-map to the Sw2 eBGP peer.
17. Review the advertised routes to Sw2
18. Make sure the new policy is applied to the Sw2 peer.
19. Review the advertised routes again. You should see the adjusted metric values.
20. On Sw2 in VRF bgpext, verify if the MED values are shown in the BGP IPv4 routing table.
21. On Agg2, define the required prefix list.
22. Configure a route-map on Agg2 that applies the requested metric to the 100.100.1.0/24 prefix and allows all other prefixes.
23. On Agg2 , apply the route-map to the Sw2 eBGP peer and make it effective.
24. On Sw2 in vrf bgpext, verify if the MED values are shown in the BGP IPv4 routing table.
25. The next-hop IP for the 100.100.1.0/24 prefix should point to Agg2.
26. Verify that all other BGP routes are still received from both Sw1 and Agg2.

This demonstrates how the MED value can be used to influence the inbound traffic.

Using Community Attributes for route control

BGP advertisements can include many attributes, one of them is the community attribute. This gives the network administrator the option to set a 'mark' or a tag on a route advertisement.

The peer router can then use this mark to take some action on the routes.

These are some 'well-known' community attributes, that provide control over the route advertisements.

Attribute	Description
internet	Advertise the prefix to all BGP neighbors.
local-as	Do not advertise the prefix outside the sub-AS.
no-advertise	Do not advertise the prefix to any BGP neighbors.
no-export	Do not advertise the prefix to any eBGP neighbors.

Currently, the Sw2 BGP partner is advertising 2 prefixes the customer:

198.18.14.0/24
198.18.15.0/24

Now the Sw2 partner company wants that the 198.18.15.0/24 prefix is only advertised to the 64500 AS, but this prefix should not be advertised upstream to the routerA and routerB systems anymore.

They could ask the administrator of 64500 to apply a route policy on their eBGP peers, but Sw2 can also control this themselves using a community attribute.

In this section, you will configure Sw2 so the route 198.18.15.0/24 is only advertised to AS 64500, but the BGP routers of AS 64500 will not advertise this route to the routerA or routerB.

27. First verify that Agg1 is advertising the 198.18.15.0/24 prefix to routerA. At the end of this section, this prefix should not be advertised anymore to routerA.
28. On Sw2, define a prefix list 'net15' that controls the prefix 198.18.15.0/24.
29. On Sw2, define a route-map 'bgp-ebgp-out'. Use a match on the 'net15' prefix list and apply the community attribute 'no-export'. Use the ? to explore the options.
30. On Sw2, define another sequence in the route-map to allow all other routes. This is required to advertise the 198.18.14.0/24 prefix.
31. On Sw2, apply the route-map 'bgp-ebgp-out' to both eBGP peers Sw1 and Agg2.

NOTE: On Sw2, BGP is running inside a VRF context, so make sure to access the VRF under the BGP context.

32. On Sw2, review the current bgp configuration

33. On Sw2, clear all BGP sessions.

Attempt to verify the community attribute on Sw1

34. On Sw1, review the BGP IPv4 details for the 198.18.15.0/24 route.

Q: Do you see the community attribute on the route with next-hop 10.x.48.5?

A: No, there is no community attribute listed.

Q: What could be wrong with the configuration?

A: In the BGP configuration, the sending of community attributes must be enabled on the peer.

35. On Sw2, update the BGP configuration, enable the 'send-community' option on both peers under the IPv4 address-family.

NOTE: The command 'send-community' without any options will enable both standard and extended attributes.

36. On Sw2, verify the BGP configuration.

37. On Sw1, review the BGP IPv4 details for the 198.18.15.0/24 route.

Q: Do you see the community attribute on the route with next-hop 10.x.48.5?

A: Yes, the community now shows the 'no-export' entry.

Q: Do you see the community attribute on the route with next-hop 10.x.0.3 (Agg2)?

A: No, this update did not receive the 'no-export' attribute.

38. Review the BGP routing table on Agg2.

Q: Do you see the community attribute on the route with next-hop 10.x.12.5?

A: Yes, the community now shows the 'no-export' entry.

Q: So Agg2 is receiving the attribute, but it is not sending it in the update to Sw1, what could be wrong?

A: Again, the iBGP peers within AS 64500 do not have the 'send-community' configured, so they are not sending this information to their peers.

Correct the iBGP peer configuration

39. On Agg2, enable Agg1 and Sw1 peers with the 'send-community' attribute.

40. On Sw1, enable Agg1 and Agg2 peers with the 'send-community' attribute.

41. On Agg1, apply the configuration to the peer group 'internal'

42. On Agg1, verify the BGP neighbor details for peer Agg2 (10.x.0.3)

Q: What do you observe for the Send Community?

A: The Send Community is set to 'both' and the ^ sign indicates it was inherited from the peer-group.

43. On Agg1, review the BGP IPv4 routing table entry for 198.18.15.0/24

Q: Do both entries have the community 'no-export' set?

A: Yes, now both entries have the community set.

44. On Agg1, review the advertised routes to routerA.

Q: What do you observe?

A: The 198.18.15.0/24 is no longer advertised. The route 198.18.14.0/24 is still advertised.

This concludes the community configuration example.

Optional Task 7: AS Path List exclusion

Objectives

- Understand how to remove an AS from the AS Path list

Steps

In this task, you will see how to remove an AS from the AS Path List.

You have a peering with the partner company on AS64600 and the partner is advertising routes 198.18.14.0/24 and 198.18.15.0/24. Based on the previous task, only 198.18.14.0/24 is advertised to the routerA and routerB.

When this route is advertised to the external provider (routerA and routerB), the partner AS (64600) should be removed from the AS Path list.

1. On Agg1, review the currently advertised routes to routerA

Q: For the 198.18.14.0/24 route, what are the AS numbers in the AS Path?

A: Both 64500 and 64600 are listed.

Q: Do you have an existing outbound route-map to routerA (10.255.101.11) ?

A: You can use these commands to review the configuration

2. On Agg1, review the contents of the route-map to routerA
3. On Agg1, modify the existing outbound route-map to routerA. You should modify the rule that handles the 'partner' routes. Add a 'set' instruction to remove the AS number 64600 for the routes that belong to AS 64600.
4. On Agg1, while in the route-map context, use 'show run current' to review the route-map configuration.
5. On Agg1, now ensure that the route-map changes are applied to the routerA BGP session.
6. Review the advertised routes to routerA.
7. On Agg2, repeat this procedure.
8. On RouterC, verify that the AS 64600 no longer appears in the AS Path for 198.18.14.0/24 prefixes.

This concludes the AS Path exclude feature.

Optional Task 8: Using a Route Reflector for iBGP

Objectives

- Understand the iBGP full mesh requirement.
- Understand how Route Reflector is an alternative to the iBGP full mesh.

Steps

In this task you will explore the BGP route reflector feature. In the current lab topology, Agg1, Agg2 and Sw1 are connected as iBGP full mesh, so they each have an iBGP peering with each other. While this is fairly simple in the current lab with only 3 routers, in larger iBGP topologies it quickly becomes very difficult to maintain this full mesh requirement between all iBGP peers.

This is where route reflectors provide a solution.

In this task you will first explore what happens when iBGP does not have a full mesh topology. Next you will configure the route reflector to correct this behavior.

Break the iBGP full mesh topology

In this section you will break the full mesh by removing the peering between Sw1 and Agg2 switches.

1. On Sw1, remove the neighbor for Agg2 (10.x.0.3)
2. On Sw1, verify the peering has been removed.
3. On Agg2, remove the neighbor for Sw1 (10.x.0.4)

Review the result of the broken topology

4. On Sw1, check the BGP IPv4 routing table, you should notice that the 198.19.102.0/24 network is no longer available in the routing table.
5. On Agg1, check the BGP IPv4 routing table for entry 198.19.102.0/24

Q: Did Agg1 receive the route?

A: Yes, since it has a direct iBGP peering with Agg2.

Q: In the BGP routing table, you should see that Agg1 has received the 198.19.102.0/24 route from the routerA, this is an eBGP peer. Why was this route not advertised to Sw1?

A: While it is correct that an eBGP route would be advertised to the iBGP peers, such as Sw1, in this case, due to route policies, the route is not the best path anymore, so it will not be advertised. Agg1 assumes that Agg2 has advertised this best route to Sw1 directly.

Configure the route reflector

In this section you will configure Agg1 to become the route reflector. This is done by making the Sw1 peer a route reflector client.

Keep in mind that the route reflector will 'reflect' routes on behalf of the route-reflector clients.

So there is no need to make Agg2 a route-reflector client as well in this topology.

Since Agg2 is a normal iBGP peer, every update Agg1 receives from Agg2 will be reflected to the configured route reflector clients, in this case Sw1.

Every update Agg1 receives from a route-reflector client (such as Sw1) will be reflected by the route reflector to the normal iBGP peers, such as Agg2 in this case.

6. On Agg1, under the IPv4 address-family, configure the Sw1 (10.x.0.4) as a route-reflector client.
7. On Sw1, verify that the external route for 198.19.102.0/24 was now received in the BGP IPv4 routing table. Review the details of the route using

This demonstrates how a route-reflector can help when an iBGP full-mesh is not possible or too complex.

Q: Do you see a value in the cluster list?

A: Yes, the route-reflector will add its cluster-id to the reflected routes.

NOTE: In this lab, a single device (Agg1) is used as the route reflector. In a redundant configuration, the 2 route reflectors that have the same roles should use the same cluster-id in their configuration.

8. Optional: on Agg1, Agg2, Sw1 and Sw2, save a new checkpoint '**asts-lab06-<yourname>**'.

This concludes the BGP Lab activity.

Lab 07: Route Redistribution

In this lab you will configure route redistribution between various routing sources.

- Perform route redistribution of static routes into OSPF
- Using route-maps to control the redistribution and control the cost
- Using prefix lists to select routes
- Perform route redistribution between OSPF and BGP
- Explore the risk of routing loops
- Control mutual redistribution with route-maps
- Use OSPF route tags

Requirements

This lab requires completion of Lab 05.

Scenario

In this setup, the routerA, routerB and routerC are no longer the Service Provider routers as seen in the BGP lab, but they belong to the same organization of a partner network.

The partner has these 2 networks that need to be reachable from the customer network:

- 10.255.21.0/24
- 10.255.22.0/24

Setup

Load checkpoint asts-lab07-1-rr on routerA, routerB and routerC, save and reboot the 3 routers.
On the 4 switches, load the checkpoint for asts-lab05-<yourname>.

Static route into OSPF

This is another scenario for your customer demonstration.

1. The customer wants to ensure that their internal router (Sw1) can reach the partner networks 10.255.21.0/24 and 10.255.22.0/24 via either Agg1 or Agg2. There should be no configuration changes on Sw1.
2. The customer wants to see that all routes that are redistributed by Agg1 are more preferred than routes that are redistributed by Agg2. Both Agg1 and Agg2 should use a higher cost than the default.
3. The customer wants to see:
 - a. Traffic to the 10.255.21.0/24 network using Agg1 as the primary path
 - b. Traffic to the 10.255.22.0/24 network using Agg2 as the primary path.
 - c. If either Agg1 or Agg2 fails or their uplink fails, the other Aggregation switch is used.
4. Next, the customer wants to see a small adjustment. For traffic destined to the 10.255.22.0/24 network, they want to ensure that the internal path cost is considered. This should apply only to the 10.255.22.0/240 network.

Prefix Lists

The customer wants you to demonstrate some examples of how to use IP prefix lists.

On PC1 Desktop > ASTS > open '**ASTS-lab 07-prefixlist-examples.txt**' and paste the example static routes into Agg1 configuration.

They want you to demonstrate these scenario's (one by one), you should only change the configuration on Agg1.

1. Scenario 1: Only the exact prefix 10.255.64.0/23 should appear in the Sw1 routing table.
2. Scenario 2: All of the routes included in the 10.255.64.0/23 prefix should appear in the Sw1 routing table

3. Scenario 3: All of the routes included in the 10.255.64.0/23 prefix, except for the summary route itself (10.255.64.0/23) should appear in the Sw1 routing table.
4. Scenario4: All of the routes included in the 10.255.64.0/23 prefix, except for the summary route itself (10.255.64.0/23) and except for host routes (/32) should appear in the Sw1 routing table.

Route Redistribution between OSPF and BGP

Load checkpoint **asts-lab07-3-rr** on routerA, routerB and routerC, save and reboot the 3 routers. Remove the static route to 10.255.21.0/24 on Agg1 and Agg2.

Single Link Redistribution

On Agg1, establish eBGP peering with routerA. Make sure that the customer OSPF routes are advertised to the partner network using eBGP and that the partner eBGP networks are accessible by the internal OSPF routers (Sw1). Sw1 should be able to reach the 10.255.21.0/24 network.

Dual Link Route Redistribution

The customer has heard about potential issues, such as route feedback, with dual links and route redistribution.

They want you to demonstrate how route feedback occurs.

1. Establish an eBGP peering between Agg2 and routerB. Configure Agg2 so it:
 - a. Advertises the OSPF routes to routerB using eBGP
 - b. Advertises the eBGP routes into OSPF.
2. Demonstrate to the customer what happens with the 10.255.21.0/24 route on Agg1 and routerA.
3. Now that the customer has seen the issue of route feedback, you should configure the customer routers Agg1 and Agg2 to prevent the route feedback using route tags. The result should be that Agg1:
 - a. Only advertises the customer OSPF routes via eBGP into routerA
 - b. Does not advertise the OSPF external routes of the partner network (that are redistributed by Agg1).

The same should also apply to Agg2 for OSPF routes that are redistributed by Agg1.

Task 1: Load the start configurations

Objectives

- Load the lab activity start configurations

Steps

The routerA, routerB and routerC will have the role of a partner network in this lab activity. You will need to setup route redistribution with the partner network to provide IP connectivity to some networks in the partner network.

You will first load the required start configuration so they can simulate a partner network.

1. On routerA, routerB and routerC load the 'asts-lab07-1-rr' checkpoint to the running configuration, save and reboot the system.

```
copy checkpoint asts-lab07-1-rr run
write mem
boot system
```

NOTE: The reboot is requested here since the virtualCX system may not always load the interfaces after a checkpoint swap. The reboot ensures that the configuration is actually applied.

2. Revert check point to asts-lab05-<yourname> on all 4 switches (Agg1/Agg2/Sw1/Sw2)

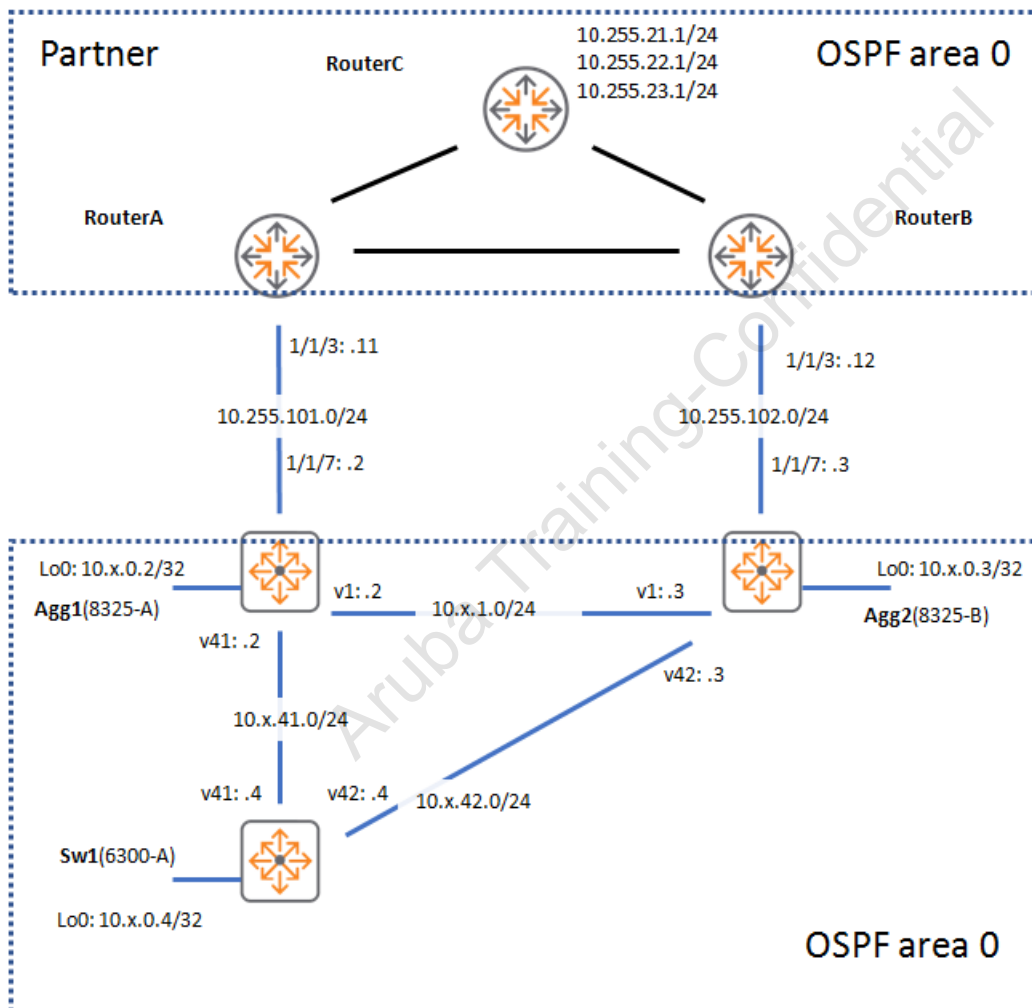
```
copy checkpoint asts-lab05-<yourname> running
```

Task 2: Route redistribution of static routes into OSPF

Objectives

- Learn how to redistribute all static routes into OSPF
- Control the OSPF costs of the redistributed routes

Lab Diagram



Steps

Redistribute static routes into OSPF

NOTE: The partner network has been configured already at this point, so only outbound configuration is required. The partner network has a static return route for 10.x.0.0/16 pointing back to the Agg1 and Agg2.

1. On Agg1, add 2 static routes to the test subnets in the partner network. Use the routerA as the next-hop IP.
2. On Agg1, attempt to ping 10.255.21.1 using your loopback0 as source interface.

NOTE: routerA has a static return route for 10.x.0.0/16, but your uplink (1/1/7) to routerA has IP address 10.255.101.2/24, which is not in that range, so a matching source address must be used.

3. On Agg1, redistribute all static routes into OSPF.
4. On Agg2, define the same static routes, but use routerB next hop (10.255.102.12)
5. Repeat the redistribution on Agg2.
6. On Sw1, check the IP routing table and metric for the 10.255.21.0 route.

Q: What is the default OSPF external type and the default metric value?

A: By default, redistributed routes are External Type 2 (E2) and they have a default cost of 25.

Q: How many route entries are shown for the 10.255.21.0/24 network?

A: The paths to Agg1 and Agg2 have the same cost, so the OSPF calculation will select both as best path and offer both to the IP routing table.

Cost Manipulation of the redistributed routes

It is possible to set a default metric under the OSPF process. This will be used as the metric for all routes that are redistributed into the routing protocol. It is an easy way to set the redistribution metric if no granular route control is required.

7. On Agg1, change the default metric to 6000.
8. On Sw1, verify the result in the IP routing table and the OSPF routing table.


```
via 10.x.42.3 interface vlan42, cost 25 distance 110
```

Q: How many routes are currently available to the 10.255.21.0/24 network?

A: Since the path to Agg1 has a higher cost, the path to Agg2 will be selected as the best path. Only this route will be proposed to the IP routing table.

9. On Sw1, review the OSPF LSDB for external routes.

Q: How many LSA entries are shown with a Link State ID of 10.255.21.0 ?

A: 2 entries, this means that the higher cost entry is in the Link State Database, but since it was not the best route, it has not been selected.

10. On Agg2, change the default metric to 7000.

11. On Sw1, verify the result of the Agg2 configuration in the IP routing table and the OSPF routing table.

Q: Did the next hop and the cost change?

A: Yes, now the Agg1 route will be the best route again.

Granular control of the redistribution using route-maps

In this section, route-maps will be used to apply more granular route control on the redistributed routes. Keep in mind that you are only configuration the outbound direction of the traffic, that is traffic that is destined to the partner network. The partner network has its own control over how traffic leaves their network.

12. The requirements in this section are:

- Traffic to destination network 10.255.21.0/24 should take the path via Agg1 if that path is up. Otherwise, traffic should take the path via Agg2.
- Traffic to destination network 10.255.22.0/24 should take the path via Agg2 if that path is up. Otherwise, traffic should take the path via Agg1.
- Use cost 1000 and 2000 for the primary and secondary paths.

13. On both Agg1 and Agg2, define prefix-lists for the required subnets.

14. On Agg1, define a route-map named 'static_to_ospf1' and apply the correct initial metrics.

15. On Agg1, apply this new route-map to the OSPF static route redistribution

16. On Sw1, verify the cost on the 2 destination subnets

Q: Did the cost change for the 2 external routes?

A: Yes, the route-map on Agg1 applied a lower cost than the default metric of Agg2 (7000). Currently Agg1 is the best path for both 10.255.21.0/24 and 10.255.22.0/24.

17. On Agg2, define a route-map 'static_to_ospf1' based on the given requirements.

18. On Agg2, apply the new route-map to the OSPF static route redistribution.

19. On Sw1, review the result. Both external routes should now be listed with metric 1000, 1 to each Aggregation switch.

OSPF External Metric Type1

With the default OSPF External Metric Type 2, the initial cost is maintained inside the OSPF AS. This is convenient when traffic to the external network must always use the same exit point, due to firewall requirements for instance.

When there is no such requirement, it is possible to send the traffic to the partner using the 'closest' exit point. This means that the internal path cost to reach the external network should be added to the external path cost. This is achieved by making the external routes use OSPF External Metric Type 1.

NOTE: OSPF will prefer external type 1 routes over external type 2 routes, even when the cost would seem higher.

Let's configure this for the 10.255.22.0/24 network:

20. On Agg1, change the route-map entry for the 10.255.22.0/24 prefix list, to set the metric type to external type 1. Leave the cost 2000 in the entry.

21. On Sw1, review the impact on the routes. Compare this with the previous output.

Q: What do you observe?

A: Previously, the path to 10.255.22.0/24 was E2 with a cost of 1000. Now the best route is E1 with a cost of 2100. This confirms that E1 routes are more preferred than E2 routes, even with a higher cost.

Q: The initial cost on Agg1 was still set to 2000. Why is the reported cost 2100?

A: The E1 cost is the initial cost plus the path cost to reach the ASBR.

22. On Sw1, review the default interface costs.

Q: What is the interface cost to reach Agg1? (SVI 41)

A: The default interface cost is 100.

23. On Agg2, change the route-map entry for the 10.255.22.0/24 to external type 1 as well.

24. On Sw1, verify that the path to Agg2 is now the best path with a metric of 1100.

Test failover impact on the route metrics

25. On Sw1, verify the cost on the 2 destination subnets and repeat the output.

26. On Agg1, disable the uplink port 1/1/7.

27. On Sw1, the 10.255.21.0/24 should now have cost 2000.

28. On Agg1, enable the uplink 1/1/7 again.

29. On Agg2, disable the uplink port 1/1/7.

30. On Sw1, the 10.255.22.0/24 route should now have cost 2100.

31. On Agg2, enable the uplink 1/1/7 again.

32. On Sw1, verify that the original costs are shown again, stop the repeat.

Understanding Prefix Lists

In this section, some example routes are provided so you can apply some prefix-list configuration to them. All these example routes are under the 10.255.64.0/23 prefix. There are some host routes and some subnet routes in the list:

```
ip route 10.255.64.0/26 10.255.101.11
ip route 10.255.64.64/26 10.255.101.11
ip route 10.255.64.128/26 10.255.101.11
ip route 10.255.64.192/26 10.255.101.11
ip route 10.255.64.1/32 10.255.101.11
ip route 10.255.64.2/32 10.255.101.11
ip route 10.255.65.0/26 10.255.101.11
ip route 10.255.65.64/26 10.255.101.11
ip route 10.255.65.128/26 10.255.101.11
ip route 10.255.65.192/26 10.255.101.11
ip route 10.255.65.1/32 10.255.101.11
ip route 10.255.65.2/32 10.255.101.11
ip route 10.255.64.0/23 10.255.101.11
```

33. Use PC1 to open an SSH connection to Agg1

34. Open 'lab 07.1 - agg1 prefixlist examples.txt' and paste it into the Agg1 configuration.

On Agg1, make changes to the existing route-map to meet these scenario's. After each scenario, verify on Sw1.

35. Scenario1: Only the exact prefix 10.255.64.0/23 should appear in the Sw1 routing table.

36. Scenario2: All the routes under the 10.255.64.0/23 prefix should appear in the Sw1 routing table.

37. Scenario3: All the routes under the 10.255.64.0/23 prefix, except for the summary route itself (10.255.64.0/23) should appear in the Sw1 routing table.

38. Scenario4: All the routes under the 10.255.64.0/23 prefix, except for the summary route itself (10.255.64.0/23) and except for host routes (/32) should appear in the Sw1 routing table.

This concludes the static route redistribution task.

Task 3: Route Redistribution between OSPF and BGP

Introduction

In the previous task, you only configured your own side, since the partner network was already configured with static routes pointing to your network.

In this task, you will configure mutual redistribution between the partner network and your network using eBGP. Both your network and the partner network are using OSPF internally, and the 2 networks will be interconnected using eBGP.

The learned eBGP routes will need to be redistributed into the internal OSPF network.

This task will walk you through some steps to complete the mutual redistribution.

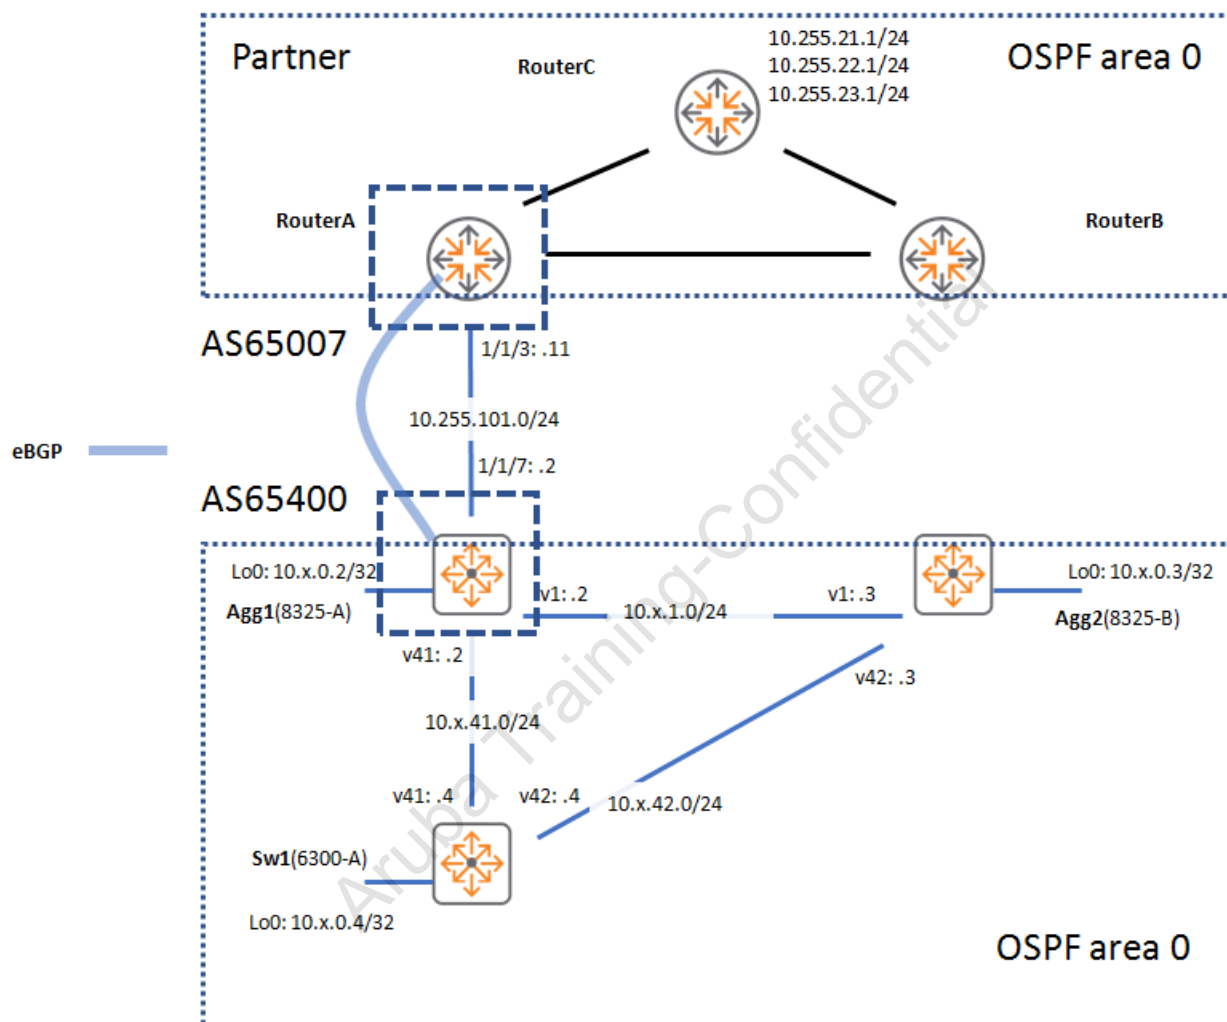
These are the high-level steps you will perform:

- Prepare the setup
- Configure eBGP peering
- Redistribute OSPF routes into BGP
- Redistribute BGP routes into OSPF
- Verify the operation

Objectives

- Redistribute routes between dynamic routing protocols

Lab Diagram



Steps

Prepare the setup

Prepare the routerA, routerB and routerC configurations for this lab.

You will first load the required start configuration so they can simulate a partner network.

1. On routerA, routerB and routerC, load the 'asts-lab07-3-rr' checkpoint to the startup configuration and reboot the system.

```
copy check asts-lab07-3-rr run
write mem
boot system
```

NOTE: The reboot is requested here since the virtualCX may not always load the interfaces after a checkpoint swap. The reboot ensures that the configuration is actually applied.

NOTE: This initial configuration contains the BGP configuration of the partner network.

1. On Agg1, remove the static route 10.255.21.0/24 to the partner network.

NOTE: Only the 10.255.21.0/24 route is removed since this route will be advertised using BGP to the partner network. The 10.255.22.0/24 route will not be used in this task.

2. On Agg2, remove the static route 10.255.21.0/24 to the partner network.

Verify current routes and OSPF neighbors

3. On Agg1, verify that there is no 10.255.21.0/24 route in the routing table.

Configure eBGP peering between Agg1 and routerA

In this section you will configure the eBGP peering.

In the real world, the partner would configure its own routers. In this lab activity the RouterA and RouterB have been preconfigured for you.

4. On Agg1, configure routerA (10.255.101.11) as an eBGP peer for AS 65007.

5. On Agg1, verify the eBGP peering state is established.

6. On Agg1, review the BGP IPv4 routing table.

Q: Are there any routes in the BGP IPv4 routing table?

A: Yes, the partner network has been configured to advertise its networks into BGP. This configuration was in the loaded checkpoint on the routerA.

7. On Agg1, review the advertised routes to routerA neighbor

Q: Are there any routes advertised by Agg1 to routerA?

A: No, Agg1 has not been configured to advertise any routes to routerA.

Redistribute local OSPF routes into BGP

To make the complete route redistribution work, several exchanges must be done.

In 1 direction (local routes to partner) this means:

- Local OSPF into BGP
- Local eBGP advertises routes to partner
- partner eBGP into partner OSPF (This step is already configured on routerA)

In the other direction (partner routes to local) this means:

- partner OSPF into partner eBGP (This step is already configured on routerA)
- partner eBGP advertises routes to local network (Already configured on routerA)
- local eBGP routes into local OSPF

In this section, the OSPF routes will be injected into BGP on the local network.

8. On Agg1, redistribute OSPF routes into the BGP IPv4 address family.

9. On Agg1, review the BGP IPv4 routing table. You should see all the local OSPF routes in the BGP routing table.

Q: What routes are now visible in the BGP IPv4 routing table?

A: All routes from both local and partner network are visible in the BGP routing table.

10. On Agg1, review the routes that are advertised to the routerA BGP peer. This now includes the local OSPF routes that were redistributed into BGP.

Redistribute BGP routes into OSPF

In the next section you will see how a simple redistribute of the BGP routes into OSPF seems to work just fine. However, when the second link will be added it will cause trouble with route feedback.

In this section you redistribute all BGP routes into OSPF on the local network.

11. On Sw1, review the current 'external' routes in the OSPF LSDB and the IP routing table.

Q: Do you see the 10.255.21.0 network?

A: No, it only exists in the Agg1 BGP and IP routing table, but since Agg1 has not redistributed the BGP routes into OSPF yet, it will not be available to Sw1.

12. On Agg1, redistribute BGP into OSPF.

13. On Sw1, review the routing table and the OSPF LSDB for the 10.255.21.0/24 route.

14. On Sw1, attempt to ping and traceroute to the partner network 10.255.21.1

Q: Did the solution work?

A: Yes, this demonstrates a basic single link route redistribution.

Task 4: Route Redistribution with multiple links

Introduction

In the previous task, there was a single link with the partner and you configured route redistribution between dynamic routing protocols OSPF and BGP.

In this task, you will explore adding a second link between the customer and the partner network.

Initially, you will see that you may break the redistribution with a routing loop due to route feedback. This will be resolved using route tags.

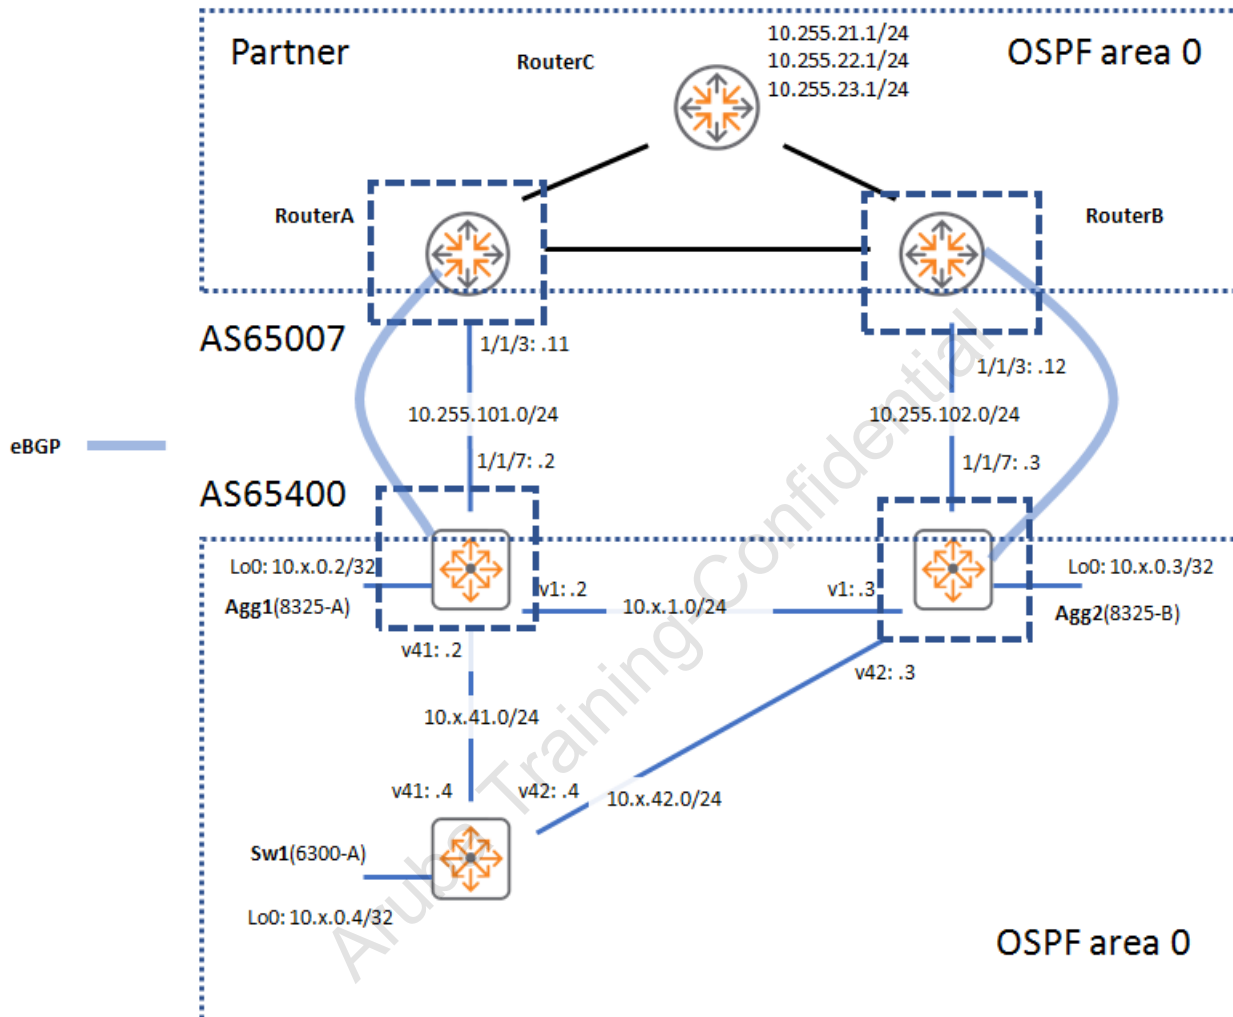
These are the high-level steps you will perform:

- Add second link - Attempt 1: Explore route feedback
- Add second link - Attempt 2: Use route tag to control route feedback

Objectives

- Configure a second link with an external network.
- Configure route-maps with route tags to identify routes.

Lab Diagram



Configure second link: Attempt 1

In this section you will see that the current configuration will reach its limits when a second link is added. First you will configure Agg2 with eBGP peering to routerB. Next the same configuration will be applied to Agg2 as you did in the previous section. You will then explore the route feedback.

NOTE: BGP is only used to learn routes from the partner network and distribute them using OSPF. There is no 'full' BGP routing topology with multiple AS systems, this is why there is no iBGP between Agg1 and Agg2.

1. On Agg2, configure the eBGP peering with routerB.
2. Verify the peering is established and that Agg2 has received BGP routes from routerB.

Review routing before the second link redistribution is configured

3. On routerA, review the 10.255.21.0 route. It is currently an OSPF route, using routerC as the nexthop IP.
4. On Agg1, list the routes that are advertised to routerA. You will review this list again in a few moments.
5. On Agg2, review the OSPF LSDB for external routes.

Q: How many external routes are there for the 10.255.21.0 network, and what is the Advertising Router?

A: Currently 1, advertised by 10.x.0.2 (agg1).

Agg2: Redistribute BGP into OSPF

On Agg2, you will now redistribute the BGP routes into OSPF and review the effect.

6. On Agg2, redistribute BGP into OSPF

Review routing after the redistribution

7. On Agg2, review the OSPF LSDB.
8. On Agg1, review the eBGP advertised routes to routerA

Q: What do you observe for the 10.255.21.0/24 prefix?

A: Since Agg2 redistributes the received BGP routes into OSPF, Agg1 will now learn the 10.255.21.0/24 prefix as an OSPF route and advertise it via BGP. This means that a route from the partner network is being advertised back to the partner network. This is also referred to as route feedback.

9. On routerA, review the routing table for 10.255.21.0

Q: Why is routerA using the BGP route instead of the internal OSPF route?

A: eBGP has a better administrative distance (20) than OSPF (110). The current situation results in Agg1 sending traffic to routerA, and routerA attempts to send it back to Agg1.

10. On Sw1, attempt to ping 10.255.21.1

Q: What would be the solution?

A: Agg2 should mark the routes from the partner network with a route tag in OSPF, so Agg1 can recognize the difference between the internal OSPF routes and the redistributed routes from the partner network.

Configure Second Link: Attempt 2 using route tags

In this section you will reconfigure the route redistribution using route tags to prevent the route feedback.

11. On Agg2, configure a route-map to set a tag on the matching routes. The tag is just a label, so any number could be used. In this case, since the partner is using BGP AS 65007, the tag 65007 will be used for these routes. Next apply this route-map to the redistributed BGP routes under the OSPF process.
12. On Agg1, configure a route-map that denies any routes with the route tag 65007. This will exclude the routes from the partner network, since these are marked with tag 65007. All other routes (the original OSPF routes) will still be accepted. Apply this route-map to the redistributed OSPF routes under the BGP context.
13. On Agg1, clear the BGP sessions to apply the changes
14. On Agg1, review the advertised routes to routerA. Compare this with the previous output of the same command you ran earlier.

Q: Is Agg1 still advertising the 10.255.21.0/24 prefix to routerA?

A: No, thanks to the use of a route tag, the route feedback has been prevented.

Customer Redistribution on Agg2

While the issue of the route feedback was already visible when Agg2 redistributed BGP routes into OSPF, the redistribution configuration was still incomplete.

On Agg2, only the BGP routes are currently redistributed into OSPF, while Agg2 is not redistributing its own OSPF routes into BGP.

This means that the partner network is only learning about the customer network via the Agg1 - routerA link at the moment.

However, Agg2 should not advertise all OSPF routes, since that would result in the route feedback again.

Configure the solution so Agg1 marks the redistributed BGP routes with tag 65007, and Agg2 denies these routes to be redistributed into BGP, while the other OSPF routes are still advertised.

15. On Agg2, verify that it is currently not advertising any routes over BGP.
16. On Agg2, verify that a simple redistribution would cause the route feedback

17. On Agg2, review the advertised routes, this currently includes 10.255.21.0/24. That will be prevented in the next steps.

Now prepare the filtering using route tags.

18. On Agg1, mark the routes when they are redistributed into OSPF
19. On Agg2, configure a route-map to deny the routes with tag 65007 and permit all other routes. Apply this to the OSPF redistribution into BGP.
20. On Agg2, clear the BGP session to apply the route-map changes.
21. On Agg2, review the advertised routes to routerB.

Q: Is the 10.255.21.0/24 prefix still advertised to the partner network?

A: No, thanks to the route tag configuration, the route feedback has been prevented.

22. On Sw1, review the IP routing table for the 10.255.21.0 entry.

Q: What is the tag listed in the output of this OSPF route?

A: The route is listed with tag 65007. When an OSPF route is marked with a tag, this information is included in the show ip route command output since the 10.05.0010 release.

Partner network configuration

You have now configured the customer side of the route redistribution. Keep in mind that the partner should apply the same configuration to prevent route feedback in the other direction.

In this lab, the partner configuration has been applied already. Feel free to explore the configuration on routerA and routerB.

23. routerA

24. routerB

25. Optionally, save a new checkpoint on all 4 devices 'asts-lab07-<your name>'.

This concludes the lab activity.

Lab 08: VRF and Route Leaking

In this lab you will configure VRFs and configure route leaking between different VRFs.

Objectives

- Configure static route leaking
- Configure route leaking using BGP

Requirements

This lab requires completion of Lab 07.

Scenario

In this demonstration, the customer expects some IOT devices in a branch location (Sw2 VRF 'branch'). The IOT sensor devices will be assigned to VLAN 62, there is an IOT management tool that will be installed in VLAN61. The VLANs 61 and 62 require inter-VLAN routing.

Since the customer does not trust the IOT devices, they want the IOT setup (SVI 61 and SVI62) in a dedicated VRF 'iot' on Sw2.

On Sw2, allow the VLANs 15,61 and 62 on the VLAN trunk 1/1/27 so that the SVI interfaces are up.

Static Route Leaking

1. On Sw2, show the customer that the VRF 'branch' SVI15 is able to communicate with the VRF 'iot' SVI61. There are no IOT devices available yet in the deployment, so it is sufficient if you can demonstrate a ping between Sw2 local SVI 15 (vrf 'branch') and SVI 61 (vrf 'iot') interfaces.
2. Next the customer wants to see that the SVI 61 of the VRF 'iot' can communicate with the SVI12 on Agg1, without making any configuration change on Agg1.
 - a. This should be demonstrated using a ping from Agg1 using source SVI12 to 10.x.61.1.
3. Once this works, demonstrate that a ping from Agg1 using source SVI 1 to 10.x.61.1 fails. The customer would like to see a control plane network trace on Sw2 that shows that the ICMP request is received in the VRF 'iot', but it cannot send a reply since the network is unreachable (from the vrf 'iot' point of view).

Dynamic Route Leaking using MP-BGP

The customer is happy the static route leaking worked well. However, they expect to have a lot of changes in the IOT routes and VLANs, so they want to ensure that routes can also be dynamically exchanged between VRFs.

1. On Sw2, remove the statically leaked routes for 10.x.61.0/24 and 10.x.12.0/24 from the vrf 'branch' and 'iot'.
2. Show MP-BGP with route distinguisher 64500:1 for VRF 'branch' and 64500:2 for VRF 'iot'.
3. Configure MP-BGP so the only routes matching a prefix-list 'branch_leak' from the VRF 'branch' are redistributed into the VRF 'iot'. The only route that should match the 'branch_leak' prefix-list is 10.x.12.0/24.
4. Configure MP-BGP so the only routes matching the prefix-list 'iot_leak' from the VRF 'iot' are redistributed into the VRF 'branch'. The only route that should match the prefix-list 'iot_leak' is 10.x.61.0/24.
5. The customer wants the routes of the VRF 'branch' marked with RT 64500:1, while the routes of VRF 'iot' should be marked with RT 64500:2.
 - a. There should be no configuration changes on Agg1.
6. Verify that a ping from Agg1 using source SVI12 can reach the 10.x.61.0/24 address in the VRF 'iot'.

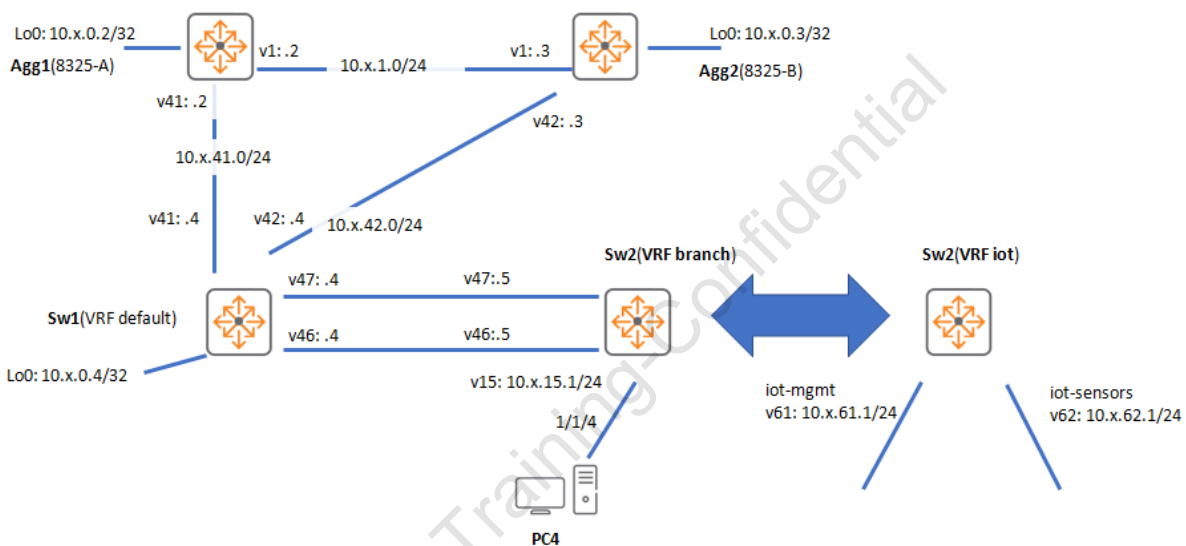
Route-map based import

The customer is pleased that the MP-BGP works fine. They just want to have more control over the route target import.

1. Update the VRF 'iot' so it does not import routes with RT 64500:1.

2. Now have VRF 'iot' import MP-BGP routes that have been marked with RT 64500:101.
3. Use a route-map to mark the 'branch_leak' prefix-list routes with RT 64500:101, instead of marking the RT at the VRF context level.
4. Verify that a ping from Agg1 using source SVI12 to 10.x.61.1/24 in the VRF 'iot' is successful.

Lab Diagram



Task 1: Prepare the base configuration

Objectives

- Prepare the VRF 'iot' on Sw2

Steps

IOT VRF Configuration

The route leaking will be performed on Sw2. Sw2 is hosting the 'branch' VRF. Now the customer is deploying IOT solutions in the branch, so they plan several VLANs for the IOT devices and sensors. They want these IOT VLANs to be completely isolated from the rest of the network, so they will deploy them in a VRF 'iot'. There is an IOT-Management subnet (10.x.61.0/24) that contains the IOT management systems. This subnet should be reachable for some subnets of the branch and main location.

This reachability will be configured using route leaking.

You will first prepare the start configuration on Sw2 to simulate the IOT network.

1. On Sw2, define the VRF 'iot'
2. On Sw2, define VLANs 61 and 62. To ensure they are 'up', they will be tagged on the port 1/1/27.

```
vlan 61
vlan 62

# vlans need to be 'up'
interface 1/1/27
  vlan trunk allowed 61,62

int vlan 61
  description iot-mgmt
  vrf attach iot
  ip address 10.x.61.1/24

int vlan 62
  description iot-sensors
  vrf attach iot
  ip address 10.x.62.1/24
int 1/1/4
  vlan access 15
```

3. On Sw2, force VLAN 15 'up' by adding it to the 1/1/27 VLAN trunk

Task 2: Explore the Environment

Objectives

- Familiarize yourself with the environment

Steps

Connectivity Checks

1. On Sw2, verify the OSPF neighbors in VRF branch. Sw1 should be a neighbor.
2. On Sw2, verify the VRF branch routing table. It is connected using SVI 46 and SVI 47 to Sw1 and should have learned the internal routes using OSPF.

Verify VRF Route Isolation

3. On Sw2, review the routing table for VRF iot.

Q: Do you see any IP addresses other than 10.x.61.0/24 and 10.x.62.0/24 ?

A: No, only IP subnets from the VRF iot are currently available.

4. Verify intra-VRF communication: On Sw2, in the VRF iot, attempt to ping 10.x.62.1 using 10.x.61.1 as the source IP address. This ping should be successful.

You will now verify that there is no communication from Sw2 VRF branch to an IP address of VRF iot.

5. On Sw2, in the VRF branch, attempt to ping to the Sw2 IP address in the VRF iot (10.x.61.1). This ping should fail and demonstrates that the VRF offers isolation.

Task 3: Static Route Leaking

Objectives

- Manually configure a route leak between 2 VRFs
- Verify the route leaking

Steps

Manually Leak Directly Connected Interface

In this section you will enable communication between directly connected interfaces on Sw2.

In the VRF `iot`, Sw2 has a directly connected link `10.x.61.0/24`

In the VRF `branch`, Sw2 has a directly connected link `10.x.15.0/24`.

You will now leak these 2 routes.

1. On Sw2, add a route to VRF `branch` for the `10.x.61.0/24` subnet. Do not use a next-hop IP, just use the original SVI interface of the source VRF as the destination interface.

NOTE: AOS-CX will automatically notice that the configured destination interface belongs to another VRF.

2. Review the IP routing table in VRF `branch`.

Q: What is the destination interface (via...) of the `10.x.61.0/24` subnet?

A: AOS-CX will report `vlan61(vrf iot)` as the destination interface.

Q: Would a ping work between the Sw2 VRF `branch` and VRF `iot` at this point?

A: No, since that requires bi-directional (return) routes.

3. On Sw2, add a static route to VRF `iot` for the `10.x.15.0/24` subnet. Do not use a next-hop IP, use the original SVI as the destination interface.
4. On Sw2, review the IP routing table of VRF `iot`. Check the `10.x.15.0/24` route.
5. On Sw2, add a static route to VRF `iot` for the `10.x.12.0/24` subnet. Do not use a next-hop IP, use the original egress SVI as the destination interface.
6. Review the IP routing table of VRF `iot`. Check the `10.x.12.0/24` route and the destination.
7. Local test1: On Sw2, ping `10.x.61.1` from the VRF `branch` using source SVI `15`.
8. Local test2: On Sw2, ping `10.x.15.1` from the VRF `iot` using source SVI `61`.
9. Both test should be successful.

10. Verify that only the 10.x.61.0/24 network is connected. On Sw2, attempt to ping using the 10.x.62.1 source IP address.

Make the leaked route available inside the VRF branch

Verify that the rest of the network does not have access to the leaked route yet.

11. From Agg1, attempt to ping 10.x.61.1 using source SVI 12.

Q: Did the ping succeed? If not, why?

A: No, check the routing table on Agg1 for route 10.x.61.0.

12. Configure Sw2 so the leaked route is distributed in OSPF in the branch VRF.

IMPORTANT: In AOS-CX, an OSPF process ID is unique within the VRF. Make sure to include the correct VRF name in the command. So 'router ospf 2' and 'router ospf 2 vrf branch' are two different OSPF processes.

13. On Agg1, review the routing table, the route 10.x.61.0 should now be in the routing table.

14. On Agg1, Attempt to ping 10.x.61.1 using source SVI 12.

Q: Did this work?

A: Yes, this should work.

15. Attempt to ping 10.x.61.1 using source SVI 1.

Q: Did this work?

A: No, since this source subnet is not known in the destination VRF iot.

Debug on Sw2

16. On Sw2, start a TCPDUMP of protocol 1 in the VRF iot.

17. On Agg1, do another attempt to ping 10.x.61.1 using source SVI 1.

18. On Sw2, review the TCPDUMP output. You should notice the incoming ICMP first (source 10.x.1.2), next an 'unreachable' message from localhost to localhost. This confirms that the packet reached the iot VRF, but there is no return route.

19. Stop the TCPDUMP on Sw2 using CTRL-C

Prepare for MP-BGP leaking

In the next section, the route leaking will be done using MP-BGP routing protocol, so the statically leaked routes must be removed.

20. On Sw2, remove the static route leaks from both VRF branch (10.x.12.0/24) and iot (10.x.61.0/24).

NOTE: The static route leak for 10.x.15.0/24 does not need to be removed from the configuration.

Task 4: MP-BGP

Objectives

- Configure BGP to perform route-leaking
- Configure route-distinguisher
- Configure and use route-targets
- Verify the BGP routing table
- Verify the route leaking

In the previous task you leaked a directly connected route using a static configuration. In this task, you will leak the same routes dynamically using BGP.

When multiple VRFs will be imported into BGP, it is important that the routes of each VRF receive a unique distinguisher to avoid a route collision on the BGP routing table. This also ensures that VRFs could have the same IP addresses without any issues in the BGP tables.

This distinguisher should be applied to each VRF.

For the Route Distinguisher, typically the main AS number is used, followed by a VRF identifier. These 2 numbers can be separated by a colon sign.

In this lab, VRF branch will be assigned VRF id 1, VRF iot will be assigned VRF id 2.

Steps

Configure Route Distinguisher and import routes into BGP routing table

1. On Sw2, configure VRF branch with route-distinguisher 64500:1
2. Configure VRF iot with route-distinguisher 64500:2
3. You want import the 10.x.12.0/24 subnet into the BGP tables, but you don't want to import any other subnets. Think about a strategy to apply this. Some questions in the process:

Q: Can you use prefix-lists to identify the required subnet? What is the source routing protocol of this route?

A: Yes, you can use a prefix list to identify the 10.x.12.0/24 route. The source routing protocol for this route is OSPF.

4. Prepare a prefix-list 'branch_leak' on Sw2 to identify the source route (10.x.12.0/24).
5. Configure a route-map 'bgp_branch_from_ospf' on Sw2 that permits the 'branch_leak' routes.
6. On Sw2, configure BGP with AS 64500 and enter the vrf branch context.

7. Redistribute the OSPF routes using the configured route-map

8. Review the BGP IPv4 routing table on Sw2

Q: Do you see any routes?

A: No, since the default command only shows the 'default' VRF routing table.

9. Review the BGP VRF blue IPv4 routing table.

Q: Do you see any routes?

A: Yes, the 10.x.121.0/24 route should be in the list. It should be the only route.

Q: What is the Route Distinguisher for this route?

A: The 10.x.12.0/24 route should be under the 64500:1 RD.

10. As a convenience, AOS-CX also includes a command to see all VRFs and all address-families. Review all the records of the BGP routing table

11. Now apply the same procedure for VRF iot. Remember that in the VRF iot, the route 10.x.61.0/24 is a connected route, not an OSPF learned route.

- Define prefix-list 'iot_leak' with the 10.x.61.0/24 subnet
- Define a routemap 'bgp_iot_from_connected' to match the prefix-list.
- Configure the BGP VRF iot context to redistribute connected routes based on the route-map

12. Verify the BGP VRF iot tables.

Q: What routes do you observe?

A: Only the 10.x.61.0/24 route should be listed.

Q: What is the RD for this route?

A: The RD should be 64500:2

13. Review the BGP configuration of both VRFs.

You have now configured a BGP database that contains 'interesting' routes. In the next section you will see how you can use these routes for the route leaking process.

Route Targets

Since the route exchange may need to be done between multiple VRFs and since the same routes may need to be exchanged based on similar rules, a concept had to be introduced to 'mark' routes.

These 'marks' are known as Route Targets (RT). Routes in the BGP routing table can receive multiple RT marks. Once these marks are in place, the administrator can choose which of these marks will be used to import the routes into the other VRFs. In the BGP routing table, these marks are saved as extended community attributes.

Route-targets can be assigned using a very simple method at the VRF level, or using a more selective method based on a route-map.

You will first use the VRF based method, then the route-map based method will be applied.

Simple method

14. On Sw2, let's review the current RT assignments on the routes (10.x.12.0/24) using the command

Q: Are there any ext-communities for this route?

A: No, currently, no route-targets have been configured.

15. In the Sw2's VRF branch context, access the IPv4 address-family and set RT 64500:1 as the export RT.

16. Review the BGP route details of the 10.x.12.0/24 route again.

Q: Are there any ext-communities for this route?

A: Yes, RT 64500:1 can now be observed on the route.

Q: Was this a selective process?

A: No, any route that is learned from the VRF branch will be assigned with this route-target. This may be fine and easy, but for more control a route-map should be used.

The next question is: what can be done with this RT? Well, this RT can be used when a VRF wants to import routes from the BGP tables.

17. On Sw2, review the VRF iot routing table.

Q: Do you currently see a route for 10.x.12.0/24?

A: No, the only leaked route should be the 10.x.15.0/24 route from the previous task.

18. On Sw2, enter the VRF iot context, access the IPv4 address-family and configure the route-target import with 64500:1

19. Review the VRF iot routing table.

Q: Do you see a route for the 10.x.12.0/24 network?

A: Yes, the route should now be imported from the BGP routing table.

Q: What is the source of the protocol?

A: BGP is listed as the routing protocol.

Q: Do you see the destination VRF in the 'via' field?

A: Yes, it shows 'vrf branch'.

20. You have now managed to make the required route from VRF branch available in the VRF iot. Now apply the same logic in the other direction to make the 10.x.61.0/24 route from VRF iot available in the VRF branch routing table. Use RT 64500:2 for the iot routes.

21. Verify by checking BGP VRF iot routes and RT

22. Review the route 10.x.61.0 in the VRF branch IP routing table, check the 'via' field of the route.

Verify connectivity

You have now succeeded to leak the routes on the Sw2 device.

23. On Agg1, attempt to ping the 10.x.61.1 IP address (this destination address is in the VRF iot), using source SVI 12.

Q: Did this work? Why?

A: It did not work. The leaked route is not available inside the VRF branch routing domain

24. On Sw2, make sure that the leaked route is made available in the OSPF routing protocol in the VRF branch.

25. Verify: On Agg1, review the IP routing table. Make sure the 10.x.61.0/24 route is listed.

26. On Agg1, attempt to ping the 10.x.61.1 IP address using source SVI 12.

Q: Did this succeed?

A: Yes, the with the completed route leaking configuration, the ping should succeed.

Route-map based RT

In the previous section you have seen how the RT was set at the VRF level, so it applied to all the routes of that VRF.

In case more selective marking is required, a route-map can be used.

In this example, the default 'export' is left in place, but it could be removed as well if required.

You will first remove the 'import' in the VRF iot and replace it with another import that will only be applied by a route-map.

27. on Sw2, enter the VRF iot IPv4 address-family and remove the current import RT (Use 'show run vrf' to see the current configuration)
28. Add a new import RT for 64500:101. There are currently no routes in the BGP table with this RT. They will be marked in the next steps.
29. Review the VRF configuration on Sw2
30. Review the Sw2's VRF iot routing table, verify that the route to 10.x.12.0/24 has disappeared.

Now you will use a route-map to mark the routes of VRF branch.

31. Review the previously created route-map that was used to select OSPF routes for BGP. The suggested name was 'bgp_branch_from_ospf'.
32. Modify the sequence that matches the 10.x.12.0/24 (branch_leak) prefix-list by adding a 'set' action. Set the extended community rt to "64500:101".

NOTE: Make sure to use the quotes (" ") when entering the RT.

NOTE: You can apply multiple Route Targets using the route map. You can use for example "64500:101 64500:201".

33. Verify the route-map configuration.
34. Review the details of the BGP VRF branch 10.x.12.0/24 entry.

Q: What are the RT marks on the route now?

A: On top of the 'default' RT that was assigned by the VRF, the route-map has added RT 64500:101 and 64500:1. to the route. Any other VRF can now use these RT to import the routes.

35. Review the VRF iot IP routing table.

Q: Is the 10.x.12.0/24 route available now?

A: Yes, since VRF iot imports RT 64500:101 and since the route has received this RT, the route will be available in the routing table.

36. Verify: On Agg1, you should now be able to reach 10.x.61.1 using source SVI 12

37. Optionally, save a new checkpoint on all 4 devices 'asts-lab08-<your name>'.

This concludes the lab activity.

Lab 09: Multicast

In this lab you will configure and troubleshoot an IP Multicast routed topology.

Objectives

- Configure Multicast Routing
- Configure VSX Active - Active PIM
- Configure and troubleshoot IGMP snooping and querier functions
- Configure IGMP ACL filtering

Requirements

This lab requires completion of Lab 08.

Introduction

To provide you with a larger multicast routing environment, in this lab topology you will be configuring Multicast Routing on the default VRF and a 'Building2' VRF (b2).

1. The Agg1 and Agg2 switches are configured as a VSX pair, and in the default VRF they provide routing for the VLAN 11 and 12 endpoint subnets.
2. Sw1 will act as a router that interconnects the 'default' VRF building and the 'b2' VRF.
3. Sw1 has a routed link SVI 41 to Agg1(default VRF) and SVI 42 to Agg2(default VRF).
4. For building 2, Sw1 has SVI 241 to Agg1(b2 VRF) and SVI 242 to Agg2(b2 VRF).
5. Sw1 has only the default VRF and simply routes between SVI 41, SVI 42, SVI 241 and SVI 242.
6. OSPF is enabled on all routers and VRFs. The environment works as a single routing domain.
7. Sw2 has the combined role of the Layer2 Access switch for both buildings.
8. PC4 will have a role in Building2 and is assigned to VLAN 212. It's default gateway is on the VSX in the VRF 'b2'.
9. PC3 will have a role in the first building. It is physically connected to Sw1 port 1/1/3 (Sw1 is limited to a routed role in this lab), in VLAN 11, and VLAN 11 connects to port 1/1/27 on Sw2.
 - This will make it appear as if PC3 is connected to Sw2 port 1/1/27.

Scenario

Setup

On PC1, open Desktop > ASTS and run **ASTS-Lab09-multicast-initial-setup.cmd** to prepare the multicast setup. This will add a new vrf 'b2' (for building2) on Agg1 and Agg2 and routed links to the Sw1.

Verify the following:

1. Sw1 has an OSPF adjacency with 10.x.200.2 (VRF 'b2' on Agg1) and 10.x.200.3 (VRF 'b2' on Agg2).
2. PC3 has an IP address in VLAN 11
3. PC4 has an IP address in VLAN 212.
4. Perform a traceroute from PC3 to PC4, there should be 3 transit hops.

Multicast routing

The customer wants a demonstration of multicast routing in their environment with the following parameters:

1. IP multicast traffic exchange must be possible on any path between PC3 and PC4.
2. All the RPs will be registered automatically on Sw1.
3. Registration is successful even when any L3 link fails on Sw1.
4. Initially Sw1 acts as the only RP to handle multicast traffic.
5. Ensure that both PC3 and PC4 can register for multicast traffic.

When reviewing the configuration, the customer noticed a strange status of the IGMP querier on SVI11 when comparing the output on Agg1 and Agg2. They want you to investigate and adjust the configuration.

6. The customer is worried that multicast traffic will be Layer2 flooded in the VLANs 11 and 212, your configuration should prevent this flooding.
7. The multicast traffic can be tested on PC4 by running '**ASTS-lab09-PC4-mcast-tx.cmd**'.
 - a. Transmit multicast traffic to 239.1.212.1.
 - b. On PC3, run '**ASTS-lab09-PC3-mcast-rx.cmd**' to start the multicast receiver and enter 239.1.212.1 to join the multicast.
8. The customer wants you to demonstrate:
 - a. The number of multicast flows registered at RP Sw1.
 - b. The number of Register and Register Stop messages on the RP and explain the purpose of these messages.
 - c. The exact active reverse path, hop by hop, of this multicast flow from PC3 to PC4, including the status on the Layer2 switched VLANs.
 - d. The active path of the multicast by entering a command on Agg2 vrf 'b2' context that shows the 239.1.212.1 multicast using 10.x.11.3 as the last hop router.

Now that multicast routing is working, they are worried that the centralized RP may become overloaded at some point in the future.

1. Adjust the configuration so that Agg1/Agg2 (vrf 'b2') would become the RP for 239.1.212.0/24 range of multicast addresses.

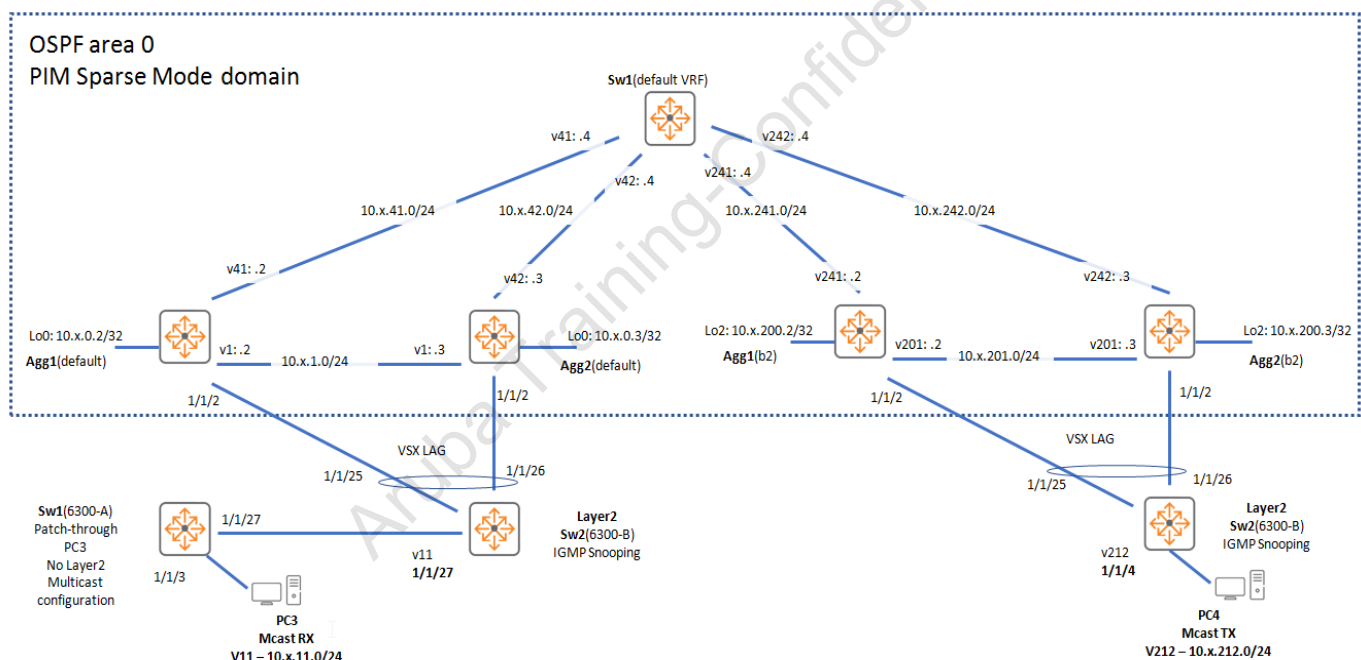
2. Demonstrate the election process and that there are no more REGISTER messages received on Sw1 when PC4 starts a new multicast stream.

VSX node reboot: The customer tests a VSX node reboot and notices it takes some time before the multicast traffic is resumed. Correct the configuration and demonstrate that a VSX node reboot would have minimal impact (less than a second) on the multicast traffic.

Filtering: The customer wants to ensure that the clients connected to SVI11 cannot join multicasts in the 239.1.2.0/24 range, since this range will be used by some security cameras. They expect you to demonstrate the filtering and show the number of dropped join packets on SVI11.

Clients connected to VLAN11 should not be able to join 239.1.3.0/24, even if the source is connected in the same VLAN11.

Lab Diagram



Task 1: Load the start configurations

Objectives

- Load the start configuration for Agg1, Agg2, Sw1 and Sw2.

Steps

Introduction

You will first load the required start configuration.

1. On PC1, open Desktop > ASTS folder and run **ASTS-Lab09-multicast-initial-setup.cmd**

These commands will:

- Configure Agg1 and Agg2 with a routed link to Sw1
- Configure VRF 'b2' on Agg1 and Agg2

NOTE: If you prefer to manually apply the settings, navigate to the PC1 > Desktop > ASTS > config folder. You can open the files `asts_cfg_lab09-init_<device>.txt` and apply the content to the respective devices.

Explore the environment

2. On Sw1, review the OSPF neighbors. There should be 2 new OSPF neighbors: 10.x.200.2 and 10.x.200.3
3. On Agg1, review the IP routing table of the VRF default. All routes should be visible, including the routes of the 'building2', since they are exchanged via OSPF by the Sw1 router. Look for the 10.x.212.0/24 route.
4. On PC3, verify the IP address. It should have an IP address in the 10.x.11.0/24 subnet.
5. On PC4, renew the IP address. It should receive an IP address in the 10.x.212.0/24 subnet. Take note of the IP address.

NOTE: .In the rest of this lab instructions and output, the PC4 IP address is listed as 10.x.212.z (z for the host IP)

6. On PC3, perform a `tracert -d` to the PC4 IP address (in 10.x.212.0/24 subnet). The trace should show:

Example on PC3 (this output shows PC4 as IP 10.x.212.z):

- Agg1 or Agg2 in VRF default
- Sw1 (10.x.41.4 or 10.x.42.4)
- Agg1 or Agg2 in VRF b2
- PC4

7. This means you have a fully routed topology with 5 logical routers:

Agg1 (default) / Agg2 (default) / Sw1 / Agg1 (b2) / Agg2 (b2).

In the next task, remember that these are the logical routers. When you need to complete a task on all logical routers, it means you need to configure:

- Agg1 (VRF default)
- Agg2 (VRF default)
- Sw1
- Agg1 (VRF b2)
- Agg2 (VRF b2)

Task 2: Configure PIM SM

Objectives

- Configure BSR election
- Configure PIM SM Interfaces
- Configure single RP

Steps

Configure the C-BSR

1. Sw1 acts as the central router, configure it as Candidate BSR using the loopback 0 IP address and BSR priority 200.
2. Enable all Layer3 interfaces for multicast on Sw1.
3. Review the enabled PIM interfaces on Sw1

Enable PIM-SM on all routing devices

On the 4 other logical routers, enable PIM SM on all the Layer3 interfaces.

NOTE: Feel free to use the ASTS_cfg_lab09-steps.txt document from PC1 > Desktop > configs. It contains these configuration snippets so you can paste this into the devices.

4. On Agg1(default)
5. On Agg1, verify the enabled PIM interfaces
6. On Agg2(default), enable the interfaces and verify the enabled interfaces.
7. On Agg1(b2) , enable the interfaces and verify the enabled interfaces.
8. On Agg2(b2), enable the interfaces and verify the enabled interfaces.
9. On Sw1, verify you have 4 PIM neighbors.
10. On Agg1(default) review the elected BSR

Q: Which system is supposed to be the elected BSR?

A: Only Sw1 has been configured as the Candidate BSR, so Sw1 (10.x.0.4) should be the elected BSR.

Q: When is the next BSR message expected?

A: The system will show the maximum timer, the BSR will send an update message every 60 seconds.

11. On Agg1, you may also review the PIM status a specific VRF or for all the VRFs

Configure Sw1 as C-RP

12. On Agg1(default), review the current RP list.

Q: Is there any RP available in the current setup?

A: No, none of the systems has been configured as a Candidate RP, so there will be no Elected RP.

13. Configure Sw1 as a candidate RP with priority 210 using the loopback 0 IP.

14. On Agg1(default), review the RP list.

NOTE: The RP Set is updated as part of the BSR message, so it may take up to 60 seconds to see the latest list. You can use the 'repeat' function to wait for the BSR update.

15. Verify that Sw1 is listed as an RP.

Q: What is the default multicast group address and group mask that an RP will handle?

A: 224.0.0.0 240.0.0.0 (224.0.0.0/4)

16. On Sw1, review the multicast summary for PIM-SM.

Q: How many PIM SM Interfaces are enabled? How many PIM SM neighbors are there ?

A: Sw1 has 5 interfaces and 4 neighbors.

Task 3: Configure IGMP and Verify Multicast operation

Objectives

- Configure IGMP Querier
- Configure IGMP Snooping
- Use mtrace
- Perform a trace of the multicast traffic between PC3 and PC4 on Sw1

Steps

IGMP Querier

The IGMP Querier is a Layer3 function that will be enabled on the client connected subnets. The IGMP querier will query if any endpoints are registering to receive multicast streams. This will be used by the PIM protocol to request the multicast forwarding in the routed network.

1. On Agg1, enable the IGMP Querier on the endpoint-facing SVIs (SVI 11,212).
2. Repeat this configuration on Agg2.
3. On Agg1, verify when the IGMP Querier election has completed, repeat the command until the SVI has either the 'Querier' or the 'Non-Querier' state.
4. Review the IGMP querier state on Agg2 as well.

Q: Why did Agg1 win the IGMP querier election for SVI 212?

A: Agg1 interface IP address is lower than the Agg2.

Q: What do you observe for SVI 11?

A: In the IGMP output, you may see that the device is Non-Querier, but the Querier IP is still listed as [this switch].

Q: What is causing this behavior?

A: For SVI 11, you have configured the 'Virtual Active Gateway' function on VSX, so both Agg1 and Agg2 have the same SVI IP address. This was configured during the VSX Lab activity. The 'Virtual Active Gateway' (using a single IP for SVI and Active Gateway) is not supported with IGMP services, so you need to configure the normal Active Gateway on SVI11 (using a unique physical IP for each node and a 3th IP address for the Active Gateway function).

Configure SVI 11 as normal Active Gateway with unique SVI IP per node.

TIP: If you would have removed the SVI IP address instead of changing it, you need to apply the SVI OSPF configuration again.

5. On Agg1, review the current configuration of SVI 11 (so you can copy/paste the active gateway commands in the next step)
6. On Agg1, disable Active Gateway IP on the SVI 11 and apply the new IP address.
7. On Agg2, apply the new IP address on SVI 11
8. On Agg1, apply the active-gateway configuration
9. On Agg1, review the IGMP Querier state is now 'Querier' for both SVI 11 and SVI 242.
10. On Agg2, review the IGMP Querier state is now 'Non-Querier' for both SVI 11 and SVI 242.

Q: Is IGMP snooping enabled for these VLANs?

A: No, IGMP snooping is a different function from the IGMP Querier, it is not enabled by default.

IGMP Snooping

11. On Agg1, review the IP IGMP snooping default status.
12. On Agg1, enable IGMP Snooping on VLANs 11 and 212.

TIP: You can use a VLAN range command for the IGMP snooping configuration.

13. On Agg1, verify the IGMP snooping status again, you should see that VLAN 11 and 242 are now enabled for IGMP snooping.
14. On Agg1, review the IGMP Querier status, it should also show that IGMP Snooping is enabled for the VLAN:
15. On Agg2, enable IGMP Snooping on VLANs 11,212 and verify the status.
16. On Sw2, enable IGMP Snooping on VLANs 11,212 and verify the status.
17. Sw2 is acting as a Layer2 access switch. Review the IP IGMP snooping information. You should use the 'repeat' function until the IGMP Querier has been detected.

Q: What is the IGMP Querier IP that was discovered by Sw2?

A: Since the IGMP Querier sends out IGMP packets, the access switch Sw2 can detect that this is the Agg1 switch SVI address.

Test Multicast and verify status on Sw1

In this section, you will use PC4 to transmit multicast test packets to destination 239.1.212.1. This traffic will be detected by the closest router and forwarded to the RP, in this case Sw1.

Since there are no listeners at this moment, Sw1 will signal that the packets no longer need to be forwarded.

Test without multicast receiver

Review status before starting the multicast transmitter

18. On Sw1, review the number of Received (RX) 'Register' and Transmitted (Tx) 'Register Stop' messages.

NOTE: The Rx Counters are listed after the Tx Counters in the output.

NOTE: Even though you have not started a multicast, you may already see some register messages. This is due to the Windows client that will send multicasts to 239.255.255.250, this is done by SSDP service and part of the UPnP discovery. You can ignore these during the lab.

19. On Sw1, review the multicast summary, look for the number of multicast flows on this RP.

Q: How many multicast flows have been registered with this RP?

A: We would expect 0, but you could notice a flow due to a the Windows client SSDP multicast transmission. Just take note of the current number.

20. **On both Agg1 and Agg2**, review the active multicast routes in the VRF b2. The expected output is that there are no active flows at the moment.

21. On PC4, run the '**ASTS-lab09-PC4-mcast-tx.cmd**' file. Enter multicast IP 239.1.212.1 as the destination IP. This will generate 1 multicast packet per second and report status every 5 seconds.

Review status after starting the multicast transmitter.

22. **On both Agg1 and Agg2**, review the active multicast routes in the VRF b2 again.

Q: What is the current states?

A: The state of the 239.1.212.1 Group Address is 'bridge'.

Q: What does this mean?

A: The multicast is not requested by any upstream multicast router, so it is not forwarded over any multicast routed interfaces. This does confirm that the presence of the multicast was detected by the PIM router on the transmitting VLAN.

23. On Sw1, review the statistics of the PIM SM RP.

Q: How many flows have registered with this RP?

A: You should notice that there is 1 additional flow registered compared to the previous output.

24. On Sw1, review the Loopback0 PIM statistics

Q: What type of packet counter has increased on the Rx Counters?

A: The RP has received a PIM register message for the new flow.

Q: What is the response from the RP?

A: Since there are no listeners, Sw1 will return a 'register-stop', the Tx Counter for 'Register Stop' should have increased.

25. On Sw1, review the current multicast routing table.

Q: Do you expect any multicast route flows in the current state?

A: Since there are no receivers active, there should be no active multicast routes.

The multicast receiver is connected to SVI11, on Agg1(VRF default) and Agg2(VRF default). These devices connect using SVI 41 and SVI 42 to Sw1. Since the receiver is currently not active, no 'join' messages should have been sent at this point.

26. On Agg1, review the IP multicast routes for the default VRF. Repeat this for the VSX-peer. You may notice the 239.255.255.250 multicast (Windows SSDP), but the 239.1.212.1 multicast should not be in the list.

Test with multicast receiver

27. On PC4, verify that the Multicast tx is still running.

28. On PC3, run the '**ASTS- lab09-PC3-mcast-rx.cmd**' and enter the multicast address that PC4 is currently using as destination (239.1.212.1).

29. PC3 IPerf should report every 5 seconds about the received packets.

NOTE: If no traffic was received, review your PIM and IGMP configurations.

NOTE: The first report may show some lost packets, the next reports should show about 5 packets.

Review the status of the multicast route

In the next steps you will explore the status of the multicast flow on the devices.

You will use this order:

- Sw2 of the multicast receiver: review the IGMP Snooping.
- Agg1/Agg2 (default VRF) of the multicast receiver: IGMP Querier and IGMP Snooping.
- Agg1/Agg2 (default VRF) of the multicast receiver: Multicast routed/bridged tables.
- Sw1: Multicast routed tables.
- Agg1/Agg2 (b2 VRF): Multicast routed tables

Sw2 - IGMP Snooping

30. On Sw2, PC3 is connected to Layer2 VLAN 11. Review the detailed IGMP snooping for VLAN11.

Q: What do you observe?

A: Sw2 IGMP snooping has detected the multicast, it is currently filtered and only forwarded to port 1/1/27 (PC3 connected port).

Agg1 - Agg2 IGMP Querier

31. On Agg1, review the active groups for the IGMP Querier on SVI 11. You should observe that your multicast group was detected by the IGMP Querier

32. On Agg1, repeat the command for the VSX peer. Agg2 should also have detected the group.

Agg1 - Agg2 IGMP Snooping

While the IGMP Querier will trigger the PIM multicast routing to forward the multicast to this Layer3 interface, the IGMP Querier does not know which Layer2 interfaces in the current VLAN should receive the multicast. This is handled by the IGMP Snooping process.

33. On Agg1, review the IGMP snooping information for VLAN11 with the detailed information.

Q: What port will Agg1 forward the multicast on inside this VLAN11?

A: Since the multicast receiver is connected to Sw2, using LAG2, the 'lag2' interface should be listed here.

34. On Agg1, repeat the command for the VSX peer to confirm it has the same information.

Agg1 - Agg2 PIM Multicast Routing

Based on the IGMP Querier, the DR PIM Router on the SVI 11 will send an upstream request to forward the multicast stream.

First review which system is the active DR for PIM on the SVI 11.

35. On Agg1, check the PIM SVI 11 interface and repeat this for the VSX peer

Q: What do you observe?

A: The currently active PIM DR on SVI 11 is Agg2, based on the highest interface address.

36. On Agg1, now review the multicast forwarding table in the default VRF. Repeat this for the VSX peer.

Q: What do you observe for your multicast stream?

A: The PIM DR (Agg2) has a 'route' entry in the forwarding tables. It receives the multicast over SVI 42. The Agg1 has a 'bridge' entry in the forwarding tables, it receives the multicast over SVI 11.

Q: What does this mean?

A: You have not configured VSX active-active PIM at the moment, so only 1 VSX node (the one with the PIM DR role on the subnet) has requested the multicast stream on the uplink router Sw1. The other node (Agg1 in this example) will still detect the multicast on the receiver subnet as 'bridged' traffic.

37. On Agg1, review the mroute details by entering the group address and the source IP of PC4 (10.x.212.z > z is the current host IP of PC4). Repeat this for the VSX peer.

Q: What do you observe?

A: Only the PIM DR (Agg2) has a Downstream interface, VLAN11. Agg1 is not forwarding this stream.

Sw1 router

38. Now investigate the status on Sw1. Review the mroute tables.

Q: What is the incoming interface?

A: SVI 241 or SVI 242.

39. On Sw1, review the details of the route.

Q: What do you observe?

A: The multicast is forwarded over SVI42. This is in line with the collected information of Agg1/Agg2 (default VRF) in the previous steps.

Agg1/Agg2 (b2 VRF)

40. On Agg1, review the mroute tables of the 'b2' VRF. Repeat for the VSX peer.

Q: What do you observe?

A: 1 node has the multicast as 'bridge' state, while the other node has the multicast as 'route' state.

41. On Agg1, review the details of the route, using PC4 IP as the source IP. Repeat this for the VSX peer.

Q: What do you observe?

A: 1 node is routing the multicast to Sw1 over the SVI 241 or SVI 242 interface.

This concludes the step by step detail review of the multicast.

Test with mtrace

AOS-CX has a multicast trace tool that allows administrators to trace the current multicast routed path of a multicast in the network. It can be run locally and for remote diagnostics.

You will first explore the routed path from the multicast receiver Last Hop Router (LHR), this is Agg2 (default VRF) in this lab.

Local traces

42. On Agg2, run mtrace with your group and enter the transmitter (PC4) source IP address. This should reveal the same path as what you explored in the previous section.

43. On Agg1, repeat this command.

Q: What is the output?

A: For Agg1 (default VRF), there is no active route, the flow will be reported in 'bridge' state.

44. On Sw1, repeat this command, this should reveal the path the Agg2 (b2 VRF)

45. On Agg1 and Agg2, repeat the command in the 'b2' VRF context. 1 node will reveal that the source is directly connected, the other node will report that the flow is in bridged state.

Remote traces

It is also possible to use mtrace remote. You will now use Agg2 (b2 VRF) to achieve the same information as the previous steps.

46. On Agg2, run the mtrace in the VRF b2, but specify the 'Last Hop Router' in the command. Use the Agg(default VRF) as the LHR

Q: What do you observe?

A: The mtrace will send the request to the 10.x.11.3 LHR, this router will perform the trace and the results are reported back. The output is the same as if it would run locally on the LHR.

47. On Agg2, run the mtrace with the Sw1 as the LHR. The output should match the output from Sw1 in the previous steps.

Optional: PIM Diagnostics Dump

In case detailed information is required on a multicast router, it is possible to perform a diagnostic dump of the PIM tables.

The output can be reviewed on the switch locally or it can be saved and uploaded to an external host for review.

48. On Sw1, perform a diagnostic dump locally.

NOTE: This will produce a lot of output. In the next step, an upload to the TFTP server will be used.

49. On Sw1, perform a diagnostic dump to a local file and upload the file to the TFTP server (10.251.x.91).

NOTE: You may need to verify that the TFTP server is running on PC1.

50. On the PC1, in the TFTP folder, open the pim-sw1.txt file. Look for 'piminfo vrf default rp'. This will list the currently registered groups with this RP.

Task 4: Configure distributed RP

Objectives

- Configure regional RPs
- Verify RP set
- Verify RP operation with network trace on Sw1

Steps

Configure regional RP

1. Configure Agg1(b2) as C-RP for the 239.1.212.0/24 multicast range.
2. Repeat this for Agg2(b2)
3. Verify on Sw1 that the RP-SET was updated. Use repeat until the list is updated, this may take a few moments.

Verify the registration traffic on the regional RP

Since the RP-SET was distributed and the RP for 239.1.212.0/24 is now inside the 'b2' zone, any new multicast transmissions to 239.1.212.0/24 should not trigger a registration with Sw1 anymore, since they will be registered with Agg1(b2) or Agg2(b2).

This means that a regional RP can reduce the control traffic between the sites.

However, when a client in the other building requests the multicast traffic, it will still pass the Sw1.

4. On Sw1, review the current Rx RP Register counters of Loopback0
5. On Agg1, review the current Rx RP Register counters of Loopback2, repeat this for the VSX peer
6. On PC4, start a new mcast transmission to 239.1.212.11.
7. On Sw1, repeat the Loopback0 counters command. The counter should not have increased.
8. On Agg1, repeat the counters command and the VSX-peer counters command. The RX Register counter should have increased.

Optional Task 5: Understand the MAC - IP Multicast relation

Objectives

- Understand the Ethernet MAC to IP Multicast relation
- Understand the risk of the mapping model

Steps

Explore the Multicast IP to Ethernet mapping

The last 23 bits of a multicast IP address are used to form the Ethernet multicast destination address. Let's review this on PC4

1. On PC4, start a mcast-tx to 239.1.1.1
2. Start Wireshark on PC4, and capture the Lab NIC traffic for a few seconds, stop the Wireshark trace.
3. Select a packet with destination IP 239.1.1.1, review the packet contents panel.

Q: What is the destination Ethernet MAC address?

A: The destination MAC is 01:00:5E:01:01:01. (IPv4mcast block).

Q: What is the IPv4 Multicast range? How many significant bits define that this is a multicast address?

A: 224.0.0.0/4, 4 bits.

Q: Only the last 23 bits of the Multicast IP are used for the Multicast Ethernet address. What does this mean?

A: This means that several IP multicast addresses are actually using the same Ethernet MAC address.

4. Start a new transmission to destination 238.1.1.1.

5. Start a new Wireshark trace and stop it after a few seconds. Stop the multicast transmission.

Q: What is the destination MAC address?

A: The destination MAC is 01:00:5E:01:01:01. (IPv4mcast block).

Q: What would happen when both 238.1.1.1 and 239.1.1.1 are transmitted on the network?

A: In case both multicast transmissions would be active, the Layer2 switches cannot differentiate between this traffic.

Task 6: VSX Multicast Failover

Objectives

- VSX PIM active-active
- Understand the VSX Multicast Routing failover
- Understand the VSX PIM Proxy
- New registrations during the failover

Steps

VSX Active Active Configuration

In the current state of the lab, when the PIM DR would go offline, it would take time for the PIM router on the other node to detect this and become the PIM DR for the subnet.

In VSX, a PIM feature that is known as 'active-active' can be used for fast failover. When the feature is enabled, both nodes will actively join the multicast stream on the upstream router. This means that the multicast traffic is actually received by both nodes.

One VSX node will still be the DR, but the other node will get the role of proxy DR. This means it will not actually forward the traffic on the client subnet, but it is 'hot standby' to take over.

In this task, you will explore the active-active feature configuration and operation.

1. On Agg1, the primary VSX node, configure active-active on both default and b2 VRF contexts for router PIM.
2. On Agg2, repeat this configuration

Understand the VSX Multicast Routing failover

In this section you will explore how VSX handles multicast traffic during a node failover.

VSX PIM active-active feature ensures that each VSX node will be actively joining upstream multicast streams.

In the lab topology, this means that both Agg1(default) and Agg2(default) PIM routers will notify Sw1 that they need to receive the multicast stream.

In the previous tasks you have seen that only the PIM DR would register with the upstream router, so that behavior will change now.

Only 1 node will actively forward the stream on the endpoint Layer3 interface, this is known as the active DR. The other node will be in 'standby' mode, this status will be the 'proxy-DR'.

You will now explore this mechanism and verify the failover by monitoring the traffic on the PC3 client.

First you need to identify your VSX node Layer3 MAC address.

3. On Agg1, take note of the MAC address.
4. On Agg1, repeat this for the VSX peer and note the MAC address of Agg2 SVI11.
5. On Agg1, review the currently active PIM DR on the multicast receiver SVI 11. Repeat this for the VSX-peer. Notice the 'Proxy DR' value.

6. On Agg1, you can also use the brief option

Q: Which switch is currently the active PIM forwarder?

A: Agg1 is currently the active forwarder, it has the 'proxy DR' value as 'false', while Agg2 has it as 'true'.

Now verify this information on the client side.

7. On PC4, start a mcast tx to 239.1.212.1.
8. On PC3, start a mcast rx for 239.1.212.1.
9. On PC3, start a Wireshark trace and stop it after a few seconds.
10. On PC3, in the Wireshark trace, select one of the received multicast packets.

Q: What is the source MAC address for this multicast frame?

A: It should be the same as the currently active VSX node forwarder.

11. On Sw1, you can now verify that both nodes have joined the multicast group. First list the group and the source address in the multicast routing table.
12. On Sw1, now review the details of this source group combination.

Q: What do you observe?

A: Previously, only SVI 42 was listed as the downstream interface. Due to the active-active configuration, both VSX nodes have now joined the multicast and both SVI 41 and SVI 42 are listed as active downstream interfaces.

Now complete your review by checking the same multicast routing entries on Agg1 and Agg2.

13. On Agg1, review the details of the group / source combination. Repeat this for the VSX peer.

Q: What do you observe?

A: Both nodes have the SVI 11 state as 'forwarding'. The Proxy DR is not really forwarding yet, but it is ready to take over when the DR would go offline.

Verify the failover

You will now verify that the 'proxy DR' is really a 'hot-standby', since it has joined the upstream PIM router for this multicast stream.

In this section you will explore the impact of a VSX node failover on the multicast routed traffic. In the previous section you have found the VSX node that is currently forwarding the traffic.

14. On the current DR Proxy node (the node with 'proxy DR: true' node), run these commands:
15. On PC3, start the Wireshark again and leave it running. You should observe incoming multicast traffic about once every second.
16. On the currently active VSX node, save the configuration and reboot the system.
17. Switch to the PC3.

Q: Did the multicast stream stop?

A: No, it continued virtually uninterrupted.

18. Stop the Wireshark trace. Review the source MAC address of the last multicast packet.

Q: What do you observe?

A: The source MAC address has now changed to the MAC address of the other VSX node. This demonstrates the seamless failover of multicast traffic with the VSX system.

19. On the 'other node' connection, stop the repeat.

Q: Did the 'proxy DR' state change?

A: Yes, this system is now the active forwarder.

Task 7: IGMP ACL

Objectives

- Understand how to configure an IGMP Querier ACL

Steps

IGMP Querier ACL

To control the IGMP joins, the IGMP Querier can be configured with an ACL. This ensures that only the authorized IP Multicast groups can be joined using the configured IGMP Querier.

1. On Agg1, define an IP ACL that denies the multicast range 239.1.2.0/24 (as destination IP), permit any other traffic and enable the vsx-sync.
2. Apply this ACL on the SVI 11 to IGMP.
3. Repeat this configuration on Agg2.

NOTE: The ACL will process new joins, it will not take effect on existing multicast flows.

Verify normal flows operate fine.

4. On PC4, configure a mcast tx to 239.1.212.3 (some new multicast address, since the filter does not apply the existing flows)
5. On PC3, use mcast rx to join 239.1.212.3. Verify you receive traffic.

Verify 239.1.2.0/24 flows are blocked.

6. On PC4, configure a mcast tx to 239.1.2.1
7. On PC3, use mcast rx to join 239.1.2.1. Verify the traffic is not received.
8. On Agg1 review the IGMP counters for VLAN11 and for the VSX-peer.

Q: Do you observe any dropped packets by ACL?

A: Yes, both Agg1 and Agg2 will show dropped packets.

IGMP Snooping ACL

The filtering can also be applied at the IGMP snooping level.

In this example, you will apply the filter on Sw2 for the range 239.1.3.0/24.

9. On Sw2, define an IP ACL igmpsnoop, permit any other traffic.

10. On Sw2, apply this 'igmpsnoop' ACL to the VLAN 11 IGMP snooping
11. On PC4, start a mcast tx to 239.1.3.1.
12. On PC3, attempt to receive mcast 239.1.3.1, it should not receive the multicast traffic.
13. On Sw2, review the IGMP snooping counters for VLAN11. You should observe that the 'Packets dropped by ACL' has increased.

IGMP debug

14. On Sw2, enable debugging for 'igmp config'.
15. On Sw2, enable terminal-monitor for severity debug and filter on '239.1.3'
16. On PC3, attempt to join mcast 239.1.3.2.
17. On Sw2, review the debug output. You should see matched by ACL messages.

This concludes the IGMP ACL task.

18. On Sw2, disable all debugging.
19. On Sw2, assign PC4 port to VLAN 12
20. On PC4, release and renew the IP address, it should get a 10.x.12.0/24 IP address.
21. Optional step: On Agg1, Agg2, Sw1 and Sw2 create a new checkpoint 'asts-lab09-<your name>'

This concludes the multicast lab.

Lab 10: Quality of Service

In this lab you will review the QOS features on the AOS-CX platform.

Objectives

- QOS Policies to remark traffic
- Apply QOS Cos and DSCP mappings
- Rate limiters and Traffic shaping

Requirements

This lab requires completion of Lab 09.

Aruba Training-Confidential


Scenario

The customer wants a demonstration of how to control the QOS marking and queuing of TCP 8000 traffic sent by PC4 to PC3.

1. To test traffic from PC4 to PC3:
 - a. On PC3, start **ASTS-Lab10-PC3-qos-rx.cmd**
 - b. On PC4, start **ASTS-Lab10-PC4-qos-tx.cmd** and enter the PC3 IP address.
2. On Sw2, the customer requires a QOS policy on port 1/1/4 for the PC4 TCP 8000 traffic to remark the DSCP mark to CS1 and the 802.1p mark to 2. They also want to see how many packets match the policy.
3. Prove to the customer the traffic remark was done.
 - a. Prove the remark when the traffic leaves port 1/1/27 of Sw2 to PC3 (to Sw1). You have to demonstrate a live packet trace in the SSH session that shows the CS1 mark for this traffic.
4. After reviewing the queue statistics on the Aggregation switches, the customer has the impression that the test traffic is not handled by the correct queue. Demonstrate to the customer the COS and DSCP mapping tables and demonstrate that traffic can be mapped based on the COS or DSCP value. The final configuration should assign the queue based on the DSCP value.
5. While the customer is now pleased that the Aggregation switches perform the correct queuing based on the DSCP mark, they have noticed that the Sw2 uplink LAG connected to the Aggregation switches is not handling the test traffic based on the correct queue, even though the DSCP trust is configured. They need to you investigate and correct the Sw2 configuration, so that the PC4 traffic that leaves Sw2 to the Aggregation switches is also placed in the correct queue (0).
6. The customer plans to transport the traffic coming from a partner company over their network. They want to ensure that the partner company application will not consume too much bandwidth in their network. On the Aggregation switches, configure a rate limiter for the test traffic (TCP 8000, tested from PC4 to PC3) that limits the inbound bandwidth from Sw2 to 20Mbps for this traffic and a burst of 20KB. The rate limiter should remain effective when either VSX node is rebooted.
7. After testing the rate limiter, the partner complains that the configured rate limiter does not allow the configured bandwidth (20Mbps), it barely reaches 5-10Mbps. They want you to adjust the Sw2 uplink configuration (the LAG to the Aggregation switches) to be in line with the Aggregation rate limiter so the full bandwidth (20Mbps) can be used.

Finally, the customer wants you to change the rate limiter from 20Mbps to 10Mbps on the Aggregation switches, and the partner should be able to use the full 10Mbps bandwidth without making any change on the Sw2 side.

Lab Diagram



Task 1: QOS Policies

Objectives

- Configure QOS Policies to remark traffic
- Explore interface and Queue statistics
- Traffic mirroring to monitor Qos marks

Steps

Review default queuing and QOS maps

1. **On PC3**, open the ASTS folder on the desktop and start **ASTS-Lab10-PC3-qos-rx.cmd**. Take note of the PC3 IP address in the output.
2. **On PC4**, open the ASTS folder on the desktop and start **ASTS-Lab10-PC4-qos-tx.cmd**. Enter the PC3 address as the target.
3. **On PC3**, verify traffic is received in the TCP8000 iperf.
4. **On both Agg1 and Agg2**, clear the int 1/1/2 statistics.
5. **On both Agg1 and Agg2**, show the interface 1/1/2 queue statistics.

NOTE: The iperf traffic will use either the path via Agg1 or Agg2. In the next sections, you only need to review the statistics on the actual Aggregation switch that handles the traffic (either Agg1 or Agg2).

Q: What do you observe?

A: Q1 is handling most of the traffic

Q: Historically, Q2 would be used for the default (best effort) traffic. Why is Q1 used now for the best effort traffic?

A: The standard has been updated, so only 1 value is used now for Background (Scavenger) traffic (queue 0 by default). Value 1 is now used for Best Effort traffic, value 2 is used for 'Excellent Effort'.

6. On Agg1, review the qos cos-map.

Q: To what local priority is the 802.1p value 2 mapped?

A: It is labeled 'Excellent Effort' and assigned to local priority 2.

Q: To what local priority is 'Best_Effort' mapped?

A: To local priority 1.

7. Review the qos dscp-map.

Q: To what local priority is CS0 (no marking / default) mapped?

A: This is considered Best Effort, it is mapped to local priority 1.

8. Review the qos queue-profiles and review the queue-profile named 'factory-default' to see the local priority to queue mapping . By default, there is a one-to-one mapping for the local priority to a queue.

Configure QOS Policy on Sw2

Define a new policy on Sw2 (the access switch that receives traffic from PC4) that will control the inbound TCP8000 traffic on port 1/1/4.

Initially you will just remark the traffic, next you will apply more controls.

9. On Sw2, define a class 'bulk-data' that matches any traffic to TCP port 8000. Enable the count option.
10. On Sw2, define a policy 'asts-qos' that assigns the class 'bulk-data' the DSCP value CS1 and PCP value 2.

Q: What does the PCP indicate?

A: The Priority Code Point is the 802.1p mark in a Layer2 tagged frame. This is used by the COS trust model.

NOTE: In a real deployment, only 1 method is typically used to mark the traffic. In this lab it is only applied to demonstrate the different QOS trust modes.

11. On Sw2, apply the policy inbound on port 1/1/4.

Verify the Remark - last hop switch

In the real world, you may need to proof the remark at the last hop switch before it reaches the endpoint. This can be achieved use a traffic mirror.

First you will review the mirror.

12. On Sw2, configure mirror session 1. Ensure the destination is set to CPU and that the mirror is enabled.

IMPORTANT: Remove any other source interfaces that may be present, such as lag255.

13. On Sw2, verify the current configuration of the mirror 1

14. On Sw2, define a policy 'qos-mirror' that matches class 'bulk-data' and applies action 'mirror 1'

15. Apply the policy on 1/1/27 in the outbound direction (remember PC3 is connected on this port via VLAN11 on Sw1 as Layer2 'pass-through').

16. Review the mirror 1, repeat the command after about 10 seconds.

```
...
```

Q: Do you see the 'Output packets' counter increase?

A: Yes, the TCP8000 packets of the iperf session are now mirrored.

17. Run a local tshark on Sw2 for a few seconds, then stop it again using CTRL-C. You will get a lot of output for a few moments.

18. Review the last packet in the trace. Scroll up to see the TCP (Transmission Control Protocol) and IP (Internet Protocol) section of the packet.

Q: What is the TCP destination port number?

A: 8000

Q: In the IP section, what is the DSCP value of the packet?

A: The packet has DSCP value CS1. This confirms that the DSCP remark was done.

Default Queue and Trust modes

19. On the Agg switch that handles the traffic (**either Agg1/Agg2**, see the previous statistics output), clear the port 1/1/2 statistics and check the queue statistics for 1/1/2.

```
# clear interface 1/1/2 statistics
```

Q: What do you observe?

A: The traffic is not automatically assigned to a local priority, so it is handled by the default queue, even though the traffic has a custom COS and DSCP mark.

20. Review the default global and interface QOS trust mode on Agg switch.

Q: What is the default mode?

A: Trust mode is set to none.

Verify the impact of the QOS trust modes.

21. Apply global trust mode 'cos' on Agg1. Clear the statistics and verify the queue assignment on the Agg switch that handles the traffic.

Q: What do you observe?

A: The switch will now use the 802.1p PCP COS value and the cos-map. Traffic will be handled by Queue 2 as a result.

NOTE: You can also set the QOS trust mode at interface level, this overrides the global level.

22. Set the global trust mode to DSCP. Clear the statistics and verify the queue assignment.

```
qos trust dscp
```

```
# clear interface 1/1/2 statistics
# show interface 1/1/2 queues
# repeat count 5
```

NOTE: QOS trust DSCP is the best practice.

23. Repeat the QOS trust dscp configuration on Sw1 and Sw2 to ensure all devices trust DSCP.

On Sw2, you will now review the local traffic handling.

24. On Sw2, clear the current session statistics of LAG255.

25. Review the queue statistics of the LAG 255 and the member ports.

Q: What do you observe?

A: AOS-CX will automatically combine the statistics of the members ports for the aggregated port.

Q: What queue is use to transmit the traffic?

A: Q1, the Best Effort queue.

26. Update the 'asts-qos' policy configuration so the traffic is not only remarked with DSCP but is also handled by Queue 0 (local priority 0). Remove the PCP action.

27. Verify the policy configuration

28. Verify that the traffic is now handled by the correct queue based on your local priority configuration.

Aruba Training-Confidential

Task 2: Rate limiter and Traffic shaping

Objectives

- Understand methods to control bandwidth

Steps

Inbound Rate Limiters

When traffic is received from a partner or untrusted source, a rate limit can be applied to control the volume of the traffic.

In this task you will explore the impact of rate limiters.

First you will configure the Aggregation switches to apply a rate limiter to the bulk-data. While doing so you will get some experience to troubleshoot some VSX synchronization scenarios.

Define the Rate Limiter on the Aggregation VSX

1. On Agg1, define the class 'bulk-data' for any to destination TCP8000. Enable the count option for this rule.
2. Define a policy 'asts-qos'.
3. Assign class 'bulk-data' with the 'cir' action. Set it to 20.000 kbps, set the cbs (burst size) to 20000. Any excessive traffic must be dropped.
4. Apply the policy to the LAG2 in the inbound direction.

Troubleshoot the VSX-sync

Review the VSX sync status

5. Review the VSX-sync.

Q: What is the Error State?

A: The system shows 'Missing reference error'

Q: What does this mean?

A: Configuration objects may contain references to other configuration objects. When a configuration is synchronized and it points to an object that was not synchronized, a reference error can be observed.

6. Let's review the peer differences.

Q: What line is different?

A: The policy line is applied on Agg1, not on Agg2

7. On Agg2, review the defined policies.

Q: What do you observe?

A: There is no policy.

8. On Agg1, update the policy with the vsx-sync command.

9. On Agg2, review the VSX sync and the current configuration of the interface LAG255.

Q: What do you observe?

A: Agg2 reports that the qos policy is being processed.

10. On Agg1, enable the vsx-sync command on the 'class ip bulk-data' class object.

11. On Agg2, review the VSX sync

Q: Are there any errors reported?

A: The Error state is 'None', all configurations have been synced successfully now.

12. On Agg2, review the current configuration of the interface LAG255 and review the peer differences.

Q: What do you observe?

A: Agg2 running configuration contains ! references about the QOS policy. The active configuration does not match the running configuration and must be reset.

Q: What does this mean?

A: Even though you have corrected the configuration issue, a policy reset command must be issued to re-apply the policy configuration.

13. On Agg2, reset the 'asts-qos' policy.

14. On Agg2, verify that the messages have disappeared from the running configuration.

You have now completed the rate-limiter configuration on both VSX nodes.

Verify the rate limiter

Now that the rate limiter is active, let's take a look at the effect.

15. On PC4, review the iperf output.(check the **SUM** value, this is the combined bandwidth of the streams).

Q: what is the reported SUM throughput in iperf?

A: Currently, there is about 1-2Mbps of traffic.

Q: How is this possible, since the configuration is set to 20000kbps?

A: The inbound rate limiter has been configured with a very small burst size, so TCP suffers from this configuration. This is typical when a service provider limits the inbound traffic into its network.

Outbound Traffic shaping

You observed that an upstream rate limiter can have significant impact on the actual TCP throughput. In case you know that such a limiter is in place, you can configure an outbound traffic shaper, so that any excess traffic will first be locally buffered. This relieves the TCP retransmissions.

16. On Sw2, configure the uplink interface LAG255 with a qos shaper of 20000kbps.

17. On PC4, verify the impact. The iperf should now run close to 20Mbps again.

Rate limit and burst size

It is also possible to adjust the inbound burst size. Increasing this inbound burst size will have a similar effect as the outbound shaper.

On the Aggregation layer, lower the inbound rate-limiter to 10Mbps.

18. On Agg1, enter the 'asts-qos' policy and change the current line

19. On PC4, review the impact.

Q: What do you observe?

A: Since the limiter is lower than the outbound shaper, the traffic is dropped on the Agg side and it is much slower again.

20. On Agg1, update the command and increase the cbs value.

21. On PC4, verify the impact. Traffic should now be around 10Mbps again.

This demonstrates the rate-limiter and traffic shaper features.

22. On PC3 and PC4, close the lperf windows.

23. Optional step: On all switches, define a new checkpoint 'asts-lab10-<yourname>'.

This concludes the lab activity.

Lab 11: Dynamic Segmentation

In this lab you will explore several AOS-CX authentication features and review common troubleshooting on the RADIUS communication.

Objectives

- Configure and troubleshoot 802.1X authentication.
- Configure and troubleshoot RADIUS tracking.
- Configure and troubleshoot device profiles.
- Configure and troubleshoot User Based Tunneling.
- Configure and troubleshoot ClearPass Downloadable User Roles.

Requirements

This lab requires completion of Lab 10.

Scenario

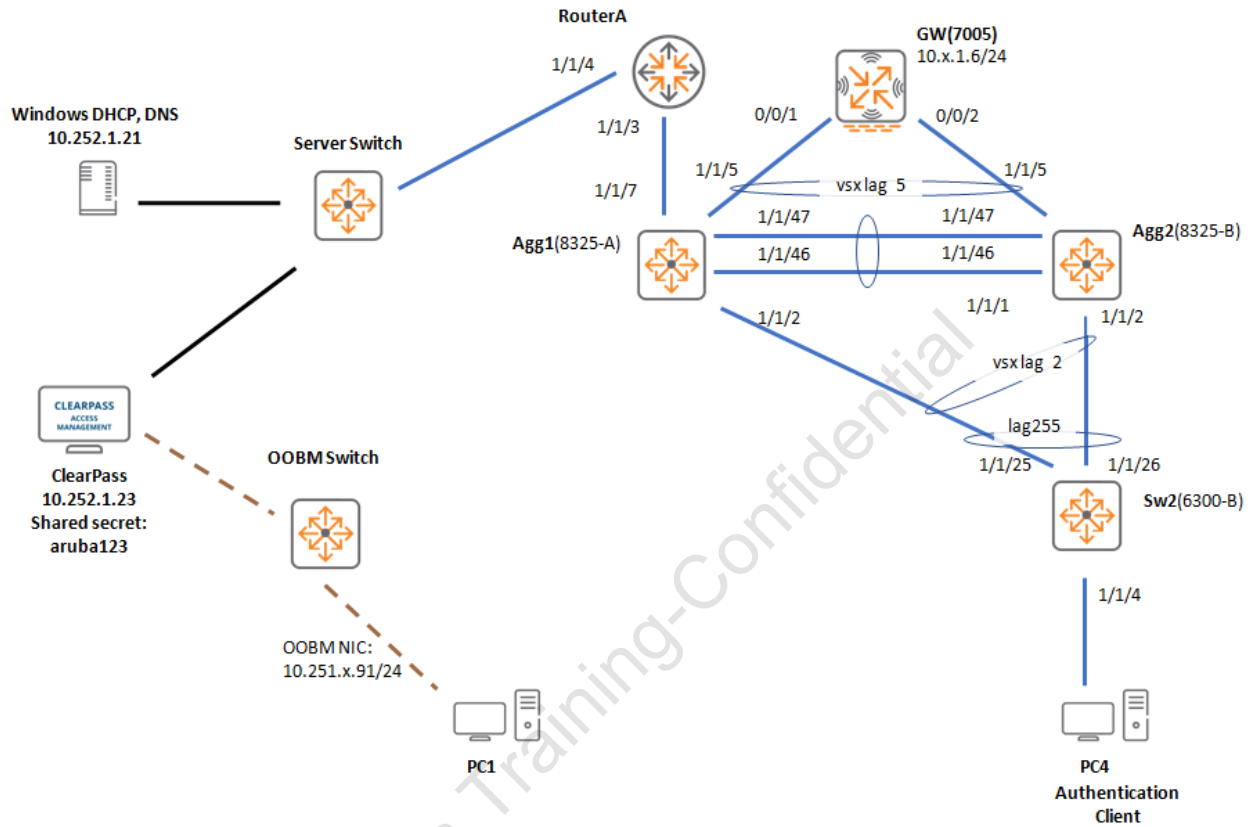
Sw2 will be used for all authentication configuration.

1. The customer wants you to configure 802.1X authentication on port 1/1/4 of Sw2:
 - Use RADIUS server 10.252.1.23
 - Use RADIUS secret aruba123.
 - Test the authentication using PC4 with credentials asts-accept / aruba123.
2. The customer plans to use User Roles or RADIUS based VLANs for both 802.1X and Mac authentication. In case a role has a missing VLAN ID or a RADIUS accept message does not contain a VLAN assignment, the device should be assigned to VLAN 19, a 'parking' VLAN that does not provide any network access.
3. The customer wants you to demonstrate several configuration errors and how they can be detected.
4. Configure PC4 to authenticate using asts-employee-vsa / aruba123.
 - Use debugging and network traces to demonstrate:
 - i. why the user cannot be authorized
 - ii. how the name of the returned Aruba-User role can be retrieved.
 - Define this role and configure it with VLAN11.
5. The customer wants to observe what happens to the existing users when the RADIUS server is unavailable. To enhance the demonstration:
 - Configure the User role with a re-authentication timer of 30 seconds.
 - Disable access to the RADIUS server (disable the uplink port on Agg1 or configure a deny ACL for the RADIUS IP)
 - Verify the behavior.
6. The customer has several critical devices. When the RADIUS server is unavailable, it is critical these devices are still able to access VLAN 11.
 - Demonstrate this after connecting a new device to port 1/1/4 of Sw2.
7. When the RADIUS server is unavailable, authenticated devices should be allowed a first re-authentication. To test this, configure a short re-authentication timeout of 30 seconds and then block the communication to the RADIUS server.
8. The customer also wants to ensure that the switch will automatically attempt to test the connection to the RADIUS server every 60 seconds when the RADIUS server is not available using the credentials radius-tracking-user / aruba123.

9. The customer has a few sites that cannot be migrated to full centralized RADIUS authentication yet. In the transition period, they plan to use local MAC authentication on the switches to assign printers to the correct VLAN.
 - Demonstrate this using port 1/1/4 on Sw2.
 - Use the OUI of the PC4 MAC address to simulate a printer.
 - Disable 802.1X on the port and make sure that the PC4 is authenticated with a User Role 'printer' and assigned to VLAN11.
10. The customer wants any devices that don't match the printer OUI MAC range are assigned the User Role 'asts-fallback' in VLAN11 (They understand that they can customize the role later).
11. The customer wants to understand what happens when a device matches both:
 - a local MAC authentication
 - 802.1X authentication
 - Configure port 1/1/4 and PC4 (that currently authenticates as printer based on the LMA) with 802.1X and demonstrate the onboarding order.
12. The customer heard good reviews about the integration of the switches with the Aruba Gateway devices to provide firewall functions.
 - Integrate the switch with the Aruba Gateway (10.x.1.6)
 - ensure that traffic from an endpoint that has received the switch User role 'employee' is handled by the Aruba Gateway using firewall policy 'authenticated'.
 - Show the debug output for:
 - i. the session setup between the switch and the Gateway
 - ii. when the user 'asts-employee-vsa' authenticates.
13. The customer wants to understand what happens when a switch user role is configured with a firewall policy that does not exist on the Aruba Gateway. Demonstrate this with:
 - the switch User Role 'employee'
 - a firewall policy of 'sales' (this policy does not exist on the Aruba gateway)
 - demonstrate the debug or log output on both the switch and the Aruba Gateway.
14. Since the traffic between the switch and the Aruba Gateway is transported using a GRE tunnel, describe any impact on the MTU for the clients.
15. Demonstrate:
 - the normal maximum MTU on a non-tunneled port (use PC3 for this test)
 - how to setup the tunneled configuration so that a tunneled client can achieve the same MTU as the non-tunneled client
 - i. the PC4 should be used as the tunneled client
 - ii. use the ping options on the PC4 client to set the 'do not fragment' option and configure the maximum payload size.
16. Since QOS is important for the customer, they are worried that any tunneled traffic may not receive the correct QOS handling in the network due to the GRE tunnel encapsulation.

- Demonstrate that a wired tunneled client (PC4) ping to the server 10.252.1.23 is assigned to queue 3.
 - i. Use the path from PC4 to the Aruba Gateway, review the queue statistics:
 - 1. on the uplink of Sw2 to the VSX Aggregation
 - 2. on the uplink from the VSX Aggregation to the Aruba Gateway
- 17. Save a pcap network trace file to PC1 that contains the GRE encapsulated traffic. This will demonstrate both the inner (user data IP header) and the outer (GRE IP header) DSCP markings.
- 18. The customer also heard how easy it can be to deploy the User Roles using ClearPass.
 - Configure the integration with ClearPass using credentials duradmin / aruba123.
- 19. The customer is aware that the integration with ClearPass is properly secured using HTTPS connections and validated certificates. Demonstrate how to:
 - Find the subject name of the ClearPass certificate that should be used in the switch configuration
 - Ensure the switch can resolve the name of the ClearPass server.
 - View debug output to view an invalid RADIUS name, after certificate authentication fails.
 - Successfully authenticate with the account 'asts-employee-dur' / aruba123.

Lab Diagram



Task 1: Prepare the configuration

Objectives

- Verify access to the ClearPass host.
- Review the RADIUS test accounts table.

Steps

Configuration

Sw2 will be the target switch for all the authentication configuration activities.

1. On Sw2, verify you can reach the ClearPass host using the in-band network.

Table with RADIUS test user accounts

The accounts in the table below are test accounts that have been defined on the RADIUS server. Some of these will be used in this lab activity, you can also use them if you want to test on your own.

IMPORTANT: This table is for your reference, you **do not** need to login at this point.

All RADIUS test user accounts have been configure with password 'aruba123'.

'vsa' stands for 'Vendor-Specific Attribute'. It will return a vendor 'Aruba-User-Role' attribute.

'dur' stands for 'Downloadable User Role'. It can be used to test the ClearPass Downloadable User Roles.

'ietf' stands for standard IETF attribute assignment.

Username	Settings returned by ClearPass
asts-employee-vsa	Aruba-User-Role: employee
asts-employee-dur	Aruba-CPPM-Role: employee (Downloadable)
asts-employee-ietf	Access-Accept with IETF Attributes VLAN Nas-Filter-Rules
asts-reject	Access-Reject
asts-accept	Access-Accept only, no attributes

Task 2: Configure basic 802.1X authentication

Objectives

- Configure the Sw2 to perform basic 802.1X authentication
- Understand the EAP frames and RADIUS packets relation

Steps

In this task you will configure the switch to perform basic 802.1X authentication, verify the access and review the normal EAP and RADIUS statistics for an 802.1X connection.

Basic RADIUS and 802.1X configuration

1. On Sw2, define the CPPM RADIUS server 10.252.1.23 with shared secret 'aruba123', assign it to radius group 'cppm'.

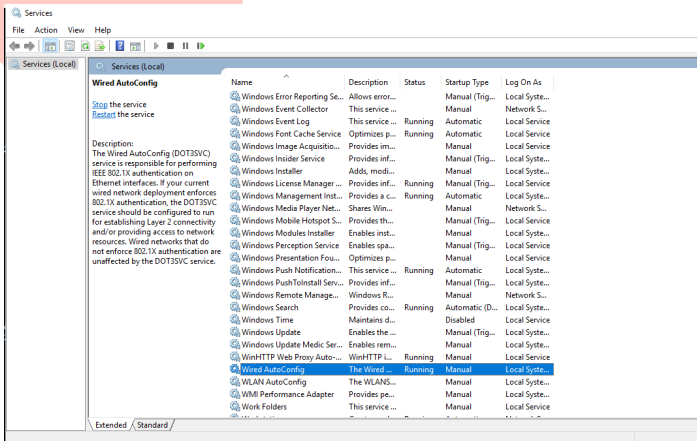
NOTE: Sw2 will connect to the CPPM host using the default VRF, not the mgmt VRF.

2. On Sw2, configure 802.1X to authenticate using the server-group 'cppm'
3. On Sw2, enable 802.1X authentication globally
4. On Sw2, on the port 1/1/4, enable 802.1X and adjust the 802.1X timers.
5. Define VLAN 19 with description, set the port VLAN for 1/1/4 to VLAN 19. This ensures that users without role or invalid roles will remain blocked in this VLAN 19.
6. Use PC1 to connect to the ClearPass server on <https://10.252.1.23>. Login using readonly/readonly. Navigate to Monitoring > Live Monitoring > Access Tracker. You can use this connection to see the authentication status on the RADIUS side.

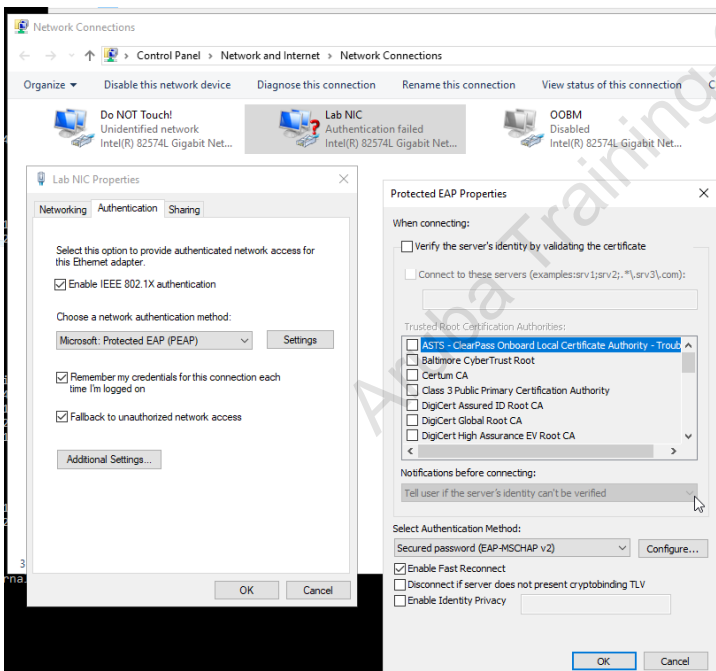
NOTE: This connection is available via the OOBM network. PC1 does not need to be connected on the lab network.

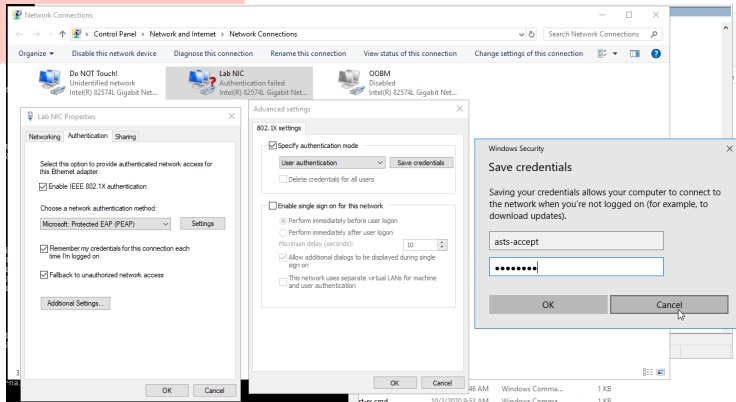
7. Configure PC4 with 802.1X authentication. You can use the user 'asts-accept' with password aruba123 to verify the connection.

NOTE: Make sure to enable and start the 'Wired autoconfig' Windows service on the PC4. This is the Windows 802.1X supplicant.



TIP: In this lab environment, the client does not trust the CPPM RADIUS EAP certificate, so you should disable the certificate validation in the Windows NIC 802.1X authentication settings.





NOTE: Make sure the option 'Fallback to unauthorized network access' is enabled.

8. On Sw2, review the authenticated clients and the details

Statistics EAP and RADIUS

In this section you will review the EAP and RADIUS statistics and explore the relation between them. First you will clear all the statistics, next the client will be re-authenticated and the updated statistics can be compared.

9. On Sw2, clear the 802.1X authenticator statistics.

10. Verify that the port statistics and the client-status have been cleared.

11. Clear the RADIUS server statistics and review that they have been cleared.

12. Enable debugging for the radius and 802.1X portaccess modules.

13. Clear the debug buffer

14. Manually re-authenticate the client

15. Review the port statistics after the clients was re-authenticated.

16. Review the RADIUS statistics

Q: Do you see a relation between the number of EAP packets and the number of RADIUS packets?

A: Yes, since the 802.1X EAP frames are encapsulated in RADIUS packets and sent to the RADIUS server, the 2 statistics will be similar.

Review the debug output for a successful authentication

17. On Sw2, review the debug buffer. Make sure you see these events:

- the initial event 'Re-Authenticate Client on Port'
- dot1x pae State transition AUTHENTICATED > REAUTHENTICATING
- several RADIUS packet exchange events will be displayed
- dot1x pae State transition REAUTHENTICATING > AUTHENTICATED

Aruba Training-Confidential

Task 3: RADIUS Troubleshooting

Objectives

- Invalid Aruba-User-Role
- Tracing RADIUS control plane
- Configure the critical role
- Configure cached re-authentication
- Understand the need for RADIUS tracking
- Tracking dead-only

Steps

Invalid Aruba-User-Role

In this section, the RADIUS server returns an Aruba-User-Role that doesn't exist in the switch configuration. You will explore how to troubleshoot this.

1. On Sw2, enable all portaccess debugging
2. On Sw2, clear the debug buffer
3. Reconfigure PC4 to authenticate with 802.1X using the 'asts-employee-vsa/aruba123' user. The RADIUS server will return the Aruba-User-Role 'employee' for this account, but this role does not exist in the switch configuration.
4. On Sw2, review the client details

Q: What is the Authorization status?

A: Invalid.

5. On Sw2, review the debug buffer. You should see a message 'Authenticated but failed to apply role'.
6. When the details of the RADIUS packet are not reported in the debug log, a trace can help. Open a second SSH connection to Sw2 and run a TCPDUMP for RADIUS traffic, leave it running.
7. In your first connection to Sw2, trigger re-authentication on port 1/1/4.
8. In the TCPDUMP session, RADIUS packets should have been captured. Use CTRL-C to stop the TCPDUMP and review the last packet (Access-Accept). You should observe Vendor-Specific Attribute for Vendor id 14823. This is the code for Aruba Networks. The Vendor Attribute 1 (the code for Aruba-User-Role) will show the name of the assigned role.
9. Close the second session (the TCPDUMP session).
10. On Sw2, configure the local user role 'employee' with access VLAN 11.

11. On Sw2, trigger re-authentication and verify that the client is now online.

Cached re-authentication

In this section you will review the cached re-authentication feature. An administrator can allow the switch to perform re-authentication of a previously authenticated user or device when the RADIUS server is not reachable.

Suppose a branch location is using port access authentication and the VPN link to the datacenter goes down. In that case the RADIUS server is no longer reachable. This will have no immediate effect on the existing sessions, until devices reach the re-authentication timer. The re-authentication timer on the port or the role can be considered a grace-period during which the client can remain online.

When the re-authentication occurs, the authentication would fail and the users would no longer be connected.

With cached re-authentication, an overall grace-period can be configured during which clients will be re-authenticated and remain online with the same settings as the previous authentication session.

12. On Sw2, modify the local user role 'employee' with a re-authentication timer of 30 seconds.

NOTE: This short re-authentication timer is for this lab purposes only and should not be applied in production environments.

13. On Sw2, enable debug and terminal monitor.
14. On Sw2, trigger a re-authentication using the command
15. On Sw2, the terminal monitor should show the messages

AUTHENTICATED > REAUTHENTICATING
REAUTHENTICATING > AUTHENTICATED

16. On Sw2, after about 30 seconds, new messages should confirm that the re-authentication occurs every 30 seconds.

Verify the operation without cached re-authentication

17. On Agg1 port 1/1/7 (uplink that supports access to RADIUS server), apply an ACL that blocks IP traffic to the RADIUS server
18. On Sw2, attempt to ping 10.252.1.23, this should fail now.
19. On Sw2, wait up to 30 seconds and review terminal monitor output.

Q: What happened?

A: The switch attempt to reauthenticate, but it does not succeed. The client transitions to UNAUTHENTICATED.

20. On Sw2, stop the terminal-monitor and review the debug buffer.

Q: What error messages do you see?

A: Messages about RADIUS server timed out.

21. On Sw2, review the RADIUS server statistics, the timeouts and retransmits value will not be 0 anymore.

22. Review the port access clients. The client should now have failed authentication.

Critical Role for Authentication

When configured, the critical role will be applied when the RADIUS server is unreachable. Customers can use this when they prefer to 'open' the network in case the RADIUS is unreachable. By default, the ports would remain 'blocked' when the RADIUS server cannot be reached, as you have seen in the previous test.

23. On Sw2, define a new user role 'asts-critical'

24. On Sw2, configure the port 1/1/4 with the critical role 'asts-critical'

25. On Sw2 bounce port 1/1/4 and review the authentication status.

Q: What role was assigned to the client?

A: The client has now been assigned the asts-critical role.

Configure the switch with cached re-authentication

In this section you will allow the switch to 'reauthenticate' a client based on the currently active session when the RADIUS server cannot be reached anymore. This feature is known as 'cached re-authentication'.

First restore the RADIUS communication.

26. On Agg1, remove the ACL application from port 1/1/7, but keep the ACL definition.

Configure cached re-authentication.

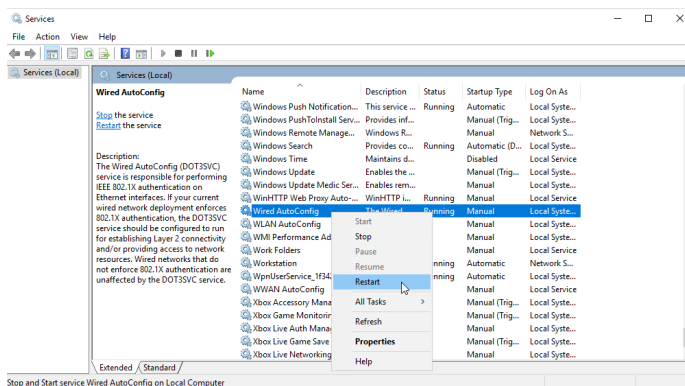
27. On Sw2, configure the port 1/1/4 with 802.1X cached re-authentication.

NOTE: It is also possible to set a cached re-authentication period to control the cache time. The default period is 30 seconds.

28. On Sw2, bounce port 1/1/4

29. On Sw2, verify the PC4 client is authenticated again

NOTE: On PC4, the Windows supplicant has a local 'quiet' timeout after some failed authentications. Restarting the 'Wired AutoConfig' service clears the timer.



30. On Sw2, enable the terminal monitor with dot1xpae filter.

31. On Agg1, block the RADIUS communication again.

32. On Sw2, review the log output.

Q: Did the client authentication change in the first minute?

A: No, the client has been re-authenticated based on the CACHED-REAUTHENTICATION.

33. After the status changed to UNAUTHENTICATED, review the authenticated clients.

Q: What happened with the client?

A: Since the cache timer has passed, the switch will no longer re-authenticate the client. The client will be assigned the critical role.

34. Cleanup: On Sw2, remove the cached re-authentication commands from the port 1/1/4.

35. Cleanup: On Agg1, remove the ACL on port 1/1/7

NOTE: Cached re-authentication must be configured per authentication method, so there is also a version under the mac authentication context.

36. On Sw2, trigger re-authentication and verify the client successfully connected.

NOTE: Restarting the PC4 'Wired AutoConfig' service can speed up the process.

Understand the need for RADIUS tracking

When the RADIUS server cannot be reached during authentication, the authentication will eventually timeout and this will result in a failed authentication or the application of the reject role to the client. At this point, it can take some time before the switch attempts to re-authenticate the client. So when a failed link to the datacenter would be restored, it would be convenient that the switch would detect the restored RADIUS communication and would start the authentication again.

This can be achieved with RADIUS tracking. By default, RADIUS tracking will continuously send requests based on a configured interval. As of 10.05 release, this can be configured with a 'dead-only' option, so the RADIUS tracking will only start when the RADIUS server is marked as unreachable based on regular authentication. This prevents unnecessary tracking authentication events on the RADIUS server.

37. On Sw2, configure global RADIUS tracking username and password (radius-tracking-user / aruba123).

38. Configure the tracking interval to 60 seconds

NOTE: This is for lab testing only. The default interval is 5 min.

39. Configure the cppm host entry. Enable tracking and set the mode to dead-only.

40. On Sw2, review the client is still authenticated as 'asts-employee-vsa'. The user 'employee' role has been set to perform re-authentication every 30 seconds.

41. On Sw2, enable debug for radius tracking

42. Use PC1 to open a second SSH connection to the Sw2, this will be referred to as the 'trace session'.

43. In the 'trace session', start a control plane TCP dump of the RADIUS traffic. Leave this session running.

44. On Agg1, apply the RADIUS ACL on port 1/1/7 to block the RADIUS communication again.

45. Switch to the original Sw2 session, review the RADIUS server details.

46. Wait about 30 seconds, the re-authentication should occur and timeout.

47. On Sw2, verify the client now has now failed authentication.

48. In the 'trace session', review the RADIUS packets that are sent by the switch.

Q: What is the User-Name value for the latest authentication requests?

A: The tracking user.

49. On Sw2, review the state in the radius-server details. It should now be reported as unreachable.
50. On Sw2, review the debug buffer for the radius module. You should notice 'DEAD' messages.
51. On Sw2, check the last event log messages that include 'radius'

Recover the link and verify the operation

52. On Agg1, remove the ACL application from the port 1/1/7.
53. Keep an eye on the 'trace session'. Wait up to 1 minute, you should see an 'Access-Accept' message for the radius-tracking-user 'Access-Request'.
54. Close the trace SSH session to Sw2.
55. Verify that after the tracking-user successfully authenticates and that the server is marked 'reachable' again.

NOTE: The Windows supplicant has a default 'Hold' timer that will start a wait period when no response was received on the EAP frames. This can be updated in a production environment using a Windows Group Policy. In this lab, you may restart the 'Wired auto-config' service on PC4 to clear the client timer.

56. Disable all debugging on Sw2

This completes the RADIUS tracking section.

Task 4: Onboarding precedence order

Objectives

- Onboarding order device profile vs aaa
- Fallback role vs reject role

Steps

Customers that are not ready for AAA based authentication can still take advantage of some automatic device access control features of AOS-CX using device profiles.

Device profiles can operate on:

- LLDP frames
- CDP frames
- MAC Address (OUI/Range or single MAC address). This is also referred to as Local MAC-Authentication. (LMA).

In this section you will see how the device profile feature interacts with AAA based authentication.

Device profiles are by default active on ports that have AAA enabled and on non-AAA enabled ports for CDP/LLDP device profiles.

However, for the Local MAC Authentication to work on a non-AAA enabled port, the 'block-until-profile-applied' command must be applied to the port. This command has no effect on a AAA enabled port.

Device Profile with Local MAC Authentication (LMA)

In this section you will first configure LMA only. This means that there will not be any RADIUS MAC-based or 802.1X authentication.

When a device matches the LMA rule (based on client MAC, LLDP or CDP information), the configured profile will be applied.

When there is no match for any rule, a 'fallback' profile can be configured. So the 'fallback' role operates as the 'default' user role for any devices that don't match an LMA rule.

You will consider the PC4 to be a 'printer' and assign the 'printer' user role based on the client MAC address.

1. On Sw2, take note of the PC4 MAC address.
2. On Sw2, disable 802.1X on the port 1/1/4.
3. On Sw2, define a mac-group named 'vm' and match the PC4 MAC OUI.

NOTE: Since PC4 is a Virtual Machine, the MAC OUI will most likely be 00:50:56 based on the Hypervisor platform default MAC range.

4. On Sw2, define a new user role 'printer' and assign it to access VLAN 11.

NOTE: VLAN 11 is used for various roles in these labs, for more advanced traffic control, a policy should be linked to the user role.

5. On Sw2, define a device profile 'dp-printer'. Link the role, the mac-group and enable the profile.
6. On Sw2, review the currently active port-access clients. There should be no active clients.
7. On Sw2, enable device profile debugging and clear the buffer.
8. On Sw2, enable the device profile feature on port 1/1/4. This is only required when there is no other authentication method enabled. (you could consider this the 'enable LMA' command on the interface)
9. On Sw2, verify the device profile was applied.

NOTE: If PC4 does not appear, generate some traffic with 'ipconfig /renew' on PC4.

NOTE: If PC4 does not appear, verify on PC4 802.1X authentication settings that the option 'Fallback to unauthorized network access' is enabled.

10. On Sw2, review the port-access clients

Q: What is the Onboarded Method?

A: The client has been onboarded using the device-profile method.

11. On Sw2, review the debug buffer messages with the INFO severity. You should see that the 'Bypass role' was set to 'printer'.

Device Profile Fallback role

Any device not matching the rules in the device profiles will be assigned the fallback role.

12. On Sw2, define a new role 'asts-fallback', assign it to access VLAN 11.
13. On Sw2, on port 1/1/4, configure the fallback role.
14. On Sw2, remove the PC MAC from the MAC Group.
15. On Sw2, verify the client has now received the fallback role.
16. Review the debug buffer, the client is marked with 'Authentication failed but with role'.

17. Revert: Assign the PC4 MAC to the MAC group again and verify the client is known as 'printer' again.

Device Profile with AAA enabled on the port: Onboarding Precedence

Device profiles can be combined with RADIUS based AAA.

In case the customer wants to provision some devices independent of the RADIUS server, the local devices profiles can be used when the RADIUS server rejects the authentication.

Since this may lead to a potential conflict between the local device profile and the RADIUS decision, the switch has an onboarding precedence. By default the RADIUS decision will have priority over the local device profile.

This section demonstrates that the RADIUS role will overrule the device profile.

You will first enable 802.1X again on the port and then connect the PC4.

18. On Sw2, enable RADIUS 802.1X auth on 1/1/4.
19. Since AAA port-access is now enabled on the port, you may remove the local MAC authentication command. (It is now 'automatically' enabled).
20. Re-authenticate 802.1X on port 1/1/4
21. Review the authenticated clients.

Q: What is the Onboarded Method?

A: dot1x. RADIUS based authentication overrules device profiles by default. This is known as the onboarding precedence. This order can be changes but typically the RADIUS server decision will overrule a local configuration.

Reject RADIUS Response

In the previous section you have seen that a valid RADIUS response will overrule the device profile feature based on the default precedence order.

In this section you will observe the behavior of device profiles when the RADIUS server returns an Access-Reject.

In case the RADIUS server returns a reject, the switch will be allowed to continue the processing of the client using the device profile feature.

22. On Sw2, enable port-access debugging and clear the current debug buffer.

23. On PC4, configure 802.1X authentication with username 'asts-reject'. This account has been configured to return a RADIUS Access-Reject.

24. On Sw2, review the resulting port-access of the client.

Q: Did the client get access to the network?

A: Yes, since it matched the device profile, this access was applied.

25. On Sw2, review the debug buffer for INFO messages. After the 'Access-Reject' RADIUS line, the 'Bypass role' message will be shown.

26. On Sw2, review the client details

Q: What is the Authentication status and the Auth Precedence status?

A: Status is Authenticated, but 802.1X shows : dot1x - Unauthenticated.

Optional: Device Profile with RADIUS Failure

In case there would be no response from the RADIUS server, AOS-CX will first check if there is any match on a device profile, in that case the device profile will be applied.

In case there is no match, the authentication 'critical' role will be applied, if configured.

This may allow for more granular 'critical' role assignments when the RADIUS infrastructure would not be reachable.

27. On PC4, configure authentication with user 'asts-employee-vsa' / aruba123.

28. On Sw2, review the port-access clients, the employee role should be applied.

29. On Agg1, apply the RADIUS ACL on port 1/1/7.

30. On Sw2, clear the debug buffer.

31. On Sw2, verify the clients. The printer role should have been applied after the re-authentication fails (this may take a few moments).

32. On Sw2, review the debug buffer INFO messages. After the RADIUS timeout, the Bypass role will be applied.

33. Cleanup: On Agg1, remove the ACL

34. Cleanup: On Sw2, bounce port 1/1/4 and verify the 'employee' role was applied. Disable all debugging.

Task 5: User-Based Tunneling

Objectives

- Configure User-Based Tunneling with the Aruba Gateway
- Understand the VLAN mapping of tunneled users
- MTU settings with UBT

Steps

Configuring UBT

In this section you will configure the UBT zone to the Aruba Gateway and review some debug logging for the connections.

1. On Sw2, disable port 1/1/4, this ensures you can first review the zone configuration before any authentication is performed.
2. Login to the Aruba Gateway using admin/aruba123
3. On the Aruba Gateway, review the current VLAN list

Q: Is VLAN 12 defined on the Aruba Gateway?

A: Yes.

The tunneled client will be assigned to VLAN 12 on the Aruba Gateway in this lab.

4. On Sw2, define the source IP interface for the UBT connection to the Aruba Gateway.
5. On Sw2, enable ubt debugging for the control plane and hw categories
6. On Sw2, configure the UBT zone.
7. On Sw2, verify the state of UBT.
8. On Sw2, review the debug buffer log for UBT.

Q: What message is reported?

A: There is an error message about 'TNS Client VLAN is not specified'

9. Configure the UBT VLAN, use VLAN 4000 as the client VLAN.

Q: Did this VLAN 4000 exist in the Aruba Gateway?

A: No. This VLAN can be created, but is not required on the Aruba Gateway.

10. On Sw2, review the UBT state. It should now be registered.
11. On Sw2, review the debug output to see the steps to connect to the Aruba GW

12. On Sw2, review the diag-dump output

Q: What is the firmware version of the Aruba Gateway?

A: The active 'SAC Firmware version' will show the current version.

13. On the Aruba GW, review the active nodes and the trace buffer. You should observe 'SW Bootstrap Req' message in the trace.

Review Normal UBT Authentication

In this section, a user will be authenticated and tunneled to the Aruba Gateway. You will review the debug logging to see the normal flow of the connection.

14. On Sw2, reconfigure the local user 'employee' role with the gateway zone, set the gateway role to 'authenticated'. Remove the reauth-period and the VLAN configuration.

15. On Sw2, bounce port 1/1/4.

16. On Sw2, review the port-access clients

17. On Sw2, review the tunneled users.

18. Review the debug buffer to see the normal steps to register the user with the Aruba Gateway.

19. On the Aruba Gateway, review the trace buffer. You should see a 'User Bootstrap Req' request.

20. On the Aruba Gateway, review the tunneled-users.

Q: To which VLAN has the client been assigned?

A: The tunneled VLAN is 4000, this is remarked to VLAN 12 on the Aruba Gateway based on the role.

21. On the Aruba Gateway, review the active users. Note the wired client MAC address.

22. On the Aruba Gateway, review the role derivation for the user

Q: What is the Role Derivation?

A: The role is derived from the Tunneled node instruction: `ROLE_DERIVATION_TUNNELED_USER_ROLE`

Q: What does this mean?

A: This method means that the Aruba Gateway received the role assignment from the switch instruction (using the PAPI session).

23. On the Aruba Gateway, review the VLAN assignment for the role authenticated

Q: What is the VLAN that has been configured on the role 'authenticated'?

A: VLAN12

Q: What would happen when the Aruba Gateway role does not contain a VLAN reference?

A: The client would still be assigned to the UBT VLAN. When this VLAN is defined and allowed on the uplink port by the Aruba Gateway, the client will be connected. In this lab environment, the VLAN 4000 does not exist on the Aruba Gateway, so the client would not be able to reach any resources.

Invalid Gateway Role

In this section you will see what happens when a switch user role refers to a gateway user role that doesn't exist on the Aruba Gateway.

On the switch, you will reconfigure the 'employee' role with a secondary role of 'sales' and review the status and debug output.

24. On Sw2, update the 'employee' role with an invalid secondary role, named 'sales'.

25. On Sw2, review the port-access clients and the detailed output.

Q: What do you observe?

A: The status is fail.

26. On Sw2, review the debug buffer entries for the UBT module.

Q: What do you observe?

A: The user bootstrap fails (nack), with a reason RC_UB_FAIL_AUTH_ENTRY_CREATE_FAILED

27. On the Aruba GW, review the last 10 entries from the error log and review the trace-buffer. The errorlog will report the invalid role name, the trace-buffer will show the failed user bootstrap.

Correct the configuration

28. On Sw2, update the user role with the gateway role 'authenticated'.

29. On Sw2, bounce port 1/1/4 and verify the user status is now success.

MTU for Tunneled users

Due to the GRE tunnel between the switch and the Aruba Gateway, the effective MTU for the clients will be reduced. When the network must support the same MTU for the tunneled clients as for the normal wired

clients, jumbo frame support must be configured on the path between the access switch and the Aruba Gateway.

In this section, you will first use PC3 to explore the default wired MTU, since it is connected to a non-tunneled port.

Next you will use PC4 (tunneled client) to verify the reduced MTU and the impact of the jumbo frames on the tunneled traffic.

30. Let's look at the expected maximum MTU for the ping. Ethernet supports 1500 bytes payload, the IP header takes 20 bytes, the ICMP header takes 8 bytes, so $1500 - 20 - 8 = 1472$ bytes in the ICMP ping payload.

Frame field	Bytes	
Default Ethernet Frame max	1514	
Ethernet Header	14	
Ethernet Payload max	$1514 - 14 = 1500$	
IP Header	20	
ICMP Header	8	
ICMP Payload max	$1500 - 20 - 8 = 1472$	

31. On PC3, ping to the default gateway IP (10.x.11.1) using the maximum MTU for default ethernet segments using the -l option to specify the payload size. Make sure to set the 'do not fragment' option in the command (-f).

32. On PC3, attempt to ping with a payload of 1473, you should see an error message.

NOTE: This PC3 is a VM. The infrastructure between the VM and the lab switch only supports a maximum MTU of 1500, so it would have no effect to change your switch access port MTU in this lab.

Now that you have verified the normal expected MTU on a locally switched port, you will repeat these checks on a tunneled client.

33. On Sw2, verify that PC4 is authenticated and tunneled with the role 'authenticated'

34. Let's look at the expected maximum MTU for the tunneled ping. The switch will first apply the UBT client VLAN tag to the client frame, this adds 4 bytes to the frame. Then the switch will encapsulate the client traffic into GRE (and GRE is transported in IP). The GRE header is 4 byte. Since each switch port has a unique GRE tunnel id, this is represented with the GRE key id of 4 bytes. This results in a maximum client ping payload of 1426 bytes.

Frame field	Bytes	Notes
Default Ethernet Frame max	1514	
Ethernet Header	14	Switch source MAC
Ethernet Payload max	$1514 - 14 = 1500$	
IP Header	20	Switch <> Aruba Gateway
GRE Header	4	
GRE Key	4	Based on switch port id
Tunneled Payload max	$1500 - 20 - 4 - 4 = 1472$	
Tunneled Frame details		
Ethernet Header	14	Original user MAC
VLAN Tag Header	4	User based tunnel client VLAN (4000 in the lab)
IP Header	20	
ICMP Header	8	
ICMP Payload max	$1472 - 14 - 4 - 20 - 8 = 1426$	

35. On PC4, attempt to ping the gateway (10.x.12.1) using a normal ping first to verify your connection.
36. On PC4, ping using the maximum calculated MTU, this should be successful.
37. On PC4, attempt to ping using an ICMP payload of 1427. This should fail.

Configure Jumbo support

38. On the Agg1 and Agg2 switches, you have enabled jumbo support already in a previous lab. The Aruba Gateway has also been enabled for jumbo frame support. These commands were executed on the Aruba Gateway and they are here for your reference only:

```
# Aruba Gateway
firewall jumbo
interface port-channel 1
jumbo
```

39. On the Sw2, configure jumbo frame support on all interfaces and on SVI 1, the source interface for the GRE traffic.
40. On PC4, attempt to ping using maximum expected MTU (ICMP payload of 1472)

Q: Did this work?

A: No

The switch needs to re-establish the connection to the Aruba Gateway for the new MTU to become effective.

- 41. On Sw2, reset the ubt zone
- 42. On Sw2, bounce the port 1/1/4

NOTE: On PC4, you may need to release and renew the IP address.

- 43. On PC4, ping using the maximum expected MTU. This should be successful now.

Default role on the Aruba Gateway

In this section you will review the default role for tunneled users on the Aruba Gateway.

Q: Is it required to define a secondary gateway role in the switch user role?

A: No, this is not required.

Q: What would be the role on the Aruba Gateway when no secondary role is define in the switch user role?

A: The Aruba Gateway has a 'default-tunneled-user' AAA profile. The initial role of this AAA profile is applied for tunneled clients that connect without secondary role.

- 44. On the Aruba Gateway, review the 'default-tunneled-user' AAA profile.

NOTE: In the lab environment, the 'Ligon' role on this AAA profile has been replaced with a custom role 'wired-default'.

These commands were issued on the Aruba Gateway and they are shown for your reference only:

```
#Aruba Gateway
user-role wired-default
access-list session allowall
vlan 12

aaa profile default-tunneled-user
initial-role wired-default
```

- 45. On Sw2, reconfigure the user-role 'employee'. Remove the secondary role.
- 46. On Sw2, review the active tunneled users
- 47. On the Aruba Gateway, verify the assigned role of the client.
- 48. On the Aruba Gateway, review the details of the client.

Q: How is the role derived?

A: The output should show 'how: ROLE_DERIVATION_INITIAL_ROLE'

Task 6: User-Based Tunneling QOS

Objectives

- Apply a Port Access Policy for the user role employee
- Trace traffic and compare inner DSCP to outer GRE tunnel DSCP markings
- Verify Queues and traffic assignment

Steps

Apply a QOS Policy to the Switch user role

In this section you will apply a classifier policy on the 'employee' user role. Although the 'employee' user role forwards the traffic to the UBT on the Aruba Gateway, the switch can still perform local inspection and control of the traffic that is sent over the tunnel.

This allows the administrator to apply QOS policies on the user traffic. The switch will automatically push the inner DSCP markings to the outer GRE tunnel IP DSCP field. This ensures that the transport network between the switch and the Aruba Gateway can also respect the QOS of the user traffic, even when the traffic is inside a GRE tunnel.

The lab steps are:

- Define a class for 'cppm' (10.252.1.23)
 - Define a class for 'any' IP traffic
 - Define a port-access policy that remarks class 'cppm' to CS3, but leaves all other traffic default.
 - Enable this port-access policy in the 'employee' user role.
1. On Sw2, define an IP class 'cppm' for any IP traffic destined to host 10.252.1.23 with the 'count' option.
 2. On Sw2, define an IP class 'any' for any IP traffic.
 3. On Sw2, define a port-access policy named 'pac-employee' (PAC > Port Access)
 - a. Add class 'cppm', apply action DSCP CS3 and local-priority 3
 - b. Add class 'any', no action is required (default action is permit).
 4. On Sw2, enable debugging for the pacpolicy module.
 5. On Sw2, clear the interface statistics of port lag255 and review the queue statistics. After the policy is applied, you will review the queue statistics again.
 6. On Sw2, assign the port-access policy 'pac-employee' to the user-role 'employee'. This change will be applied on the fly by the switch.
 7. On Sw2, review the details of the port-access clients. You should notice the Access policy details with the 'cppm' class.
 8. On Sw2, review the debug buffer entries for the pacpolicy module.

9. On Sw2, review the port-access policy hitcounts

10. On PC4, ping 10.252.1.23.

11. On Sw2, review the hitcounts again.

Q: Do you see the hitcount increase?

A: Yes.

12. On Sw2, review the queue statistics of the LAG255 uplink

Q: Are there any packets handled by Queue 3?

A: Yes, this means that the outer GRE packets were processed in the correct queue.

Optional: Trace traffic and compare inner DSCP to outer GRE DSCP markings

In this optional section you may perform a network trace of the Sw2 uplink traffic and verify that the port-access policy has successfully applied the DSCP mappings to the user traffic.

You will also verify that the inner DSCP markings are pushed to the outer GRE IP header.

13. On PC4, start a continuous ping to 10.252.1.23

14. On Sw2, review mirror session 1. It should be enabled and have destination set to CPU.

15. On Sw2, define a new policy that will mirror the class 'any' to mirror1.

16. On Sw2, apply the policy on the uplink LAG255

17. On Sw2, enable Tshark to dump the CPU mirror to a local file

18. On Sw2, use CTRL-C to stop the capture of packets after about 2 seconds.

19. On Sw2, copy the file with TFTP to the PC1 host using the 'mgmt' VRF

20. On PC1, open the file from the TFTP folder using Wireshark.

NOTE: There is an example trace file for your reference in the ASTS folder: "**ASTS-Lab11-ubt-gre-qos-mark-icmp.pcap.pcapng**".

21. On PC1, look for an ICMP packet to destination IP 10.252.1.23.

Q: What is the DSCP value of the inner IP header?

A: CS3.

Q: What is the DSCP value of the GRE IP header?

A: CS3. This demonstrates that the DSCP value is propagated to the GRE IP header.

22. Look for the ICMP reply packet (source 10.252.1.23)

Q: What is the DSCP value of the inner IP header?

A: The CPPM host sends the ICMP reply using the same DSCP mark as the received ICMP packet.

Q: What is the DSCP value of the GRE IP header?

A: CS3. The Aruba Gateway will also apply this inner DSCP value to the outer GRE tunnel DSCP value.

Cleanup

23. On Sw2, remove the policy in and out from the LAG255

24. On Sw2, disable all debugging.

Aruba Training-Confidential

Task 7: Downloadable User Roles

Objectives

- Review and troubleshoot issues related to Downloadable User Role configuration
- Review and troubleshoot certificate and authentication issues
- Review the normal authentication status

Steps

ClearPass DUR account configuration

1. On Sw2, disable port 1/1/4.
2. On Sw2, enable all debugging for portaccess and pacpolicy (port access policy). Clear the debug buffer.
3. On Sw2, update the radius configuration to support downloadable user roles. The ClearPass credentials you may use are admin / aruba123 .

NOTE: These credentials are invalid, but that will allow you to troubleshoot the connection.

4. On PC4, configure the 802.1X credentials asts-employee-dur / aruba123.
5. On Sw2, enable port 1/1/4
6. On Sw2, review the port-access clients
7. On Sw2, review the port-access client details

Q: What do you observe?

A: The authorization status is 'Not Ready'

ClearPass Certificate Validation

8. On Sw2, review the last 10 events. You should see a certificate verification failure error message.
9. On Sw2, review the debug buffer for errors.

Q: What errors do you notice?

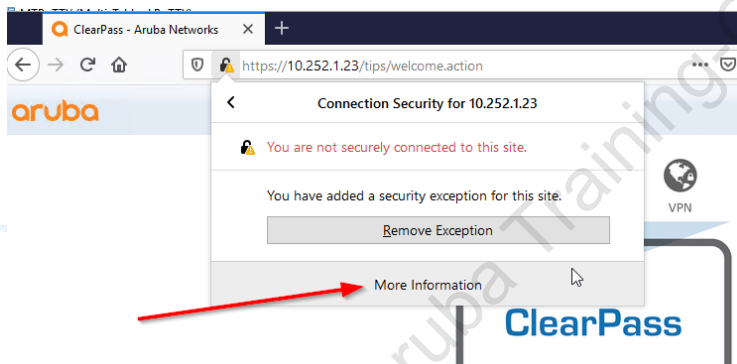

A: There should be an error about the certificate subject name that doesn't match the peer hostname

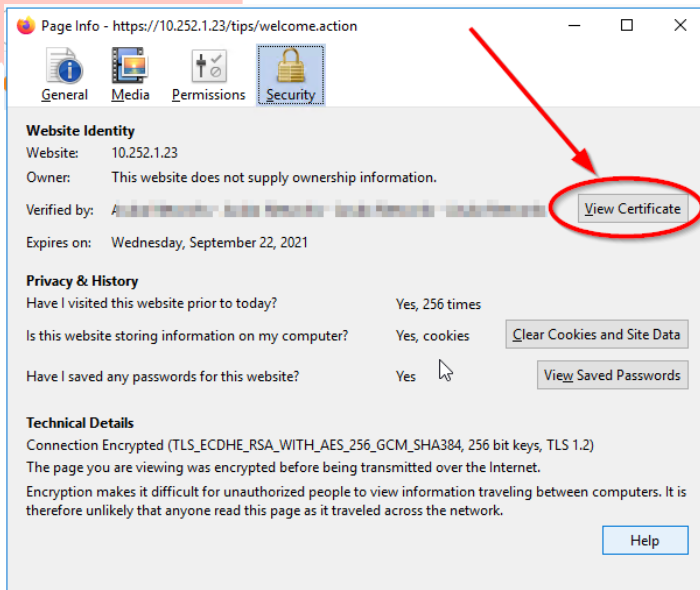
10. On PC1, use Firefox to open an HTTPS connection to the CPPM host, accept the certificate warning. You don't need to login to ClearPass, just access the login page.

`https://10.252.1.23`

NOTE: You could use any browser, the screenshots are based on Firefox.

11. On PC1, review the web certificate details to find the certificate subject name (common name).





Q: What is the Common Name (Subject name) of the installed certificate?

A: This may depend on the remote lab environment. This lab guide assumes it is cppm.arubatraining.com

12. On Sw2, review the RADIUS definition

Q: What is the server name for your RADIUS server in the switch configuration?

A: 10.252.1.23

Q: Does this match the certificate subject name?

A: No, the certificate subject name is based on a DNS name, not an IP address.

13. On Sw2, verify your name resolution works for this hostname. Verify you have DNS host 10.252.1.21 defined and that you can ping the hostname.

NOTE: If there is no DNS server available, it is possible to define a local host entry on the switch, for example: ip dns host cppm.arubatraining.com 10.252.1.23

14. On Sw2, add a new RADIUS server entry with the correct hostname.

NOTE: This is a change compared to the AOS-Switch platform, where an IP address for the RADIUS host was accepted when using Downloadable User

Roles.

15. Assign the new RADIUS entry to the correct group and remove the previous entry.

16. On Sw2, bounce port 1/1/4 and wait for the port-access client to show in the output.

Q: Was the authorization successful this time?

A: No.

17. On Sw2, review the debug buffer. You should observe a message about 'Cert chain failed crypto validation'.

Q: What does this mean?

A: While the subject name matches this time, the switch does not have sufficient information to validate that the certificate was signed by the correct root certificate authority. This requires a 'Trust Anchor' profile in the switch configuration with the root CA certificate that has signed the ClearPass web server certificate.

18. On Sw2, disable port 1/1/4 to stop the authentication attempts.

Install the Trust Anchor certificate on the Switch

In this section you will install the Trust Anchor profile on the switch.

ClearPass will automatically publish the Root Certificate of its own server certificate. You will now connect and download this certificate.

19. On PC1, open a web browser and navigate to this URL.

```
http://10.252.1.23/.well-known/aruba/clearpass/https-root.pem
```

NOTE: This is an HTTP connection, not HTTPS!

NOTE: On PC1, you can also use the **ASTS-Lab11-cppm-get-root-cert.cmd** script. It uses CURL to get the same page contents.

20. On PC1, you can now copy the text from this page.

21. On Sw2, import the root TA certificate to validate the ClearPass web certificate.

22. Paste the copied root cert text and press CTRL-D. Confirm the installation and exit the TA profile.

IMPORTANT: You must 'exit' the context to complete the import process!

23. On Sw2, review the ta profiles

24. On Sw2, review the details of the cppm profile

Now that you have imported the certificate, retry the PC4 authentication

25. On Sw2, clear the debug buffer, enable port 1/1/4 and trigger re-authentication.

26. On Sw2, review the port-access client details.

Q: Did the authorization succeed?

A: No, the status is 'Not Ready'.

27. On Sw2, review the last 20 events.

Q: Was the certificate verified and accepted?

A: Yes, there is an event that confirms that the certificate was verified and accepted. This means that the subject name matches the radius host and that the certificate could be validated based on the TA profile.

28. On Sw2, review the debug log.

Q: What do you observe?

A: The certificate passed the crypto validation, but there is a 'parsing xml failed' error message.

Q: What does this mean?

A: For downloadable user roles, ClearPass returns the internal object name of the ClearPass enforcement policy and version number. The switch attempts to download this using an HTTPS GET. When the XML result does not contain the correct commands, the 'parsing xml failed' message will be shown.

29. **Option1:** Review the downloaded XML file from the /tmp folder on the switch in the shell context.

NOTE: In general, shell access should only be used under supervision of Aruba support.

Q: What do you observe?

A: The XML body refers to wrong credentials to access the server. This indicates that the saved ClearPass credentials are incorrect.

This is the end of Option1.

30. **Option2:** Run CURL to simulate the XML file download. On the PC1, open Desktop > ASTS and run "**ASTS-Lab11-cppm-dur-test.cmd**". This will connect to CPPM on the same URL as the switch would and will show the same output XML.

NOTE: The default script contains the invalid credentials admin/aruba123 to show the error. After you have seen the error, you may edit the script with the credentials duradmin/aruba123 and test it again.

This is the end of Option2.

31. On Sw2, correct the clearpass credentials for the radius host to 'duradmin/aruba123'. (The previous credentials admin/aruba123 were invalid)

Review normal DUR authentication and role

Now that the ClearPass configuration has been updated, you will explore the output of a normal authentication with a downloadable user role.

32. On Sw2, bounce port 1/1/4 and trigger re-authentication.
33. On Sw2, review the port-access clients
34. On Sw2, review the port-access client details
35. On Sw2, review the role details that resulted from the XML file download.

36. Optional step: Define a new checkpoint 'asts-lab11-<yourname>' on the 4 switches.

This concludes the Lab Activity!

Lab 12: Network Security Features

In this lab you will, explore various security features to protect the network and the network devices.

Objectives

- Understand ACL TCAM resources.
- Configure and troubleshoot Control Plane ACLs.
- Understand the Control Plane Policy configuration.
- Configure and troubleshoot DHCPv4 Snooping.
- Configure and troubleshoot ARP Inspection.

Requirements

This lab requires completion of Lab 11.

Scenario

TCAM Resource Usage

The customer plans to use ACL on the switches, but they want to understand how they can monitor the TCAM resource limits and what is the impact when they use object-groups.

You must show the customer how to monitor the TCAM resource usage and configure an example with object-groups to exhaust the TCAM tables and show the error messages.

You can use the file ASTS-Lab12-acl-snippets.txt and apply an example ACL on routed port 1/1/7 of Agg1. Make sure to remove the ACL after the demonstration is done.

Hitcounts

The customer wants to track traffic to the ClearPass host.

1. Configure a policy with a class 'cppm' that counts the packets sent to the host 10.252.1.23.
2. Apply the policy on Sw2 to ports 1/1/4 and 1/1/27.
3. Use ping to test from PC4 to 10.252.1.23 and verify the hit count increases.
4. Observe the hit counts on the policy for the port 1/1/7 are also increasing.
5. Explain what happens and correct the configuration, so that the port 1/1/27 and 1/1/4 each have their own distinct hit counts.

Control Plane Security

The customer has observed that any device on the OOBM subnet may attempt to connect to the Sw1 OOBM IP address. They want to protect the OOBM IP address of the switch and ensure that only the hosts 10.251.x.91 (PC1) and 10.251.x.200 (NetEdit) are allowed to connect to the OOBM IP address. All other devices should be blocked. The block should prevent ping to the Sw1 OOBM IP.

The customer has experienced that it can be very tricky to configure ACLs on a remote system (they locked themselves out of a remote site switch). They want you to demonstrate that the switch can automatically revert the configuration after some time if the configuration is not explicitly confirmed by the administrator. For the demonstration set this to 60 seconds.

Now that the customer has seen that the OOBM IP address can be secured and configurations can be automatically reverted, they notice that the in-band VRF default is still open. PC4 can simply open a web-browser and attempt to open an HTTPS connection to the Agg1 or Agg2 IP addresses.

Limit the in-band access so only the SSH port (TCP 22) and ICMP are allowed on the VRF 'default'. Diagnose and correct any OSPF routing issues and PIM peering issues.

Control Plane Policing

The customer has heard that AOS-CX protects the control plane against loops or too many packets for a control plane process. Demonstrate how:

1. Control Plane Policing works and point out any packets have been dropped for this feature when PC1 attempts to send too many ICMP packets to Sw1 in VLAN1.
 - a. Use the command 'ping 10.x.1.4 -t 1 -w 1 -n 1000' to attempt to send 1000 requests, with an interval of 1 millisecond and a maximum wait time of 1 millisecond.
2. Repeat this test using the OOBM IP and explain what happens.

DHCP Snooping

The customer has had several issues in the past with clients that were running their own DHCP servers. On Sw2, they want to ensure that clients in VLAN 11 can only receive an IP address from the correct DHCP server 10.252.1.22 on the uplink LAG port.

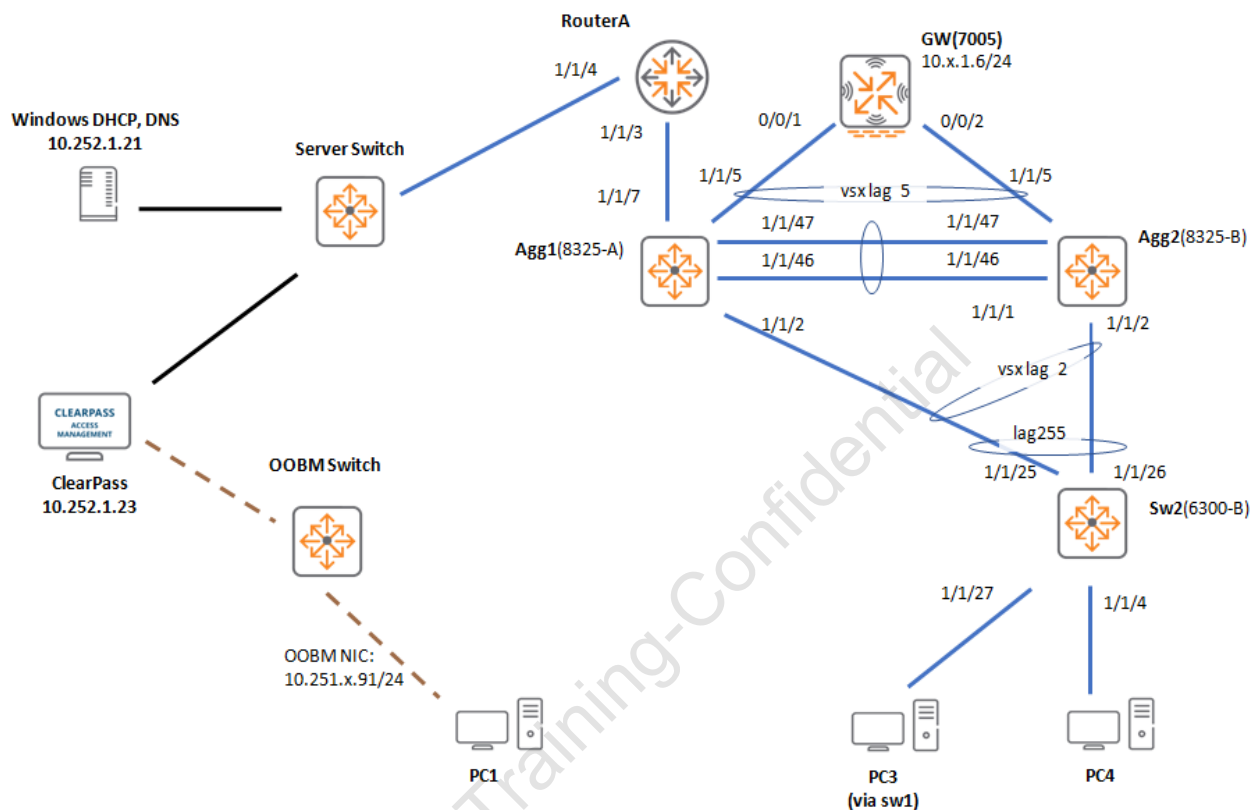
Demonstrate the logs or debug logs showing that DHCP server packets are dropped. The customer also wants to see the list of the snooped client DHCP addresses.

ARP Inspection

Since a security audit has revealed that it is very easy to spoof ARP packets to perform Man-in-the-middle attacks, the customer wants to ensure that only valid ARP packets are entering the network.

On Sw2, they want to ensure that on the client ports (1/1/4 and 1/1/27) ARP packets will only be accepted based on the learned IP address of the port. ARP traffic from clients with a static IP address should be blocked.

Lab Diagram



Task 1: ACL and Resource usage

Objectives

- Understand the resource TCAM usage.
- Understand how hitcounts on policies operate.

Steps

Configuration

In this section you will explore the status of an ACL that reaches the ACL TCAM limit. The switch has extensive TCAM resources, but by using the object-group feature it is possible to quickly exhaust the available resources.

IMPORTANT: While the object group feature is used in this task to reach the limit, this also comes as a warning when using object groups in a production environment. Make sure you understand the impact of using object-groups in ACL statements.

Prepare the a routed ACL

1. On Agg1, review the current resource usage using 'show resources'
2. On Agg1, define a new access-list 'servers', allow all traffic and apply it outbound on 1/1/7.
3. Review the resource usage.

Q: Did the usage change?

A: There was a minor change for the Egress TCAM entries, it incremented with 2.

Update the ACL using object-groups

Tip: On PC1 > Desktop > ASTS folder, you can find these commands in the file 'Lab12-acl-snippets.txt'

4. On Agg1, define an object-group of type 'ip' for 'servers'. Define 10 entries.
5. On Agg1, define an object-group of type 'ip' for the 'clients'. Define 10 entries.
6. Define an object-group of type 'port' named 'srvtcp'. Define 5 ports.
7. Access the access-list 'servers', add a new line
8. Review the resources. It may take 10-20 seconds to update the counters.

Q: How many new entries were added ?

A: 500 new entries were added.

Q: Why are there this many resources in use?

A: The object-group results in: 10 IPs x 10 IPs x 5 Ports = 500 rules, each consuming some entries in the TCAM.

9. Add the same entry to the ACL again (as admin you realize this would have no effect, but the switch will apply the rules anyway)
10. On Agg1, use 'show resources' to verify 500 entries were added.
11. On Agg1, add the line 1 more time. This will exhaust the TCAM resources.
12. On Agg1, review the event log. You should see a TCAM error.
13. On Agg1, review the running configuration of port 1/1/7

Q: What do you observe?

A: The switch reports that the user configuration does not match the active configuration. This is an indication that the switch failed to program the ACL or policy in the hardware.

14. Use the command 'show access-list configuration'. This will also reveal that there is a mismatch between the active and user configuration.
15. Remove the ACL from port 1/1/7

Hitcounts and Statistics

When policies with classes or ACLs contain a 'count' option, the administrator should understand that only a single counter is maintained for the entry, even when the ACL or policy is applied to multiple objects.

In this section you will observe this behavior.

You will use PC4 to send a ping to the RADIUS server. By applying the same policy with a counter on the Sw2 ports 1/1/4 (PC4) and 1/1/27 (PC3), you can observe the counter operation. The lab assumes that the user 'asts-employee-dur' is connected, so traffic is not tunneled to the Aruba Gateway.

16. On Sw2, review your existing classes. You should have a class 'cppm' with a counter enabled on it.
17. On Sw2, define a policy named 'asts-count1' and add the class 'cppm'. No action need to be set.
18. On Sw2, apply the policy on interfaces 1/1/1 and 1/1/27 inbound.
19. On Sw2, review the current hitcounts of the policy.
20. On PC4, send a single ping to 10.252.1.23. This ping should be counted on 1/1/4.
21. On Sw2, review the hitcounts.

Q: Did the hitcounts on the 1/1/4 and 1/1/27 increase?

A: Yes, both increased with one, since they use the same counter.

Optional: Unique Policy per interface for unique counters

In this optional section you will define a new policy and verify that each policy now has its own counters.

22. Define a new policy 'counter2' on Sw2, add the same class 'cpm' to it.
23. Replace the policy 'counter1' on port 1/1/27 with the 'counter2'.
24. Clear the all hitcount statistics
25. On PC4, repeat the test ping to 10.252.1.23.
26. On Sw2, verify that the counters are now unique.

Task 2: Control Plane ACL

Objectives

- Understand how to configure Control Plane ACL per VRF.
- Understand that COPP ACL applies to both management and control protocols.

Steps

In this section you will evaluate the Control Plane ACL feature. This can be used to apply an IP filter to all the traffic that is sent to the control plane.

This feature is configured per VRF and it can also be applied to the 'mgmt' VRF.

In the first section you will explore the 'mgmt' VRF application. In the second example, the default VRF is configured and you will see the impact on control protocols, using OSPF as an example.

COPP ACL on VRF mgmt

By default, there is no COPP ACL, so any device that has IP reachability to the 'mgmt' IP address of the devices could attempt to access it.

1. On Agg1, attempt to reach Sw1 on the 'mgmt' VRF (10.251.x.4) using SSH or a ping. This should succeed.

NOTE: There is no need to actually login, just verify that the device is reachable.

2. On Sw1, configure an ACL named 'copp-mgmt'.
3. Add a rule to permit any traffic from source IP 10.251.x.91/32 (PC1) and 10.251.x.200.32 (NetEdit) and enable the count option.
4. Add a rule to deny traffic from any to any with the count action.
5. Apply the ACL to the control-plane on the VRF mgmt.
6. On PC1, verify you can still access Sw1 using SSH.

NOTE: In case you made a mistake, you can use the remote lab console connection to the devices.

7. On Agg1, attempt to reach Sw1 again on the 'mgmt' VRF.

Q: Did this succeed?

A: The traffic should have been blocked by the Control Plane ACL.

8. Review the hitcounts on Sw1

Q: Where there any hitcounts on the deny rule?

A: Yes, the ping from Agg1 should be counted with the deny rule.

Prepare for failure

In many real deployments, you may not have immediate console access to the devices, so changing the in-band management can be dangerous. AOS-CX has a convenient checkpoint rollback feature that can be used for these situations.

NOTE: Make sure you are connected via PC1 using SSH to Sw1 OOBM IP. If you are connected to the console, you will not observe the disconnect.

9. On Sw1, schedule an automatic rollback in 1 minute. This command cannot be executed in the configuration mode.

10. On Sw1, introduce an error in the mgmt ACL as the first rule (you should be able to do it during one minute)

Q: What do you observe?

A: The SSH session seems to hang since the traffic is blocked now. You need to wait up to 1 minute, then the auto-rollback will revert the configuration.

11. After about 1 minute, use PC1 to open the SSH connection to Sw1 again and review the event log to see the checkpoint messages.

NOTE: To confirm the changes, use the command 'checkpoint auto confirm'. To revert to the saved checkpoint without waiting for the timer, use the command 'no checkpoint auto confirm'.

COPP ACL on VRF default

While the 'mgmt' VRF will only be used for management protocols, the other VRFs on the switch may handle management, control and dataplane traffic.

The Control Plane ACL does not handle dataplane traffic, but it does control all the control plane traffic. This means that both management and control protocols will be verified by the Control Plane ACL.

In this section you will explore the impact on your active control plane protocols.

12. On Sw1, review your existing OSPF and PIM neighbors. There should be several OSPF and PIM neighbors.

13. Define a new ACL named 'copp-vrf-default' on Sw1.
14. Add a rule for any SSH and any ICMP traffic.
15. Apply the policy to the Control Plane of the VRF 'default' on Sw1.
16. On Agg1, attempt to ping and SSH to the Sw1(10.x.41.4) using the default VRF.

Q: Did this succeed?

A: Yes, both ICMP and SSH are allowed in the default VRF

17. Review the hitcounts on the control-plane ACL on Sw1.

18. Now review your OSPF neighbors and PIM neighbors

Q: Are there any OSPF neighbors? Why?

A: There are no more OSPF neighbors, since the OSPF IP protocol is also blocked by the control-plane ACL.

19. On Sw1, clear the OSPF interface statistics and review the SVI41 counters

Q: Do you see any inbound counters increase?

A: No, all inbound packets are filtered by the control plane ACL.

20. Correct this problem in the ACL so OSPF and PIM can operate again.

21. Verify the OSPF and PIM neighbors are now connected.

This concludes the Control Plane ACL Task.

Task 3: Control Plane Policing

Objectives

- Understand Control Plane Policing
- Review the Default policy
- Review the statistics and dropped packets for the Control Plane Policy.
- Understand the default rule of the Policy.

Steps

AOS-CX has a default Control Plane Policy active that will protect the management processor from being overloaded by excessive packets or frames.

It is best practice to leave this policy in place. Changes should only be considered if the default policy would have some negative effect to your production environment.

NOTE: Contact Aruba TAC in case assistance is required.

In this section you will test the default Control Plane Policy.

1. Open a connection to Sw1, review the copp-policy 'default' configuration.

Q: How many icmpv4-unicasts will be accepted per second?

A: 225 packets per second and a burst of 225.

2. Open a connection to Agg1, review the copp-policy 'default'

Q: How many icmpv4-unicasts will be accepted per second? Is that different from the 6300?

A: 1000 packets per second. Yes, every platform can have its own policy limits.

3. On the Sw1, review the copp-policy statistics. Take note how many icmp-unicast-ipv4 packets have passed and have been dropped.

NOTE: During the VSX and Layer2 lab activity you have configured a loop. You may still see many arp-broadcast drops. You may not have realized it at that time, but the default Control Plane Policing protected the switch during that loop.

You will now use fping (fast ping) on PC1 to attempt to overload the Sw1.

4. On PC1, open a command prompt and send 1000 pings, 1 every ms (-t) and wait maximum 1ms (-w) to the VLAN1 IP address of Sw1 (1.x.1.4).

Q: What do you observe?

A: The initial 450 packets get responses thanks to the 'burst' configuration. Then the limiter starts and only a limited number of responses are processed. The exact number will vary based on the speed of the remote lab VM systems. Repeat the command a few times if necessary.

5. On Sw1, review the copp-policy statistics. You should observe icmp-unicast-ipv4 drops in the output.

VRF mgmt behavior

While the policy is effectively handled by the hardware ASIC, this also means that the rate limiter has no effect on the 'mgmt' VRF of the switch.

6. On PC1, repeat the fping command but use the 'mgmt' IP (10.251.x.4).

Q: What do you observe?

A: There are no packets dropped since there is no limiter on the 'mgmt' VRF.

NOTE: In case an administrator has made changes to the 'default' COPP Policy, the changes to the policy can be easily reverted using 'copp-policy default revert'.

Optional Task 4: Monitoring Control Plane with NAE

Objectives

- Use NAE to monitor the COPP Policy

Steps

This task will show how you can use NAE to monitor the Control Plane Policy statistics.

1. On PC1, look for the file '**ASTS-Lab12-copp.3.1.py**' file in the ASTS folder.
2. Use PC1, open a web connection to Sw1 (10.251.x.4)
3. Upload the copp.3.1.py NAE script and define an Agent named 'copp-agent'. You can use the default parameters.
4. In the agent details, in the graph window, use the 'Configure Chart' icon to modify the items you want to see in the graph.
5. Select 'Custom Monitoring'
6. Disable the default Monitor items (uncheck them)
7. Enable **icmp-unicast-ipv4 traffic passed** and **icmp-unicast-ipv4 traffic dropped** and click save.
8. On PC1, run the fping to the VLAN 1 IP address of Sw1 (10.x.1.4)
9. Observe how NAE shows the graphs after a few seconds.

Task 5: DHCP Snooping

Objectives

- Understand how to configure DHCP snooping
- Understand the trusted and untrusted ports

Steps

DHCP Snooping is the recommended solution when the network must be protected against unauthorized DHCP servers.

The feature can be enabled per VLAN.

When it is enabled, every DHCP packet will be screened by the snooping process. DHCP client packets can only be forwarded to ports that have been configured as trusted ports. By default all ports are untrusted.

PC4 will be your test PC in VLAN 11 to test the DHCP snooping options.

1. Open a connection to Sw2.
2. On Sw2, review the default dhcpv4 snooping settings.
3. On Sw2, enable DHCPv4 snooping globally and on VLAN 11

TIP: If you would need to enable this on several VLANs, a VLAN range can be used.

4. Open a connection to PC4, release your ip address.
5. On PC4, attempt to renew your IP address.

Q: Did this work?

A: No, the client did not get an IP address anymore.

6. On Sw2, review the dhcpv4 snooping statistics.

Q: What do you observe?

A: Client packets were dropped.

Q: What is the reason for the drop?

A: The reason states: destination on untrusted port.

7. On Sw2, enable DHCPv4 debugging and clear the debug buffer.

8. On Sw2, make the LAG255 uplink trusted for DHCP snooping.
9. On PC4, attempt to renew your IP. This should now succeed.
10. On Sw2, review the DHCPv4 snooping statistics.
11. On Sw2, review the debug buffer to see the normal DHCP process. You should observe the discover/offer/request/ack phases of the DHCP process.

Optional: Authorized Server

DHCPv4 can also be configured to check the source IP on the uplink connection.

12. Start with a mistake: Configure 10.252.1.22 as an authorized DHCP server on Sw2.
13. On PC4, attempt to release and renew the IP address, this should fail.
14. On Sw2, review the statistics and note the Reason for the server drop.
15. On Sw2, review the event log. A WARN packet should be observed.
16. On Sw2, review the debug logs. You should see a message with '
|DHCPV4_BOOTP_REPLY packet was dropped at port:lag255.Reason - Reply packet came
from unauthorized server:10.252.1.21'
17. On Sw2, add the correct DHCP server with IP 10.252.1.21.
18. Verify that PC4 can now obtain an IP address using DHCP.

DHCP Bindings Database

Based on the DHCP Snooping, the switch will also learn the client MAC and client IP.

19. On Sw2, review the DHCPv4 Bindings database

Q: What would be other uses for this database?

A: This binding database can be used by other security features, such as ARP Inspection or IP source lockdown to verify the correct L2 MAC to L3 IP relation for a given switch port.

20. Disable the debugging

Task 6: ARP Inspection

Objectives

- Configure ARP Inspection

Steps

Within the client VLAN, endpoints could use simple tools to spoof the ARP tables of the routing switches and other clients on the local subnet.

ARP Inspection can be used to verify the incoming ARP packets so only valid ARP packets will enter the network.

ARP inspection can be enabled per VLAN. All ports in the VLAN will be considered untrusted by default, meaning that the incoming ARP packets will be verified by the ARP Inspection.

The switch will need a database to know on which switch port an IP/MAC combination is valid or invalid. The DHCP Snooping database will be used for this verification process.

Devices with a static IP will not be learned by the DHCPv4 snooping, so they will not be in the DHCPv4 snooping database. The result will be that the ARP packets send by these systems will also be rejected.

These ports can be configured as 'trusted' ports, so incoming ARP packets on these ports will not be verified by the ARP Inspection feature.

You will use PC4 in VLAN 11 to test the feature.

Configure ARP Inspection

1. On PC4, verify that you have received a DHCP address in VLAN 11.
2. On PC4, verify that you can ping to the default gateway on 10.x.11.1.
3. On PC4, review the arp table, check the entry for 10.x.11.1
4. On Sw2, review the DHCPv4 snooping binding database. PC4 should be present with its MAC and IP address on the port 1/1/4. This will be used to check incoming ARP packets.
5. On Sw2, enable ARP Inspect on VLAN 11.
6. On Sw2, enable arp inspection debugging
7. On PC4, clear the ARP cache using the command 'arp -d *', next attempt to ping the gateway on 10.x.11.1.

NOTE: Make sure you are using an elevated command prompt as administrator for the arp command.

Q: Did this work?

A: No, there is no response.

8. On PC4, review the ARP cache again, there should be no ARP entry for the gateway.

9. On Sw2, investigate the ARP Inspect output, review the arp inspect statistics for VLAN11.

Q: What do you observe?

A: There are several dropped entries.

10. On Sw2, review the debug buffer

Q: What do you observe?

A: Validation failed for arp packet on VLAN 11, untrusted port lag255.

11. Review the ARP Inspect interface output on Sw2.

Q: What is wrong here?

A: The client gateway (VSX Agg) is connected on port LAG255, the IP/MAC will not be learned by DHCPv4 snooping, but the port is by default set to 'Untrusted', so ARP packets are dropped.

12. Configure the uplink port on Sw2 as trusted for ARP Inspect. ARP packets will now be accepted on the LAG255 without further inspection.

13. On PC4, you should now be able to reach the network again.

14. On Sw2, review the debug buffer again to see the normal operation.

Optional: Verify IP Binding database dependency

The ARP Inspection feature relies on the IP Binding database to verify the ARP packets. When the IP Binding database is not populated, the ARP Inspect will not find a match, so the ARP packets will be dropped. This means that DHCP Snooping configuration must be enabled on the same VLAN as the ARP Inspect VLAN to be effective.

In this section you will disable the DHCP Snooping on VLAN 11. ARP Inspection will then blocks ARP traffic since there is no match in the IP Binding database.

15. On Sw2, review the IP Binding database.

16. On Sw2, disable DHCP snooping for VLAN 11. Verify the IP Binding database is empty.

17. On Agg1 and on Agg2, clear the arp cache, since the PC4 entry is still present in this cache.

18. On Sw2, clear the debug buffer.

19. On PC4, clear the arp cache and ping to the gateway.

Q: Did the ping succeed?

A: No, the ARP Inspect seems to be correctly configured, but since the IP Binding database is not populated, ARP packets are dropped.

20. On Sw2, review the debug buffer.

This demonstrates the dependency on the IP Binding database.

21. On Sw2, enable DHCP snooping again for VLAN 11

22. On PC4, release and renew the IP address and verify you can ping the gateway.

Verify ARP Inspection filtering

Now you will verify that only DHCP enabled clients are able to send ARP packets.

You will do this by configuring PC4 with a static IP address. Since the static IP is not learned by ARP Inspect, the ARP traffic of the PC will be dropped.

23. On PC4, configure the LAB Nic with a static IP 10.x.11.200/24 and GW 10.x.11.1.

24. On PC4, attempt to ping 10.x.11.1. This should fail.

25. On Sw2, review the debug buffer.

Q: What do you observe?

A: a WARN message will be shown 'Validation failed for arp pkt'

26. This demonstrates that only clients that have passed the DHCPv4 process will be able to use the ARP process to access the network.

Endpoint with static IP

In case an endpoint in the VLAN, such as a printer for example, has a static IP address, the port can be configured as trusted. Since PC4 now has a static IP address, the port 1/1/4 can be used to test this scenario.

27. On Sw2, configure port 1/1/4 as ARP Inspect trusted port.

28. On PC4, you should now be able to access the network.

29. Cleanup: Disable debugging, make port 1/1/4 untrusted again.

30. On PC4, make the Lab NIC a DHCP client again and verify you can ping the gateway.

31. Optional step: Make a new checkpoint asts-lab12-<your name>.

This concludes the Lab Activity.

Aruba Training-Confidential

Lab 13: IPv6

In this lab you will, explore a basic IPv6 deployment.

Objectives

- Configure OSPFv3.
- Configure SLAAC.
- Configure DHCP stateless address assignment.
- Configure DHCP stateful address assignment.

Requirements

This lab requires completion of Lab 12.

Aruba Training-Confidential

Scenario

The customer wants to start with a pilot IPv6 deployment. They will be using the IPv6 prefix 2001:db8:x::/48 for the address assignments. Demonstrate the following IPv6 features.

OSPFv3

Configure the loopback IPs for:

- Agg1: 2001:db8:x:0::2/128
- Agg2: 2001:db8:x:0::3/128
- Sw1: 2001:db8:x:0::4/128

Configure OSPFv3 between Agg1, Agg2 and Sw1. Make sure there are no DR/BDR elections on these links.

You must configure Agg1 and Agg2 with a global IPv6 address for SVI1 and enable it for OSPFv3:

- SVI1 (Agg1: 2001:db8:x:1::2/64 Agg2: 2001:db8:x:1::3/64)

The links SVI41 and SVI42 between Sw1 - Agg1 and Sw1 - Agg2 should not have a global unicast address but still allow for IPv6 routing.

SLAAC

Configure Agg1 and Agg2 SVI11 with:

- Agg1: 2001:db8:x:11::2/64
- Agg2: 2001:db8:x:11::3/64

This subnet should be advertised in OSPF, but Agg1 and Agg2 should not be able to establish an adjacency on this subnet.

Configure a VSX active gateway on SVI11 for IPv6 2001:db8:x:11::1 using MAC address 02:00:00:00:00:01.

Ensure that clients on SVI11 receive an IPv6 address based on SLAAC. The customer wants to see the SLAAC address assignment on PC3 (connected on Sw2 1/1/27).

DHCPv6 - Stateless

Show the customer DHCPv6 stateless address configuration on SVI12. Sw1 should be configured with a DHCPv6 pool and a DNS host (note: this DNS host is not active).

- range 2001:db8:x:12::100 2001:db8:x:12::200 prefix-len 64
- dns-server 2001:db8:252:1::21

Configure Agg1 and Agg2 SVI 12 for stateless DHCPv6 on prefix 2001:db8:x:12::/64. Agg1 and Agg2 should be configured as DHCPv6 relay, using the Sw1 loopback IPv6 address as the DHCP server address.

Assign Sw2, port 1/1/4 (connected to PC4) to VLAN12. PC4 should receive:

- an IP address in the 2001:db8:x:12::/64 subnet
- a DNS server address of 2001:db8:252:1::21

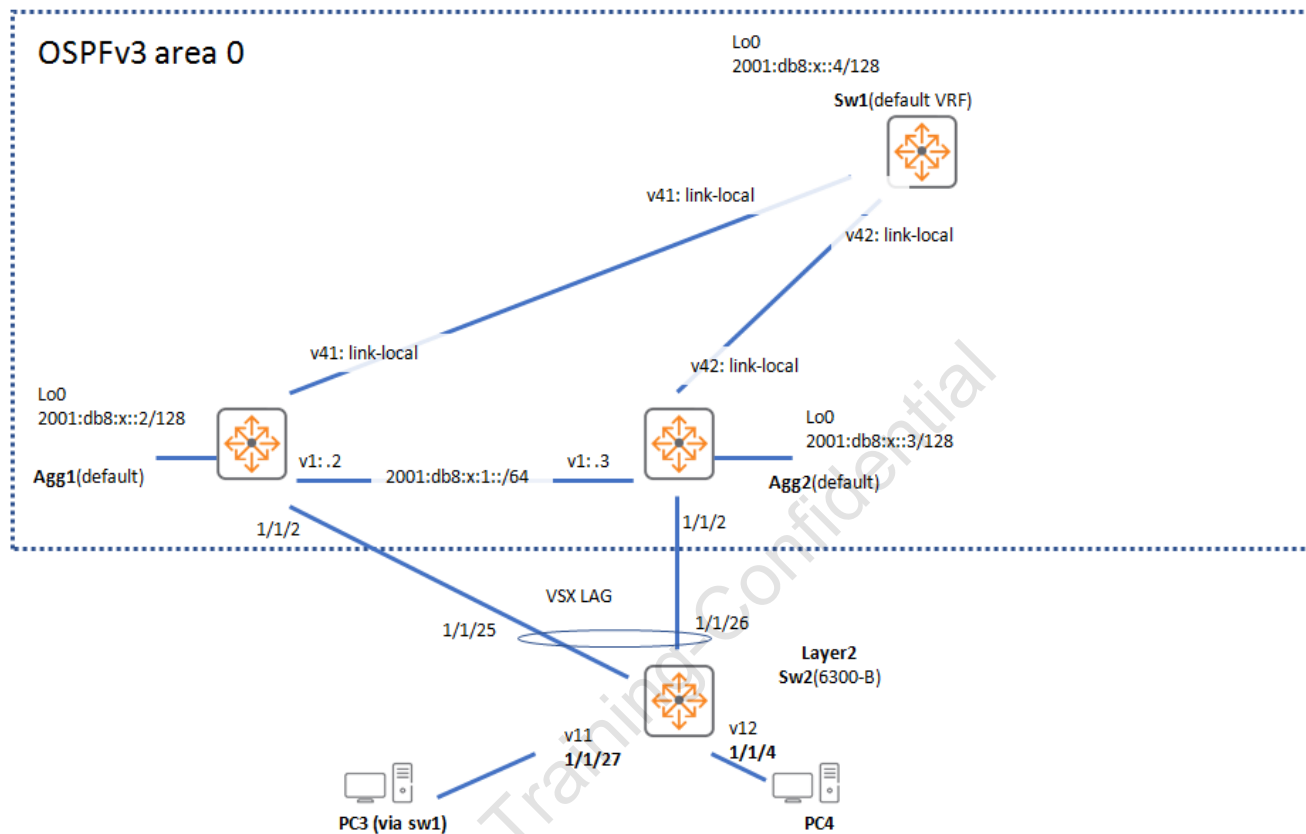
DHCPv6 - Stateful

Although the customer likes the concept of stateless addressing, they are still considering tracking the active client IPv6 addresses and to have a centralized list of active systems. Demonstrate how to change the SVI12 configuration to stateful DHCP services.

They want to be able to see the assigned IPv6 addresses on the Sw1 DHCP server.

Aruba Training-Confidential

Lab Diagram



Task 1: Configure IPv6 links and OSPFv3

Objectives

- Loopback Interface
- Interface configuration
- OSPFv3 configuration
- Passive interface (default)

Steps

You will start this task by configuring a basic OSPFv3 connection between Agg1, Agg2 and Sw1. First you will prepare the IPv6 loopback interface and SVI interfaces. Next, OSPFv3 will be configured and enabled.

Agg1 - Agg2 Link

1. On Agg1, configure loopback 0 with IPv6 address 2001:db8:x:0::2/128.
2. Configure SVI 1 with IPv6 address ipv6 address 2001:db8:x:1::2/64.
3. Review your IPv6 interfaces.
4. Review the IPv6 interface VLAN1.

Q: What is the status of your IPv6 address?

A: After the Duplicate Address Detection has completed, the status will be reported as VALID

5. On Agg2, configure the loopback0 with 2001:db8:x:0::3/128.
6. On Agg2, configure SVI 1, use 2001:db8:x:1::3/64 as the host IP.
7. On Agg1, verify connectivity using ping6 to Agg2 SVI1 IPv6 address. This should be successful.
8. On Agg1, review the IPv6 neighbor list. This should reveal the Layer2 address of the neighbor.
9. On Agg1, clear the neighbor cache.

NOTE: This command applies to both IPv4 and IPv6 neighbor entries.

10. On Agg1, verify the neighbor list is empty.
11. On Agg1, repeat the ping and verify it was learned in the IPv6 neighbor list.

OSPFv3 on Agg1 and Agg2

In this section you will configure OSPFv3 between Agg1 and Agg2 on SVI1. The router-id is a 32bit number. In this lab, the IPv4 Loopback address will be used as the OSPFv3 router-id.

12. On Agg1, configure OSPFv3 with area 0. Assign the router-id to the IPv4 loopback0 address.
13. On Agg1, enable Loopback0 and SVI1 for OSPFv3.
14. On Agg2, configure OSPFv3 with area 0. Assign the router-id to the IPv4 loopback0 address.
15. On Agg2, enable Loopback0 and SVI1 for OSPFv3.
16. On Agg1, review the OSPFv3 neighbors.

Q: What is the IPv6 neighbor address ? Does it match the configured IPv6 address?

A: No, it does not match the configured IPv6 address. OSPFv3 uses the link-local address for the adjacency with the neighbors.

17. On Agg1, review the OSPFv3 routes.
18. On Agg1, review the IPv6 routing table.
19. Check the event log. Look for events generated by the hpe-routing daemon. This should also show the link-local address for the Adjacency setup.

Configure Sw1 as OSPFv3 router

In this section you will configure Sw1 as OSPFv3 router connected to Agg1 with SVI 41 and Agg2 with SVI 42.

These SVI interfaces will be enabled for IPv6, but they will only use the link-local IPv6 address, so no global unicast address needs to be configured on them.

20. On Sw1, configure loopback 0 with IPv6 address 2001:db8:x:0::4/128.
21. On Sw1, enable SVI 41 and SVI 42 for IPv6.
22. On Sw1, configure OSPFv3 with area 0. Assign the router-id to the IPv4 loopback0 address.
23. On Sw1, enable Loopback0 for OSPFv3.
24. On Sw1, enable SVI 41 and SVI 42 for OSPFv3 and set the network type to point to point.
25. On Sw1, review the OSPFv3 interfaces.
26. On Agg1, configure the link to Sw1: enable SVI 41 for IPv6 and OSPFv3.
27. On Agg2, configure the link to Sw1: enable SVI 42 for IPv6 and OSPFv3.
28. On Sw1, review the OSPFv3 neighbors, both Agg1 and Agg2 should be listed.

29. On Sw1, review the IPv6 routing table, you should see the Loopback IPv6 addresses of Agg1 and Agg2.
30. On Sw1, attempt to use ping6 to reach the Agg1 IPv6 loopback address using your local IPv6 loopback address as the source. This should be successful.

Aruba Training-Confidential

Task 2: Configure IPv6 on client subnet using SLAAC

Objectives

- Configure VSX Active Gateway
- Configure the endpoint subnet for SLAAC
- Verify client connectivity

Steps

In this section you will configure VLAN 11 for the PC3 client host. The Agg1 and Agg2 will be configured to provide the clients with an IPv6 address based on SLAAC.

Aggregation configuration

1. On Agg1, configure SVI 11 with IPv6 address 2001:db8:x:11::2/64.
2. On Agg1 SVI11, enable OSPFv3 area 0 and set it as OSPFv3 passive interface.
3. On Agg1 SVI11, configure VSX Active Gateway for SVI 11.
4. On Agg1 SVI11, enable the IPv6 ND Router Advertisements (no ipv6 nd suppress-ra)
5. On Agg2, configure SVI 11 with IPv6 address 2001:db8:x:11::3/64.
6. On Agg2, complete the SVI 11 configuration: OSPFv3 and RA.

Client verification

7. On the PC3, review the IPv6 address. The client should have received an IPv6 address in the SVI 11 subnet 2001:db8:x:11::/64. Make sure, that IPv6 checkbox is marked in Lab NIC interface's properties.
8. On the PC3, verify IPv6 connectivity with the Agg1 and Agg2 IPv6 loopback addresses using ping -6.
9. On PC3, use tracert -d -6 to the Sw1 IPv6 Loopback (2001:db8:x::4)
10. On Agg1, review the IPv6 neighbors. You should observe the PC3 in VLAN11 on LAG1.

Task 3: Configure IPv6 on client subnet using DHCP

Objectives

- Configure DHCPv6 Server and DHCPv6 relay
- Configure the endpoint subnet for DHCP Stateless
- Configure the endpoint subnet for DHCP Stateful

Steps

In this section you will configure SVI 12 for the PC4 client using DHCPv6.

DHCPv6 can be used in 'stateless' mode. This means that the client will locally generate an IPv6 address based on the RA, but the client will still send a DHCP message to discover options. This allows the network administrator to centralize common options such as the DNS server IP address.

DHCPv6 can also be used in 'stateful' mode. This means that the client will not generate an address based on the RA, but it will request an IPv6 address from the DHCPv6 server. The DHCPv6 server will have an overview of the connected IPv6 addresses on these subnets.

First you will configure Sw1 as a DHCPv6 server with a pool for VLAN 12.

Next you will configure Agg1/Agg2 SVI 12 with a DHCP relay option to Sw1.

The SVI12 will then be configured for stateless DHCPv6. After this has been verified, stateful DHCPv6 will be configured.

Configure PC4

PC4 is currently connected to VLAN11. Configure it in VLAN 12

11. On Sw2, change the port VLAN id of port 1/1/4 to VLAN 12
12. On PC4, change the 802.1X login user to username asts-accept (aruba123) and check that IPv6 is enable under Lab NIC interface's properties.

NOTE: The RADIUS server will return a basic 'accept' for this account. Since no other settings are provided, the client will be assigned to the port VLAN id, in this case VLAN 12.

13. On Sw2, review the MAC addresses on port 1/1/4. The client MAC address should now be learned in VLAN 12.
14. On PC4, use ipconfig to verify it has received an IPv4 address in VLAN 12.

Configure DHCPv6 Server

In this section you will configure a DHCPv6 server on Sw1 with a pool for VLAN 12.

15. On Sw1, enter the DHCPv6 context in the default VRF, make it authoritative and define a pool for VLAN 12
16. On Sw1, define a range range 2001:db8:x:12::100 2001:db8:x:12::200 for the pool 12 and set the DNS server address option.

NOTE: The DNS address is an example in this lab. There is no real IPv6 DNS server active in the lab.

17. On Sw1, exit the pool level and enable the DHCPv6 server.
18. On Sw1, review the settings.

Configure SVI12 for Stateless DHCPv6

In this section, Agg1 and Agg2 SVI 12 will be configured to support IPv6 with the Stateless DHCPv6 option.

19. On Agg1, enable DHCPv6 relay globally
20. On Agg1, set the IPv6 SVI address, and enable OSPFv3
21. On Agg1 SVI 12, enable DHCP relay. The destination address is the Sw1 Loopback IP address.
22. On Agg1 SVI 12, configure the stateless DHCPv6. This is done by setting the 'OTHER' option flag in the RA messages. The ND Prefix is still required, since the client will generate an address in this range.
23. Review the DHCPv6 status.
24. On Agg2, repeat this configuration

Verify Stateless DHCPv6

25. On Agg1, enable dhcp-relay debugging and clear the buffer.
26. On PC4, disable and enable the 'Lab NIC'.
27. On PC4, review the address details. The IPv6 address should be a generated address in the 2001:db8:x:12::/64 range. The DNS Server list should now also include the IPv6 DNS address for 2001:db8:252:1::21.
28. On Agg1, review the debug buffer. You should see the DHCPv6 relay flow.
29. On Sw1, review the DHCPv6 lease database.

Q: Are there any active leases?

A: No, since the client is in stateless mode, it will not request an IP address from the DHCP server, only the options.

30. On Sw1, use the diag-dump to see DHCPv6 server details.

Q: What message type to you see?

A: The client has used a DHCPINFORMATION-REQUEST to obtain the option information.

Configure SVI 12 for Stateful DHCPv6

In this section you will update the SVI 12 on Agg1 and Agg2 so the clients will operate in stateful DHCPv6 mode.

31. On Agg1 SVI 12, set the default option for prefixes to no-advertise.

32. On Agg1 SVI12, enable the stateful DHCP with the 'MANAGED' flag in the RA options.

33. On Agg2, repeat this configuration on SVI 12.

Verify Stateful DHCPv6

In this section you will verify the stateful DHCPv6 operation on PC4.

34. On PC4, disable and enable the 'Lab NIC' interface.

35. On PC4, verify the address. You should now see an IPv6 address in the range 2001:db8:x:12::100-200.

36. On Sw1, review the DHCPv6 lease database

37. On Sw1, review the diag-dump of the DHCPv6 server. There should be 4 packets.

Q: What packet types do you see?

A: DHCP SOLICIT, DHCP ADVERTISE, DHCP CONFIRM and DHCP REPLY messages

38. On Agg1, review the debug buffer. You should see the same 4 packets being exchanged between the client and the DHCP server, resulting in 8 messages in total.

39. On Agg1, disable the debugging.

40. Optional step: On all devices, make a new checkpoint 'asts-lab13-<your name>'.

This concludes the basic IPv6 configuration.

Aruba Training-Confidential

Lab 14: Trouble Tickets

In this lab activity several trouble tickets are introduced in the lab environment.

Aruba Training-Confidential

Task 1: Prepare the Troubleshooting Setup

Introduction

The troubleshooting tickets use the final state of the training labs as the base for the trouble tickets. You must have completed at least Lab01. In Lab01 you should have configured the 'oobm' checkpoint and you should have run the 'ASTS-Lab01-deploy-all-configs.cmd' file.

This lab assumes you have completed at least this Lab01 successfully, if this is not the case, you should perform Lab01 again.

The tickets in this lab will always start by loading the '**asts-lab13-done**' checkpoint (saved by the deploy all script in Lab1), and then apply some additional configuration, or introduce some configuration issues into that topology.

NOTE: While you could 'troubleshoot' by checking the differences between the current configuration and the saved checkpoint, you are encouraged to use troubleshooting skills and techniques using show, debug and diagnostics commands.

NOTE: This guide only contains the Tickets information. You can open the document 'ASTS - Lab - 14 - TroubleTickets Hints and Solutions' if you need some example show or debug commands. The Solution Guide contains the actual solution information and detailed show or debug output. Make sure to consult the Solution Guide if you need additional information.

1. On **RouterA**, load asts-base-l3 checkpoint to the running configuration, save the configuration and reboot.
2. On Agg1, verify there is a checkpoint named '**asts-lab13-done**'. Repeat this check on Agg2, Sw1 and Sw2.
3. On all 4 switches, load the asts-lab13-done checkpoint, save the configuration and reboot.

```
copy checkpoint asts-lab13-done running
write mem
boot system
```

NOTE: When loading a new Ticket, the PC may seem to lose its network access. Typically, renewing the IP address (ipconfig /release and /renew) or resetting the Lab NIC solves this issue.

NOTE: If you see any 'Failure' message when loading checkpoint or during the deploy script, stop working and reboot the switch. After the switch is rebooted, you may try to load the configuration or ticket again.

It may also help to manually copy the checkpoint that causes issues to the startup config and reboot. Example (Make sure you **do not save the running-configuration** when asked)

```
copy checkpoint asts-lab13-done startup  
boot system
```

IMPORTANT: Make sure you do not save the running configuration when rebooting!

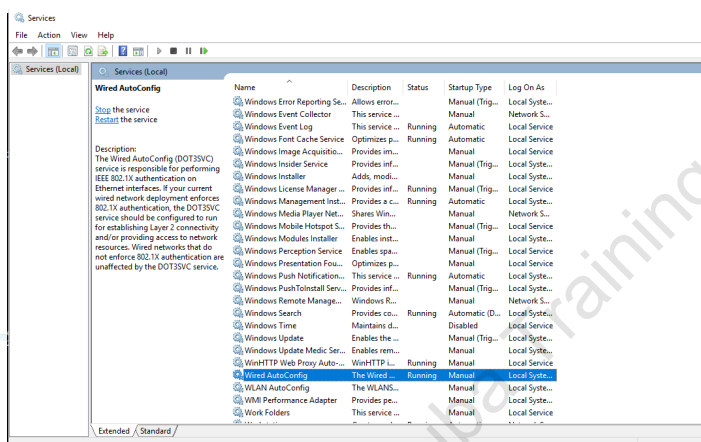
Aruba Training-Confidential

Trouble Ticket Setup 01

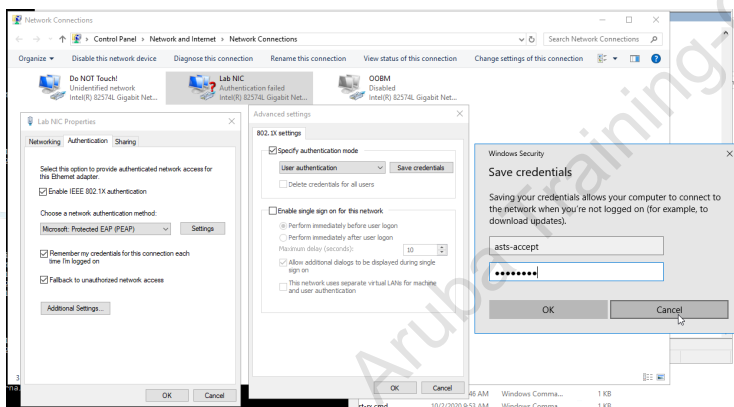
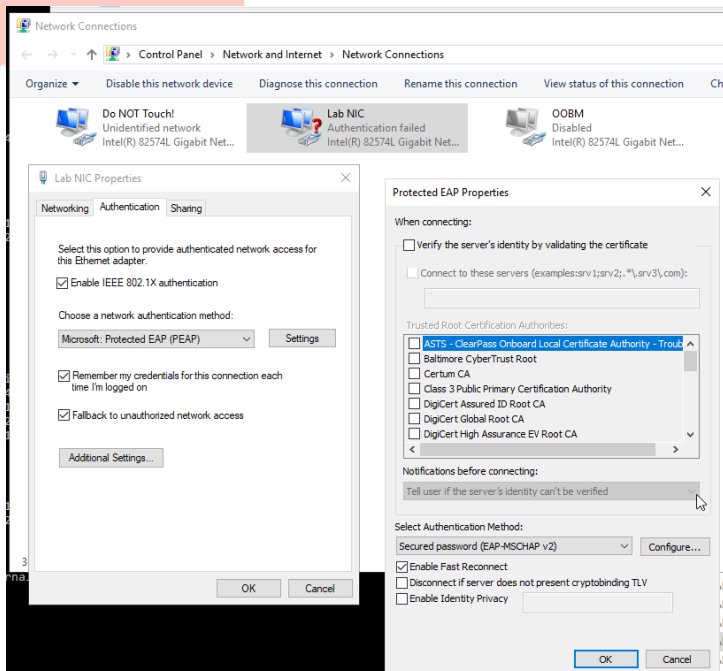
Steps

1. On PC1, on the desktop, open folder ASTS. Run **ASTS-Lab14-TT01.cmd**
2. PC4 should authenticate with 802.1X user asts-accept / aruba123. This should be successful (no issues). PC4 should be able to ping 10.252.1.21 (no issues).
3. In case you did not complete Lab11, these screenshots show an example of the PC4 802.1X configuration.

NOTE: Make sure to enable and start the 'Wired autoconfig' Windows service on the PC4. This is the Windows 802.1X supplicant.



TIP: In this lab environment, the client does not trust the CPPM RADIUS EAP certificate, so you should disable the certificate validation in the Windows NIC 802.1X authentication settings.



NOTE: Make sure the option 'Fallback to unauthorized network access' is enabled.

Trouble Ticket 01.01

4. The customer reports that when Sw2 uplink to Agg1 fails, PC4 can no longer reach 10.252.1.21.
5. On PC4 start a continuous ping to 10.252.1.21.

```
ping 10.252.1.21 -t
```

6. When the uplink to Agg1 is down, when PC4 attempts to re-authenticate, it is no longer possible to authenticate. You can disable/enable the Lab NIC on PC4 to test this re-authentication issue.
7. To simulate the issue: On Sw2, shutdown port 1/1/25 (to Agg1), the ping on PC4 should stop.

```
# Sw2  
interface 1/1/25  
shutdown
```

8. The customer wants you to find the issue and correct the issue.

Trouble Ticket 01.02

Note: This ticket requires Issue 01.01 to be resolved first.

9. PC4 can reach 10.252.1.21 with a ping, but packet drops are observed when the ping rate is increased.
10. On PC4, open a command prompt and ping 10.252.1.21. This should work.
11. On PC4, in the command prompt, run **fping 10.252.1.21 -t 50 -w 10 -c**. This will show intermittent drops.

The customer wants you to identify what is wrong and correct the configuration so that the requested fping has no drops.

Trouble Ticket Setup 02

Steps

1. On PC1, on the desktop, open folder ASTS. Run **ASTS-Lab14-TT02.cmd**
2. PC3 should be able to ping 10.252.1.21 (no issues).

Trouble Ticket 02.01

3. In a previous POC, the customer has seen in that the traffic failover for a VSX LAG is very fast. They have now configured their own network. When they attempt to test this failover between Sw2 and the Agg1/Agg2 it seems to take several seconds before the traffic is working again.
4. The customer is using this test procedure:

PC3 has a continuous ping running to 10.252.1.21.

On Sw2, disable 1/1/25 > verify slow failover of traffic on PC3> restore the link 1/1/25 and wait about 10 seconds.

On Sw2, disable 1/1/26 > verify slow failover of traffic on PC3> restore the link 1/1/26 and wait about 10 seconds.

5. This is an example of the slow ping failover

```
Reply from 10.252.1.21: bytes=32 time=3ms TTL=124
Reply from 10.252.1.21: bytes=32 time=4ms TTL=124
Reply from 10.252.1.21: bytes=32 time=2ms TTL=124
Request timed out.
Request timed out.
Request timed out.
Reply from 10.252.1.21: bytes=32 time=5ms TTL=124
Reply from 10.252.1.21: bytes=32 time=3ms TTL=124
Reply from 10.252.1.21: bytes=32 time=3ms TTL=124
Reply from 10.252.1.21: bytes=32 time=3ms TTL=124
```

6. Repeat the test if the first failover was still fast.

Trouble Ticket Setup 03

Steps

1. On PC1, on the desktop, open folder ASTS. Run **ASTS-Lab14-TT03.cmd**

Trouble Ticket 03.01

2. The customer is rolling out a new network, using a VSX pair as the Aggregation switches. The Access switches will be connected using a LAG to the VSX pair. The customer plans to use zero touch provisioning for the deployment of new switches. Now the customer is deploying a new switch (Sw2) and has an issue with a PC3 connected to this switch, it does not seem to be able to connect to the network.
3. The customer has already tested several steps:

On Sw2, they have verified that the links are up and connected properly using

```
show lldp neighbor
```

On Sw2, they have used verified that both interfaces of the LAG are UP as Forwarding state.

```
show lacp interfaces
```

On Sw2, they have verified IP connectivity to the Agg1/Agg2 and both pings were successful.

```
ping 10.x.1.2  
ping 10.x.1.3
```

On Agg1, Agg2 and Sw2, they have verified that the VLAN11 exists and is enabled on the LAG.

```
show VLAN 11
```

4. Still, PC3 (assigned to VLAN11) is unable to access the network.
5. The customer wants you to investigate why the LAG seems ok, but the PC3 still cannot access the network.

Trouble Ticket Setup 04

Steps

1. On PC1, on the desktop, open folder ASTS. Run **ASTS-Lab14-TT04.cmd**

Trouble Ticket 04.01

2. The customer has routed SVI 41 between Sw1 - Agg1 and routed SVI 42 between Sw1 - Agg2. The customer is using OSPF as the routing protocol to exchange routing information between these 3 devices. These routed SVIs 41 and 42 should always be up, independent of the VSX ISL status. The customer is using BFD between the OSPF routers for fast failure detection.
3. However, the customer has noticed that the OSPF adjacencies between Sw1 and the VSX pair are not stable when the ISL is broken. Within a minute after the ISL fails, the OSPF adjacency between Sw1 and Agg2 goes offline. Also, it can take more than a minute to restore the OSPF adjacency after the ISL is restored.

The customer is using this test procedure:

4. On Sw1, verify that there is an OSPF adjacency with Agg1 (SVI 41) and Agg2 (SVI 42). You may need to wait a few moments after the trouble ticket checkpoint was loaded.

```
show ip ospf neighbors
```

5. On Agg1, disable lag256

```
interface lag256  
shutdown
```

6. On Sw1, review the OSPF neighbors, the adjacency to Agg2 (SVI 42) will go offline within a minute.

```
show ip ospf neighbors
```

7. On Agg1, enable lag256

8. On Sw1, review the OSPF neighbors, the adjacency to Agg2 (SVI 42) takes significant time before it is established again.

```
show ip ospf neighbors
```

9. The customer wants you to investigate what is wrong with these links and apply the corrections. The result should be that the OSPF adjacencies are not affected by an ISL failure.

Trouble Ticket Setup 05

Steps

1. On PC1, on the desktop, open folder ASTS. Run **ASTS-Lab14-TT05.cmd**

Trouble Ticket 05.01

2. Customer has recently deployed a VSX pair with L3 uplink connections to Sw1.
3. Agg1 connects with SVI41 to Sw1, Agg2 connects with SVI42 to Sw1.
4. During the setup, the routed connections and OSPF adjacencies were verified and found correct.
5. Recently some best practices were applied to the network and the customer observed that Sw1 does not have an OSPF adjacency with Agg2 anymore. They want to understand what happened and how to resolve it.

Trouble Ticket Setup 06

Steps

1. On PC1, on the desktop, open folder ASTS. Run **ASTS-Lab14-TT06.cmd**

Trouble Ticket 06.01

2. Customer has recently deployed a VSX pair using 1/1/45, 1/1/46 and 1/1/47 for the ISL LAG256.
3. The VSX is configured with L3 uplink connections to Sw1. Agg1 connects with SVI41 to Sw1, Agg2 connects with SVI42 to Sw1. This is an OSPF routed topology.
4. The customer did not have spare ports available for a direct VSX keepalive, so they have configured in-band Keepalive detection. The reported Keepalive state is 'Keepalive-Established', so Agg1 and Agg2 have successfully established the Keepalive channel.
5. However, when the customer tests a failure of the ISL link, the secondary node remains active, so the Keepalive does not seem to work and a split-brain condition occurs.

note: You should not configure a direct Keepalive since the customer has stated that there are no more direct interfaces available.

The customer has tested with these steps:

6. On Agg2, after loading the checkpoint, verify that the linkup-delay timer has passed and that VSX is working ok. There should be no interfaces 'Disabled by VSX'

```
show vsx status
show interface brief
```

7. On Agg2, verify the Keepalive is Established

```
show vsx status keepalive
```

8. On Agg1, disable lag256ore

```
interface lag256
shutdown
```

9. On Agg2, verify interfaces are still up (should have been disabled by VSX)

```
show interface brief
```

10. Enable the LAG256 and repeat the above process to verify.
11. They want to understand what happened and they want you to implement a correction so the split-brain detection operates correctly.

Trouble Ticket Setup 07

Authentication

Steps

1. On PC1, on the desktop, open folder ASTS. Run **ASTS-Lab14-TT07.cmd**

Trouble Ticket 07.01

2. The customer wants to setup 802.1X authentication on Sw2 with ClearPass Downloadable User Roles, but they are running into issues.
3. They have tried various configurations and made several configuration updates, but they have lost track. They want you to correct their configuration and they want to see a successful authentication on PC4 with a Downloadable User Role applied to the user.
4. The account that should be used for 802.1X on PC4 is **asts-employee-dur / aruba123**
5. This table shows the required configuration.

Setting	Value	Notes
ClearPass	10.252.1.23	RADIUS
ClearPass DNS name	cppm.arubatraining.com	
ClearPass RADIUS secret	aruba123	
ClearPass DUR username	duradmin	
ClearPass DUR password	aruba123	
DNS	10.252.1.21	

Trouble Ticket Setup 08

In this section, several RADIUS server configuration errors will be explored.

For some tickets, you don't need to apply a fix, you only need to identify the issue on the RADIUS server.

Steps

1. On PC1, on the desktop, open folder ASTS. Run **ASTS-Lab14-TT08.cmd**
2. On PC4, configure 802.1X authentication with user **asts-employee-dur / aruba123**. This should be successful.
3. note: You may need to bounce port 1/1/4 on Sw2 after the checkpoint is loaded.
4. In the following tickets, you will configure PC4 to login with different user accounts, each of them contains some issue that results in network access issues. The customer wants you to investigate and correct when requested.

Trouble Ticket 08.01

5. Configure PC4 to authenticate using account **asts-error1 / aruba123**
6. The customer reports that the client does not come online on Sw2, even though ClearPass Access Tracker reported a success authentication.
7. You must investigate and ensure the user is authenticated and the correct role can be applied. The PC4 does **not** need to get an IP address.

Trouble Ticket 08.02

8. Configure PC4 to authenticate using account **asts-error2 / aruba123**
9. The customer reports that the client does not come online on Sw2, even though ClearPass Access Tracker reported a success authentication.
10. You must investigate and ensure the user is authenticated and the correct role can be applied. The PC4 should get an IP address and be able to access the network.

Trouble Ticket 08.03

11. Configure PC4 to authenticate using account **asts-error3 / aruba123**
12. Customer wants to explore the ClearPass Downloadable User roles. Their ClearPass administrator is not very familiar with the Downloadable User Roles and they are running into an issue when this user authenticates.

13. They are convinced that there is an issue in the ClearPass configuration for this user, but they want you to identify what they did wrong, since access tracker just shows a success message. You should troubleshoot on the switch and **identify** what is wrong in the ClearPass configuration, **you do not need to correct this issue**.

Trouble Ticket 08.04

1. Configure PC4 to authenticate using account asts-error4 / aruba123
2. The customer expects that PC4 gets connected to VLAN11 using this user account. However, the device does not get an IP address in VLAN11, even though authentication seems to work fine.
3. The customer expects you to identify the issue and correct the configuration so the PC4 will be connected to VLAN 11 when it authenticates using this account.

Trouble Ticket 08.05

note: This ticket is very similar to ticket 08.03. If you have completed ticket 08.03, you may skip this ticket.

4. Configure PC4 to authenticate using account asts-error5 / aruba123
5. Customer wants to explore the ClearPass Downloadable User roles. Their ClearPass administrator is not very familiar with the Downloadable User Roles and they are running into an issue when this user authenticates.
6. They are convinced that there is an issue in the ClearPass configuration for this user, but they want you to identify what they did wrong, since access tracker just shows a success. You should troubleshoot on the switch and **identify** what is wrong in the ClearPass configuration, **you do not need to correct this issue**.

Trouble Ticket 08.06

7. Configure PC4 to authenticate using account asts-error6 / aruba123
8. The customer is testing a new device that should come online in VLAN 12, but when they configure the device, it does not come online.
9. The customer expects you to identify the issue and correct the configuration so the PC4 will be connected to VLAN 12 when it authenticates using this account.

Trouble Ticket 08.07

10. Configure PC4 to authenticate using account asts-error7 / aruba123
11. The customer is testing authentication between the AOS-CX switch and a RADIUS policy that is also used for some legacy switches in the network.
12. The PC4 fails to be authorized on the Sw2. The customer expects you to identify the issue and correct the configuration so the PC4 will be authenticated. PC4 does not need to get an IP address, it just needs to be online as 'success' authenticated.

Aruba Training-Confidential

Trouble Ticket Setup 09

In this section you will need to troubleshoot a routing and multicast routing setup.

Steps

1. On PC1, on the desktop, open folder ASTS. Run **ASTS-Lab14-TT09.cmd**
2. On PC4, release and renew the IP address. The PC has been assigned to VLAN212 and should have an IP address in the 10.x.212.0/24 subnet.
3. On PC4, run **asts-lab09-PC4-mcast-tx.cmd** to send multicast traffic to 239.1.1.100.
4. On PC3, renew the IP address, it should receive an IP address in the 10.x.11.0/24 subnet.
5. On PC3, run **asts-lab09-PC3-mcast-rx.cmd** to attempt to receive multicast traffic for 239.1.1.100. This will not work and must be resolved as part of this ticket.

Trouble Ticket 09.01

6. The customer is migrating from a static route configuration to OSPF. They have reconfigured some IP addresses of routed links and they have configured OSPF between the routers.
7. They are also implementing a multicast application and they are attempting to setup multicast between VLAN212 as the transmitter (PC4) and V11 as the receiver (PC3).
8. They are experiencing reachability issues between the VLAN11 and VLAN212 subnets, they report that PC3 cannot even ping PC4.
9. You must ensure that the reachability between PC3 and PC4 is restored and that the multicast transmission by PC4 to 239.1.1.100 can be received by PC3.

This concludes the trouble tickets activity.

Lab 14: Trouble Tickets Hints and Solutions

This document contains the hints and solutions for the Trouble Tickets in the Lab Guide.

Aruba Training-Confidential

Trouble Ticket Hints and Solutions

Refer to the Solution Guide document for the solution steps.

In the Lab Guide version, only the Hints are shown.

Trouble Ticket 01.01

Hints

```
On Sw2
show interface brief
show spanning-tree
```

Solution

Trouble Ticket 01.02

Hints

```
# Use PC4 with tracer -d 10.252.1.23 to attempt to detect where drops occur
On devices where response seems unstable, review the interface stats and configuration where
traffic comes in and goes out.
```

Solution

Trouble Ticket 02.01

Hints

```
On Sw2
show interface brief
show lacp interfaces
```

Solution

Trouble Ticket 03.01

Hints

```
# On Agg1
show interface brief
show lacp interfaces multi-chassis
```

Solution

Trouble Ticket 04.01

Hints

```
# On Agg2
show interface brief
show ip interface brief
show vlan
show vlan <id>
```

Solution

Trouble Ticket 05.01

Hints

```
# On Sw2
show ip ospf neighbors
show ip ospf interfaces

# On Agg2
show ip ospf neighbors
show ip ospf interfaces
```

Solution

Aruba Training-Confidential

Trouble Ticket 06.01

Hints

```
# On Agg2
show vsx status
show vsx status keepalive
show vsx configuration keepalive
trace
show ip ospf interfaces
```

Solution

Aruba Training-Confidential

Trouble Ticket 07.01

Hints

```
# On Sw2
ping <ip>
ping <name>
show port-access clients
show port-access clients detail
show aaa server-groups
debug portaccess all

clear radius-server statistics authentication
show radius-server statistics authentication
show aaa authentication port-access dot1x authenticator interface 1/1/4 port-statistics
show aaa authentication port-access interface 1/1/4 client-status
```

Solution

Trouble Ticket 08.01

Hints

```
# On Sw2
show port-access clients
show port-access clients detail
```

Solution

Aruba Training-Confidential

Trouble Ticket 08.02

Hints

```
# On Sw2
show port-access clients
show port-access clients detail
```

Solution

Aruba Training-Confidential

Trouble Ticket 08.03

Hints

```
# On Sw2
debug portaccess all
show debug sev err
show event -r -n 10 -s err
```

Solution

Aruba Training-Confidential

Trouble Ticket 08.04

Hints

```
# Sw2 Review the status
show port-access clients
show port-access clients details

# Sw2 Enable debugging
debug portaccess all
clear debug buffer
port-access reauthenticate interface 1/1/4
show debug buffer severity err
```

Solution

Trouble Ticket 08.05

Hints

```
show port-access clients
show port-access clients detail

debug portaccess all
clear debug buffer
port-access reauthenticate interface 1/1/4
show debug buffer
```

Solution

Trouble Ticket 08.06

Hints

```
show port-access clients
show port-access clients detail

debug portaccess all
clear debug buffer
port-access reauthenticate interface 1/1/4
show debug buffer
```

Solution

Trouble Ticket 08.07

Hints

```
show port-access clients  
show port-access clients detail
```

Solution

Aruba Training-Confidential

Trouble Ticket 09.01 - Routing and Multicast

Hints

```
# PC3
ping to PC4
# Agg1/2
show ip route 10.x.212.0

show ip route 10.x.212.0
show ip ospf neighbors
show ip pim neighbors
show ip pim interface brief
show ip ospf interface brief
```

Solution

This concludes the trouble tickets activity.

Aruba Training-Confidential

3333 Scott Blvd, Santa Clara, CA 95054 TEL:
408.227.4500 | FAX: 408.227.4550
www.ARUBANETWORKS.com