



Aruba Advanced Switching Troubleshooting and Solutions

Solution Guide



Aruba Advanced Switching Troubleshooting and Solutions

20.41

Solutions Guide

December 2020

Aruba Advanced Switching Troubleshooting and Solutions Solutions 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-
Solutions-v20.41 December
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	13
Objectives.....	13
Steps	13
Load initial configuration to the router devices	13
OOBM Network Access.....	13
Push Base configuration to your devices	15
TASK 3: VERIFY DEVICE ACCESS.....	16
Objectives.....	16
Device IP Address Table	16
STEPS.....	16
Verify in-band connectivity.....	17
OOBM Connection tests.....	18
Configuration file formats.....	18
Table configuration backup and restore on Management PC.....	19
Table checkpoint configuration (Optional).....	20
Course labs 'known good' checkpoints.....	21
OPTIONAL LAB 02: TROUBLESHOOTING TOOLS	23
REQUIREMENTS.....	23
OBJECTIVES	23
TASK 1: INTRODUCTION TO SHOW COMMANDS AND SUPPORT FILES	24
Objectives.....	24
Steps	24
Show commands and support files review	24
Interface statistics.....	25
Support Logs	27

Accounting logs	27
Event log	28
Terminal Monitor	30
TASK 2: DIAGNOSTIC TOOLS	33
Objectives	33
Steps	33
Explore dropped frames - Jumbo frames	33
Explore dropped frames - VLAN mismatch	36
TASK 3: DEBUGGING OPTIONS	38
Objectives	38
Steps	38
Review default logging and logging destination	38
Review normal ARP operation	39
Review missing ARP reply operation	40
TASK 4: TRAFFIC MIRRORING AND PACKET CAPTURES	42
Objectives	42
Steps	42
Data plane mirror - Local mirror	42
Dataplane mirror - Remote mirror over GRE tunnel	43
Dataplane mirror - TSHARK - Mirror to cpu - Live	46
TSHARK - Mirror to cpu - PCAP	48
Class-based mirroring	49
TCPDUMP	50
LAB 03: MONITORING AND AUTOMATION TOOLS	53
REQUIREMENTS	53
OBJECTIVES	53
SCENARIO	54
NetEdit	54
RESTAPI	54
NAE	54
SNMP	54
TASK 1: NETEDIT	56
Objectives	56
Steps	56
Discovery issues	58
Diagnostics page	61
NetEdit Insights	62
NetEdit Debugging	62
Collect NetEdit logs	65
TASK 2: RESTAPI	66
Objectives	66
Steps	66
Using the built-in Swagger interface	66
Wildcard	66
Escape Characters	68
Filtering output	68
Configuration Depth	69
REST Sessions and clearing REST sessions on AOS-CX	69
TASK 3: NAE	71
Objectives	71
Steps	71
Introduction	71

NAE Script Hierarchy.....	71
NAE Script Debugging.....	72
NAE Script deployment.....	74
Clearing NAE data.....	74
TASK 4: SNMP CONFIGURATION.....	76
Objectives.....	76
Steps.....	76
SNMP VRF Support.....	76
SNMPv2 communities.....	76
Configure SNMPv3 Support for Airwave.....	77
Troubleshooting SNMP packets.....	78
SNMP Events.....	78
Diagnostics SNMP.....	79
LAB 04: VSX AND LAYER2 TECHNOLOGIES.....	80
REQUIREMENTS.....	80
OBJECTIVES.....	80
SCENARIO.....	81
VSX.....	81
VSX LAG - Aruba Gateway.....	81
VSX LAG - Sw2.....	81
VSX - DHCP Relay.....	81
VSX - Virtual Active Gateway.....	82
Client in VLAN11.....	82
VSX Best Practices.....	82
VSX Linkup Delay.....	82
Spanning-Tree.....	82
Loop Protection.....	83
TASK 1: CONFIGURE VSX.....	85
Objectives.....	85
Steps.....	85
VSX Configuration.....	85
VSX SVI Active Gateway - SVI 1.....	87
VSX SVI Active Gateway - SVI 12.....	87
VSX LAG to Aruba Gateway.....	90
VSX LAG to Sw2.....	92
VSX LAG LACP Status.....	93
VSX DHCP Relay Services.....	96
Impact of Primary node SVI Shutdown.....	98
Impact of Primary node SVI Shutdown - DHCP Relay.....	98
Impact of Primary node SVI Shutdown - Traffic.....	98
Virtual Active Gateway vs Active Gateway.....	100
Verify the configuration and troubleshoot the Layer2 connection.....	102
TASK 2: APPLY SOME VSX BEST PRACTICES.....	106
Objectives.....	106
Steps.....	106
MTU Configuration.....	107
Uplink delay-timer.....	108
Validate the linkup delay-timer.....	108
Review impact on the VSX LAG interfaces.....	109
Linkup-delay exclusion.....	111
VSX LAG Hashing algorithm.....	113
TASK 3: CONFIGURING STP WITH VSX.....	115

Objectives.....	115
Steps.....	115
Configure STP.....	115
VSX STP BPDU Transmissions.....	116
STP features - Root-guard.....	117
Restore the configuration.....	119
VSX configuration consistency check.....	120
Configure and review STP on the Access switch Sw2.....	120
STP Loop-guard.....	121
Validating STP features - Loop-guard.....	121
Configure the access ports.....	122
Topology Change Notifications.....	124
BPDU-Guard configuration.....	125
TASK 4: CONFIGURING LOOP-PROTECTION IN THE ACCESS LAYER.....	127
Objectives.....	127
Steps.....	127
Loop Setup.....	127
Configure the network loop without Loop Protect.....	127
Enable the loop.....	128
Loop Protection configuration.....	129
Restore the configuration.....	131
LAB 05: REVIEW OF LAYER3 BASICS.....	132
REQUIREMENTS.....	132
OBJECTIVES.....	132
SCENARIO.....	133
OSPF.....	133
VSX Active Forwarding.....	133
OSPF Key-chain and tuning (Optional).....	133
Policy Based Routing (PBR).....	133
PBR Path selection.....	134
TASK 1: CONFIGURE VSX WITH OSPF.....	136
Objectives.....	136
Steps.....	136
OSPF Configuration on Agg1 and Agg2.....	136
TASK 2: DEBUGGING OSPF.....	139
Objectives.....	139
Steps.....	139
Adjacency Debugging.....	139
Link Troubleshooting.....	142
Review the normal OSPF adjacency setup.....	145
Invalid Timers.....	146
Invalid authentication.....	148
Router-id Conflict on Peer.....	150
OSPF Point to Point networks.....	151
TASK 3: VSX ACTIVE FORWARDING.....	153
Objectives.....	153
Steps.....	153
Active Forwarding - Understand the need for Active Forwarding (optional).....	156
Testing.....	158
Active Forwarding.....	158
OPTIONAL TASK 4: OSPF KEY-CHAIN AND MAX-METRIC ON STARTUP.....	161
Objectives.....	161

Steps	161
Configuration	161
Activation	162
Migration - Attempt 1	162
TCPDUMP on Agg1	164
Migration - Think about the migration steps	165
Migration - Attempt 2	165
Max-metric on Startup	169
TASK 5: POLICY BASED ROUTING (PBR).....	172
Objectives.....	172
Steps	172
Prepare the setup.....	172
Option 1: Manual Configuration.....	172
Option2: Saved Script configuration.....	175
OSPF best path verification.....	175
Test assigned link and failover without PBR applied.....	176
Test link failover.....	177
Enhance the failover detection for the SVI using BFD	177
Test link failover.....	179
Policy Based Routing Configuration.....	180
Verification.....	181
Verify the 'other' traffic.....	182
Cleanup	182
Required: Post-lab checkpoint.....	182
LAB 06: BGP	184
REQUIREMENTS.....	184
SCENARIO	185
iBGP	185
eBGP.....	185
BGP Advertisements	185
BGP Peer Groups.....	186
Controlling Transit Traffic	186
Outbound Traffic Route Control.....	186
Inbound Traffic route control.....	187
AS Number filtering	188
iBGP peering	188
TASK 1: PREPARE THE SETUP.....	190
Objectives.....	190
Steps	190
Router Configurations.....	190
TASK 2: CONFIGURE IBGP AND EBGP PEERING	191
Objectives.....	191
Steps	191
Troubleshooting iBGP peering	191
Incorrect iBGP session IP address.....	191
Configure iBGP peering between the other switches.....	194
Configure eBGP peering with routerA and routerB	195
Troubleshoot remote AS number	195
Next hop self.....	197
Solution: configure next-hop self.....	199
Solution: Reset the BGP peering.....	200
eBGP between Sw1 and Sw2 (acting as External BGP Partner).....	201

<i>Advertise default route into OSPF</i>	204
<i>Advertise a route to eBGP peer</i>	204
TASK 3: USING BGP PEER GROUPS	211
<i>Objectives</i>	211
<i>Steps</i>	211
<i>Prepare the Peer-group</i>	211
<i>Migrate the existing peer to the peer-group</i>	211
TASK 4: CONTROLLING TRANSIT TRAFFIC.....	213
<i>Objectives</i>	213
<i>Steps</i>	213
<i>Prevent transit AS</i>	213
<i>Allow selective transit AS</i>	216
TASK 5: OUTBOUND TRAFFIC ROUTE CONTROL	218
<i>Objectives</i>	218
<i>Steps</i>	218
<i>Apply default local preference</i>	218
<i>Apply selective local preference</i>	222
<i>AS multipath-relax</i>	225
<i>Optional: Test link failure and route failover</i>	227
<i>Using weight locally on a router</i>	228
<i>Soft reconfiguration inbound</i>	232
<i>Soft reconfiguration inbound impact on route refresh</i>	235
TASK 6: INBOUND TRAFFIC ROUTE CONTROL	238
<i>Objectives</i>	238
<i>Steps</i>	240
<i>Apply AS-Path length for outbound advertisements</i>	240
<i>Use MED for inbound route preference</i>	242
<i>Prepare the link between Agg2 and Sw2</i>	243
<i>Apply the MED values</i>	244
<i>Using Community Attributes for route control</i>	247
<i>Correct the iBGP peer configuration</i>	253
OPTIONAL TASK 7: AS PATH LIST EXCLUSION.....	256
<i>Objectives</i>	256
<i>Steps</i>	256
OPTIONAL TASK 8: USING A ROUTE REFLECTOR FOR IBGP	259
<i>Objectives</i>	259
<i>Steps</i>	259
<i>Break the iBGP full mesh topology</i>	259
<i>Review the result of the broken topology</i>	260
<i>Configure the route reflector</i>	261
LAB 07: ROUTE REDISTRIBUTION.....	263
REQUIREMENTS	263
SCENARIO	264
<i>Setup</i>	264
<i>Static route into OSPF</i>	264
<i>Prefix Lists</i>	264
<i>Route Redistribution between OSPF and BGP</i>	265
<i>Single Link Redistribution</i>	265
<i>Dual Link Route Redistribution</i>	265
TASK 1: LOAD THE START CONFIGURATIONS	266
<i>Objectives</i>	266
<i>Steps</i>	266

TASK 2: ROUTE REDISTRIBUTION OF STATIC ROUTES INTO OSPF	267
OBJECTIVES	267
Steps	268
Redistribute static routes into OSPF	268
Cost Manipulation of the redistributed routes	269
Granular control of the redistribution using route-maps	271
OSPF External Metric Type1	273
Test failover impact on the route metrics.....	275
Understanding Prefix Lists.....	276
TASK 3: ROUTE REDISTRIBUTION BETWEEN OSPF AND BGP	279
Introduction.....	279
Objectives.....	279
Steps	281
Prepare the setup.....	281
Verify current routes and OSPF neighbors.....	281
Configure eBGP peering between Agg1 and routerA.....	281
Redistribute local OSPF routes into BGP.....	283
Redistribute BGP routes into OSPF	285
TASK 4: ROUTE REDISTRIBUTION WITH MULTIPLE LINKS.....	287
Introduction.....	287
Objectives.....	287
Configure second link: Attempt 1.....	288
Review routing before the second link redistribution is configured.....	289
Agg2: Redistribute BGP into OSPF.....	291
Review routing after the redistribution.....	291
Configure Second Link: Attempt 2 using route tags.....	293
Customer Redistribution on Agg2.....	294
Partner network configuration.....	297
LAB 08: VRF AND ROUTE LEAKING	299
OBJECTIVES	299
REQUIREMENTS.....	299
SCENARIO	300
Static Route Leaking	300
Dynamic Route Leaking using MP-BGP.....	300
Route-map based import.....	300
TASK 1: PREPARE THE BASE CONFIGURATION.....	302
Objectives.....	302
Steps	302
IOT VRF Configuration.....	302
TASK 2: EXPLORE THE ENVIRONMENT	304
Objectives.....	304
Steps	304
Connectivity Checks.....	304
Verify VRF Route Isolation	304
TASK 3: STATIC ROUTE LEAKING	306
Objectives.....	306
Steps	306
Manually Leak Directly Connected Interface.....	306
Make the leaked route available inside the VRF branch	308
Debug on Sw2.....	309
Prepare for MP-BGP leaking.....	310
TASK 4: MP-BGP.....	311

Objectives.....	311
Steps.....	311
Configure Route Distinguisher and import routes into BGP routing table.....	311
Route Targets.....	314
Simple method.....	314
Verify connectivity.....	317
Route-map based RT.....	318
LAB 09: MULTICAST	323
OBJECTIVES	323
REQUIREMENTS	323
Introduction.....	323
SCENARIO	324
Setup	324
Multicast routing	324
TASK 1: LOAD THE START CONFIGURATIONS	326
Objectives.....	326
Steps	326
Introduction.....	326
Explore the environment.....	326
TASK 2: CONFIGURE PIM SM.....	329
Objectives.....	329
Steps	329
Configure the C-BSR.....	329
Enable PIM-SM on all routing devices.....	330
Configure Sw1 as C-RP	334
TASK 3: CONFIGURE IGMP AND VERIFY MULTICAST OPERATION	336
Objectives.....	336
Steps	336
IGMP Querier	336
IGMP Snooping	340
Test Multicast and verify status on Sw1	344
Test without multicast receiver	344
Test with multicast receiver	348
Review the status of the multicast route.....	348
Test with mtrace	356
Optional: PIM Diagnostics Dump.....	358
TASK 4: CONFIGURE DISTRIBUTED RP	360
Objectives.....	360
Steps	360
Configure regional RP.....	360
Verify the registration traffic on the regional RP.....	361
OPTIONAL TASK 5: UNDERSTAND THE MAC - IP MULTICAST RELATION	364
Objectives.....	364
Steps	364
Explore the Multicast IP to Ethernet mapping	364
TASK 6: VSX MULTICAST FAILOVER.....	366
Objectives.....	366
Steps	366
VSX Active Active Configuration	366
Understand the VSX Multicast Routing failover.....	366
Verify the failover.....	371
TASK 7: IGMP ACL.....	374

Objectives.....	374
Steps.....	374
IGMP Querier ACL.....	374
IGMP Snooping ACL.....	375
IGMP debug.....	377
LAB 10: QUALITY OF SERVICE	378
OBJECTIVES	378
REQUIREMENTS.....	378
SCENARIO	379
TASK 1: QOS POLICIES	381
Objectives.....	381
Steps	381
Review default queuing and QOS maps	381
Configure QOS Policy on Sw2	383
Verify the Remark - last hop switch	384
Default Queue and Trust modes	386
Verify the impact of the QOS trust modes.....	387
On Sw2, you will now review the local traffic handling.....	388
TASK 2: RATE LIMITER AND TRAFFIC SHAPING.....	391
Objectives.....	391
Steps	391
Inbound Rate Limiters	391
Define the Rate Limiter on the Aggregation VSX.....	391
Troubleshoot the VSX-sync.....	391
Verify the rate limiter.....	395
Outbound Traffic shaping	395
Rate limit and burst size	396
LAB 11: DYNAMIC SEGMENTATION	397
OBJECTIVES	397
REQUIREMENTS.....	397
SCENARIO	398
TASK 1: PREPARE THE CONFIGURATION	402
Objectives.....	402
Steps	402
Configuration	402
Table with RADIUS test user accounts	402
TASK 2: CONFIGURE BASIC 802.1X AUTHENTICATION	403
Objectives.....	403
Steps	403
Basic RADIUS and 802.1X configuration	403
Statistics EAP and RADIUS	406
Review the debug output for a successful authentication	409
TASK 3: RADIUS TROUBLESHOOTING	411
Objectives.....	411
Steps	411
Invalid Aruba-User-Role	411
Cached re-authentication	414
Verify the operation without cached re-authentication.....	415
Critical Role for Authentication	417
Configure the switch with cached re-authentication	418
Understand the need for RADIUS tracking.....	420

Recover the link and verify the operation	424
TASK 4: ONBOARDING PRECEDENCE ORDER	426
Objectives	426
Steps	426
Device Profile with Local MAC Authentication (LMA)	426
Device Profile Fallback role	429
Device Profile with AAA enabled on the port: Onboarding Precedence	430
Reject RADIUS Response	431
Optional: Device Profile with RADIUS Failure	432
TASK 5: USER-BASED TUNNELING	434
Objectives	434
Steps	434
Configuring UBT	434
Review Normal UBT Authentication	437
Invalid Gateway Role	440
MTU for Tunneled users	442
Configure Jumbo support	445
Default role on the Aruba Gateway	446
TASK 6: USER-BASED TUNNELING QOS	449
OBJECTIVES	449
STEPS	449
Apply a QOS Policy to the Switch user role	449
Optional: Trace traffic and compare inner DSCP to outer GRE DSCP markings	452
TASK 7: DOWNLOADABLE USER ROLES	456
Objectives	456
Steps	456
ClearPass DUR account configuration	456
ClearPass Certificate Validation	457
Install the Trust Anchor certificate on the Switch	460
Review normal DUR authentication and role	465
LAB 12: NETWORK SECURITY FEATURES	468
OBJECTIVES	468
REQUIREMENTS	468
SCENARIO	469
TCAM Resource Usage	469
Hitcounts	469
Control Plane Security	469
Control Plane Policing	470
DHCP Snooping	470
ARP Inspection	470
TASK 1: ACL AND RESOURCE USAGE	472
Objectives	472
Steps	472
Configuration	472
Prepare the a routed ACL	472
Update the ACL using object-groups	473
Hitcounts and Statistics	476
Optional: Unique Policy per interface for unique counters	478
TASK 2: CONTROL PLANE ACL	480
Objectives	480
Steps	480
COPP ACL on VRF mgmt	480

Prepare for failure.....	481
COPP ACL on VRF default	483
TASK 3: CONTROL PLANE POLICING.....	486
Objectives.....	486
Steps	486
VRF mgmt behavior.....	489
OPTIONAL TASK 4: MONITORING CONTROL PLANE WITH NAE	490
Objectives.....	490
Steps	490
TASK 5: DHCP SNOOPING	493
Objectives.....	493
Steps	493
Optional: Authorized Server	495
DHCP Bindings Database	496
TASK 6: ARP INSPECTION	498
Objectives.....	498
Steps	498
Configure ARP Inspection	498
Optional: Verify IP Binding database dependency	501
Verify ARP Inspection filtering	503
Endpoint with static IP	504
LAB 13: IPV6	506
OBJECTIVES	506
REQUIREMENTS	506
SCENARIO	507
OSPFv3.....	507
SLAAC.....	507
DHCPv6 - Stateless.....	507
DHCPv6 - Stateful	508
TASK 1: CONFIGURE IPV6 LINKS AND OSPFv3.....	510
Objectives.....	510
Steps	510
Agg1 - Agg2 Link.....	510
OSPFv3 on Agg1 and Agg2	512
Configure Sw1 as OSPFv3 router	514
TASK 2: CONFIGURE IPV6 ON CLIENT SUBNET USING SLAAC	517
Objectives.....	517
Steps	517
Aggregation configuration.....	517
Client verification	517
TASK 3: CONFIGURE IPV6 ON CLIENT SUBNET USING DHCP	521
Objectives.....	521
Steps	521
Configure PC4.....	521
Configure DHCPv6 Server	522
Configure SVI12 for Stateless DHCPv6	523
Verify Stateless DHCPv6.....	524
Configure SVI 12 for Stateful DHCPv6.....	526
Verify Stateful DHCPv6	526
LAB 14: TROUBLE TICKETS HINTS AND SOLUTIONS	529
TROUBLE TICKET HINTS AND SOLUTIONS	530

<i>Trouble Ticket 01.01</i>	530
<i>Trouble Ticket 01.02</i>	533
<i>Trouble Ticket 02.01</i>	538
<i>Trouble Ticket 03.01</i>	541
<i>Trouble Ticket 04.01</i>	543
<i>Trouble Ticket 05.01</i>	546
<i>Trouble Ticket 06.01</i>	549
<i>Trouble Ticket 07.01</i>	551
<i>Trouble Ticket 08.01</i>	559
<i>Trouble Ticket 08.02</i>	562
<i>Trouble Ticket 08.03</i>	565
<i>Trouble Ticket 08.04</i>	568
<i>Trouble Ticket 08.05</i>	570
<i>Trouble Ticket 08.06</i>	572
<i>Trouble Ticket 08.07</i>	577
<i>Trouble Ticket 09.01 - Routing and Multicast</i>	581

LAB 14: TROUBLE TICKETS 590

TASK 1: PREPARE THE TROUBLESHOOTING SETUP	591
<i>Introduction</i>	591
TROUBLE TICKET SETUP 01	593
<i>Steps</i>	593
<i>Trouble Ticket 01.01</i>	594
<i>Trouble Ticket 01.02</i>	595
TROUBLE TICKET SETUP 02	596
<i>Steps</i>	596
<i>Trouble Ticket 02.01</i>	596
TROUBLE TICKET SETUP 03	597
<i>Steps</i>	597
<i>Trouble Ticket 03.01</i>	597
TROUBLE TICKET SETUP 04	598
<i>Steps</i>	598
<i>Trouble Ticket 04.01</i>	598
TROUBLE TICKET SETUP 05	599
<i>Steps</i>	599
<i>Trouble Ticket 05.01</i>	599
TROUBLE TICKET SETUP 06	600
<i>Steps</i>	600
<i>Trouble Ticket 06.01</i>	600
TROUBLE TICKET SETUP 07	601
<i>Steps</i>	601
<i>Trouble Ticket 07.01</i>	601
TROUBLE TICKET SETUP 08	602
<i>Steps</i>	602
<i>Trouble Ticket 08.01</i>	602
<i>Trouble Ticket 08.02</i>	602
<i>Trouble Ticket 08.03</i>	602
<i>Trouble Ticket 08.04</i>	603
<i>Trouble Ticket 08.05</i>	603
<i>Trouble Ticket 08.06</i>	603
<i>Trouble Ticket 08.07</i>	604
TROUBLE TICKET SETUP 09	605
<i>Steps</i>	605

<i>Trouble Ticket 09.01</i>	605
LAB 0: TESTING LAB CONNECTIVITY (OPTIONAL)	1
OVERVIEW	1
<i>Objectives</i>	1
TASK 1: ARUBA TRAINING LAB ACCESS	2
<i>Objectives</i>	2
<i>Steps</i>	2
TASK 2: TESTING CONNECTIVITY	4
<i>Objectives</i>	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
<i>Objectives</i>	10
<i>Steps</i>	10
TASK 2: LOAD THE INITIAL CONFIGURATION	13
<i>Objectives</i>	13
<i>Steps</i>	13
<i>Load initial configuration to the router devices</i>	13
<i>OOBM Network Access</i>	13
<i>Push Base configuration to your devices</i>	15
TASK 3: VERIFY DEVICE ACCESS	16
<i>Objectives</i>	16
<i>Device IP Address Table</i>	16
STEPS	16
<i>Verify in-band connectivity</i>	17
<i>OOBM Connection tests</i>	18
<i>Configuration file formats</i>	18
<i>Table configuration backup and restore on Management PC</i>	19
<i>Table checkpoint configuration (Optional)</i>	20
<i>Course labs 'known good' checkpoints</i>	21
OPTIONAL LAB 02: TROUBLESHOOTING TOOLS	23
REQUIREMENTS	23
OBJECTIVES	23
TASK 1: INTRODUCTION TO SHOW COMMANDS AND SUPPORT FILES	24
<i>Objectives</i>	24
<i>Steps</i>	24
<i>Show commands and support files review</i>	24
<i>Interface statistics</i>	25
<i>Support Logs</i>	27
<i>Accounting logs</i>	27
<i>Event log</i>	28
<i>Terminal Monitor</i>	30
TASK 2: DIAGNOSTIC TOOLS	33
<i>Objectives</i>	33
<i>Steps</i>	33
<i>Explore dropped frames - Jumbo frames</i>	33
<i>Explore dropped frames - VLAN mismatch</i>	36

TASK 3: DEBUGGING OPTIONS	38
Objectives	38
Steps	38
Review default logging and logging destination	38
Review normal ARP operation	39
Review missing ARP reply operation	40
TASK 4: TRAFFIC MIRRORING AND PACKET CAPTURES	42
Objectives	42
Steps	42
Data plane mirror - Local mirror	42
Dataplane mirror - Remote mirror over GRE tunnel	43
Dataplane mirror - TSHARK - Mirror to cpu - Live	46
TSHARK - Mirror to cpu - PCAP	48
Class-based mirroring	49
TCPDUMP	50
LAB 03: MONITORING AND AUTOMATION TOOLS	53
REQUIREMENTS	53
OBJECTIVES	53
SCENARIO	54
NetEdit	54
RESTAPI	54
NAE	54
SNMP	54
TASK 1: NETEDIT	56
Objectives	56
Steps	56
Discovery issues	58
Diagnostics page	61
NetEdit Insights	62
NetEdit Debugging	62
Collect NetEdit logs	65
TASK 2: RESTAPI	66
Objectives	66
Steps	66
Using the built-in Swagger interface	66
Wildcard	66
Escape Characters	68
Filtering output	68
Configuration Depth	69
REST Sessions and clearing REST sessions on AOS-CX	69
TASK 3: NAE	71
Objectives	71
Steps	71
Introduction	71
NAE Script Hierarchy	71
NAE Script Debugging	72
NAE Script deployment	74
Clearing NAE data	74
TASK 4: SNMP CONFIGURATION	76
Objectives	76
Steps	76
SNMP VRF Support	76

SNMPv2 communities	76
Configure SNMPv3 Support for Airwave	77
Troubleshooting SNMP packets	78
SNMP Events	78
Diagnostics SNMP	79
LAB 04: VSX AND LAYER2 TECHNOLOGIES	80
REQUIREMENTS	80
OBJECTIVES	80
SCENARIO	81
VSX	81
VSX LAG - Aruba Gateway	81
VSX LAG - Sw2	81
VSX - DHCP Relay	81
VSX - Virtual Active Gateway	82
Client in VLAN11	82
VSX Best Practices	82
VSX Linkup Delay	82
Spanning-Tree	82
Loop Protection	83
TASK 1: CONFIGURE VSX	85
Objectives	85
Steps	85
VSX Configuration	85
VSX SVI Active Gateway - SVI 1	87
VSX SVI Active Gateway - SVI 12	87
VSX LAG to Aruba Gateway	90
VSX LAG to Sw2	92
VSX LAG LACP Status	93
VSX DHCP Relay Services	96
Impact of Primary node SVI Shutdown	98
Impact of Primary node SVI Shutdown - DHCP Relay	98
Impact of Primary node SVI Shutdown - Traffic	98
Virtual Active Gateway vs Active Gateway	100
Verify the configuration and troubleshoot the Layer2 connection	102
TASK 2: APPLY SOME VSX BEST PRACTICES	106
Objectives	106
Steps	106
MTU Configuration	107
Uplink delay-timer	108
Validate the linkup delay-timer	108
Review impact on the VSX LAG interfaces	109
Linkup-delay exclusion	111
VSX LAG Hashing algorithm	113
TASK 3: CONFIGURING STP WITH VSX	115
Objectives	115
Steps	115
Configure STP	115
VSX STP BPDU Transmissions	116
STP features - Root-guard	117
Restore the configuration	119
VSX configuration consistency check	120
Configure and review STP on the Access switch Sw2	120

STP Loop-guard	121
Validating STP features - Loop-guard	121
Configure the access ports	122
Topology Change Notifications	124
BPDUGuard configuration	125
TASK 4: CONFIGURING LOOP-PROTECTION IN THE ACCESS LAYER	127
Objectives	127
Steps	127
Loop Setup	127
Configure the network loop without Loop Protect	127
Enable the loop	128
Loop Protection configuration	129
Restore the configuration	131
LAB 05: REVIEW OF LAYER3 BASICS	132
REQUIREMENTS	132
OBJECTIVES	132
SCENARIO	133
OSPF	133
VSX Active Forwarding	133
OSPF Key-chain and tuning (Optional)	133
Policy Based Routing (PBR)	133
PBR Path selection	134
TASK 1: CONFIGURE VSX WITH OSPF	136
Objectives	136
Steps	136
OSPF Configuration on Agg1 and Agg2	136
TASK 2: DEBUGGING OSPF	139
Objectives	139
Steps	139
Adjacency Debugging	139
Link Troubleshooting	142
Review the normal OSPF adjacency setup	145
Invalid Timers	146
Invalid authentication	148
Router-id Conflict on Peer	150
OSPF Point to Point networks	151
TASK 3: VSX ACTIVE FORWARDING	153
Objectives	153
Steps	153
Active Forwarding - Understand the need for Active Forwarding (optional)	156
Testing	158
Active Forwarding	158
OPTIONAL TASK 4: OSPF KEY-CHAIN AND MAX-METRIC ON STARTUP	161
Objectives	161
Steps	161
Configuration	161
Activation	162
Migration - Attempt 1	162
TCPDUMP on Agg1	164
Migration - Think about the migration steps	165
Migration - Attempt 2	165
Max-metric on Startup	169

TASK 5: POLICY BASED ROUTING (PBR)	172
Objectives	172
Steps	172
Prepare the setup	172
Option 1: Manual Configuration	172
Option2: Saved Script configuration	175
OSPF best path verification	175
Test assigned link and failover without PBR applied	176
Test link failover	177
Enhance the failover detection for the SVI using BFD	177
Test link failover	179
Policy Based Routing Configuration	180
Verification	181
Verify the 'other' traffic	182
Cleanup	182
Required: Post-lab checkpoint	182
LAB 06: BGP	184
REQUIREMENTS	184
SCENARIO	185
iBGP	185
eBGP	185
BGP Advertisements	185
BGP Peer Groups	186
Controlling Transit Traffic	186
Outbound Traffic Route Control	186
Inbound Traffic route control	187
AS Number filtering	188
iBGP peering	188
TASK 1: PREPARE THE SETUP	190
Objectives	190
Steps	190
Router Configurations	190
TASK 2: CONFIGURE IBGP AND EBGP PEERING	191
Objectives	191
Steps	191
Troubleshooting iBGP peering	191
Incorrect iBGP session IP address	191
Configure iBGP peering between the other switches	194
Configure eBGP peering with routerA and routerB	195
Troubleshoot remote AS number	195
Next hop self	197
Solution: configure next-hop self	199
Solution: Reset the BGP peering	200
eBGP between Sw1 and Sw2 (acting as External BGP Partner)	201
Advertise default route into OSPF	204
Advertise a route to eBGP peer	204
TASK 3: USING BGP PEER GROUPS	211
Objectives	211
Steps	211
Prepare the Peer-group	211
Migrate the existing peer to the peer-group	211
TASK 4: CONTROLLING TRANSIT TRAFFIC	213

Objectives.....	213
Steps.....	213
Prevent transit AS.....	213
Allow selective transit AS.....	216
TASK 5: OUTBOUND TRAFFIC ROUTE CONTROL.....	218
Objectives.....	218
Steps.....	218
Apply default local preference.....	218
Apply selective local preference.....	222
AS multipath-relax.....	225
Optional: Test link failure and route failover.....	227
Using weight locally on a router.....	228
Soft reconfiguration inbound.....	232
Soft reconfiguration inbound impact on route refresh.....	235
TASK 6: INBOUND TRAFFIC ROUTE CONTROL.....	238
Objectives.....	238
Steps.....	240
Apply AS-Path length for outbound advertisements.....	240
Use MED for inbound route preference.....	242
Prepare the link between Agg2 and Sw2.....	243
Apply the MED values.....	244
Using Community Attributes for route control.....	247
Correct the iBGP peer configuration.....	253
OPTIONAL TASK 7: AS PATH LIST EXCLUSION.....	256
Objectives.....	256
Steps.....	256
OPTIONAL TASK 8: USING A ROUTE REFLECTOR FOR iBGP.....	259
Objectives.....	259
Steps.....	259
Break the iBGP full mesh topology.....	259
Review the result of the broken topology.....	260
Configure the route reflector.....	261
LAB 07: ROUTE REDISTRIBUTION.....	263
REQUIREMENTS.....	263
SCENARIO.....	264
Setup.....	264
Static route into OSPF.....	264
Prefix Lists.....	264
Route Redistribution between OSPF and BGP.....	265
Single Link Redistribution.....	265
Dual Link Route Redistribution.....	265
TASK 1: LOAD THE START CONFIGURATIONS.....	266
Objectives.....	266
Steps.....	266
TASK 2: ROUTE REDISTRIBUTION OF STATIC ROUTES INTO OSPF.....	267
OBJECTIVES.....	267
Steps.....	268
Redistribute static routes into OSPF.....	268
Cost Manipulation of the redistributed routes.....	269
Granular control of the redistribution using route-maps.....	271
OSPF External Metric Type1.....	273
Test failover impact on the route metrics.....	275

<i>Understanding Prefix Lists</i>	276
TASK 3: ROUTE REDISTRIBUTION BETWEEN OSPF AND BGP	279
<i>Introduction</i>	279
<i>Objectives</i>	279
<i>Steps</i>	281
<i>Prepare the setup</i>	281
<i>Verify current routes and OSPF neighbors</i>	281
<i>Configure eBGP peering between Agg1 and routerA</i>	281
<i>Redistribute local OSPF routes into BGP</i>	283
<i>Redistribute BGP routes into OSPF</i>	285
TASK 4: ROUTE REDISTRIBUTION WITH MULTIPLE LINKS.....	287
<i>Introduction</i>	287
<i>Objectives</i>	287
<i>Configure second link: Attempt 1</i>	288
<i>Review routing before the second link redistribution is configured</i>	289
<i>Agg2: Redistribute BGP into OSPF</i>	291
<i>Review routing after the redistribution</i>	291
<i>Configure Second Link: Attempt 2 using route tags</i>	293
<i>Customer Redistribution on Agg2</i>	294
<i>Partner network configuration</i>	297
LAB 08: VRF AND ROUTE LEAKING	299
OBJECTIVES	299
REQUIREMENTS	299
SCENARIO	300
<i>Static Route Leaking</i>	300
<i>Dynamic Route Leaking using MP-BGP</i>	300
<i>Route-map based import</i>	300
TASK 1: PREPARE THE BASE CONFIGURATION.....	302
<i>Objectives</i>	302
<i>Steps</i>	302
<i>IOT VRF Configuration</i>	302
TASK 2: EXPLORE THE ENVIRONMENT.....	304
<i>Objectives</i>	304
<i>Steps</i>	304
<i>Connectivity Checks</i>	304
<i>Verify VRF Route Isolation</i>	304
TASK 3: STATIC ROUTE LEAKING	306
<i>Objectives</i>	306
<i>Steps</i>	306
<i>Manually Leak Directly Connected Interface</i>	306
<i>Make the leaked route available inside the VRF branch</i>	308
<i>Debug on Sw2</i>	309
<i>Prepare for MP-BGP leaking</i>	310
TASK 4: MP-BGP.....	311
<i>Objectives</i>	311
<i>Steps</i>	311
<i>Configure Route Distinguisher and import routes into BGP routing table</i>	311
<i>Route Targets</i>	314
<i>Simple method</i>	314
<i>Verify connectivity</i>	317
<i>Route-map based RT</i>	318

LAB 09: MULTICAST	323
OBJECTIVES	323
REQUIREMENTS	323
<i>Introduction</i>	323
SCENARIO	324
<i>Setup</i>	324
<i>Multicast routing</i>	324
TASK 1: LOAD THE START CONFIGURATIONS	326
<i>Objectives</i>	326
<i>Steps</i>	326
<i>Introduction</i>	326
<i>Explore the environment</i>	326
TASK 2: CONFIGURE PIM SM	329
<i>Objectives</i>	329
<i>Steps</i>	329
<i>Configure the C-BSR</i>	329
<i>Enable PIM-SM on all routing devices</i>	330
<i>Configure Sw1 as C-RP</i>	334
TASK 3: CONFIGURE IGMP AND VERIFY MULTICAST OPERATION	336
<i>Objectives</i>	336
<i>Steps</i>	336
<i>IGMP Querier</i>	336
<i>IGMP Snooping</i>	340
<i>Test Multicast and verify status on Sw1</i>	344
<i>Test without multicast receiver</i>	344
<i>Test with multicast receiver</i>	348
<i>Review the status of the multicast route</i>	348
<i>Test with mtrace</i>	356
<i>Optional: PIM Diagnostics Dump</i>	358
TASK 4: CONFIGURE DISTRIBUTED RP	360
<i>Objectives</i>	360
<i>Steps</i>	360
<i>Configure regional RP</i>	360
<i>Verify the registration traffic on the regional RP</i>	361
OPTIONAL TASK 5: UNDERSTAND THE MAC - IP MULTICAST RELATION	364
<i>Objectives</i>	364
<i>Steps</i>	364
<i>Explore the Multicast IP to Ethernet mapping</i>	364
TASK 6: VSX MULTICAST FAILOVER	366
<i>Objectives</i>	366
<i>Steps</i>	366
<i>VSX Active Active Configuration</i>	366
<i>Understand the VSX Multicast Routing failover</i>	366
<i>Verify the failover</i>	371
TASK 7: IGMP ACL	374
<i>Objectives</i>	374
<i>Steps</i>	374
<i>IGMP Querier ACL</i>	374
<i>IGMP Snooping ACL</i>	375
<i>IGMP debug</i>	377
LAB 10: QUALITY OF SERVICE	378

OBJECTIVES	378
REQUIREMENTS	378
SCENARIO	379
TASK 1: QOS POLICIES	381
Objectives	381
Steps	381
Review default queuing and QOS maps	381
Configure QOS Policy on Sw2	383
Verify the Remark - last hop switch	384
Default Queue and Trust modes	386
Verify the impact of the QOS trust modes	387
On Sw2, you will now review the local traffic handling	388
TASK 2: RATE LIMITER AND TRAFFIC SHAPING	391
Objectives	391
Steps	391
Inbound Rate Limiters	391
Define the Rate Limiter on the Aggregation VSX	391
Troubleshoot the VSX-sync	391
Verify the rate limiter	395
Outbound Traffic shaping	395
Rate limit and burst size	396
LAB 11: DYNAMIC SEGMENTATION	397
OBJECTIVES	397
REQUIREMENTS	397
SCENARIO	398
TASK 1: PREPARE THE CONFIGURATION	402
Objectives	402
Steps	402
Configuration	402
Table with RADIUS test user accounts	402
TASK 2: CONFIGURE BASIC 802.1X AUTHENTICATION	403
Objectives	403
Steps	403
Basic RADIUS and 802.1X configuration	403
Statistics EAP and RADIUS	406
Review the debug output for a successful authentication	409
TASK 3: RADIUS TROUBLESHOOTING	411
Objectives	411
Steps	411
Invalid Aruba-User-Role	411
Cached re-authentication	414
Verify the operation without cached re-authentication	415
Critical Role for Authentication	417
Configure the switch with cached re-authentication	418
Understand the need for RADIUS tracking	420
Recover the link and verify the operation	424
TASK 4: ONBOARDING PRECEDENCE ORDER	426
Objectives	426
Steps	426
Device Profile with Local MAC Authentication (LMA)	426
Device Profile Fallback role	429
Device Profile with AAA enabled on the port: Onboarding Precedence	430

<i>Reject RADIUS Response</i>	431
<i>Optional: Device Profile with RADIUS Failure</i>	432
TASK 5: USER-BASED TUNNELING	434
<i>Objectives</i>	434
<i>Steps</i>	434
<i>Configuring UBT</i>	434
<i>Review Normal UBT Authentication</i>	437
<i>Invalid Gateway Role</i>	440
<i>MTU for Tunneled users</i>	442
<i>Configure Jumbo support</i>	445
<i>Default role on the Aruba Gateway</i>	446
TASK 6: USER-BASED TUNNELING QOS	449
OBJECTIVES	449
STEPS.....	449
<i>Apply a QOS Policy to the Switch user role</i>	449
<i>Optional: Trace traffic and compare inner DSCP to outer GRE DSCP markings</i>	452
TASK 7: DOWNLOADABLE USER ROLES.....	456
<i>Objectives</i>	456
<i>Steps</i>	456
<i>ClearPass DUR account configuration</i>	456
<i>ClearPass Certificate Validation</i>	457
<i>Install the Trust Anchor certificate on the Switch</i>	460
<i>Review normal DUR authentication and role</i>	465
LAB 12: NETWORK SECURITY FEATURES.....	468
OBJECTIVES	468
REQUIREMENTS.....	468
SCENARIO	469
<i>TCAM Resource Usage</i>	469
<i>Hitcounts</i>	469
<i>Control Plane Security</i>	469
<i>Control Plane Policing</i>	470
<i>DHCP Snooping</i>	470
<i>ARP Inspection</i>	470
TASK 1: ACL AND RESOURCE USAGE.....	472
<i>Objectives</i>	472
<i>Steps</i>	472
<i>Configuration</i>	472
<i>Prepare the a routed ACL</i>	472
<i>Update the ACL using object-groups</i>	473
<i>Hitcounts and Statistics</i>	476
<i>Optional: Unique Policy per interface for unique counters</i>	478
TASK 2: CONTROL PLANE ACL.....	480
<i>Objectives</i>	480
<i>Steps</i>	480
<i>COPP ACL on VRF mgmt</i>	480
<i>Prepare for failure</i>	481
<i>COPP ACL on VRF default</i>	483
TASK 3: CONTROL PLANE POLICING.....	486
<i>Objectives</i>	486
<i>Steps</i>	486
<i>VRF mgmt behavior</i>	489
OPTIONAL TASK 4: MONITORING CONTROL PLANE WITH NAE	490

Objectives.....	490
Steps.....	490
TASK 5: DHCP SNOOPING.....	493
Objectives.....	493
Steps.....	493
Optional: Authorized Server.....	495
DHCP Bindings Database.....	496
TASK 6: ARP INSPECTION.....	498
Objectives.....	498
Steps.....	498
Configure ARP Inspection.....	498
Optional: Verify IP Binding database dependency.....	501
Verify ARP Inspection filtering.....	503
Endpoint with static IP.....	504
LAB 13: IPV6.....	506
OBJECTIVES.....	506
REQUIREMENTS.....	506
SCENARIO.....	507
OSPFv3.....	507
SLAAC.....	507
DHCPv6 - Stateless.....	507
DHCPv6 - Stateful.....	508
TASK 1: CONFIGURE IPV6 LINKS AND OSPFv3.....	510
Objectives.....	510
Steps.....	510
Agg1 - Agg2 Link.....	510
OSPFv3 on Agg1 and Agg2.....	512
Configure Sw1 as OSPFv3 router.....	514
TASK 2: CONFIGURE IPV6 ON CLIENT SUBNET USING SLAAC.....	517
Objectives.....	517
Steps.....	517
Aggregation configuration.....	517
Client verification.....	517
TASK 3: CONFIGURE IPV6 ON CLIENT SUBNET USING DHCP.....	521
Objectives.....	521
Steps.....	521
Configure PC4.....	521
Configure DHCPv6 Server.....	522
Configure SVI12 for Stateless DHCPv6.....	523
Verify Stateless DHCPv6.....	524
Configure SVI 12 for Stateful DHCPv6.....	526
Verify Stateful DHCPv6.....	526
LAB 14: TROUBLE TICKETS HINTS AND SOLUTIONS.....	529
TROUBLE TICKET HINTS AND SOLUTIONS.....	530
Trouble Ticket 01.01.....	530
Trouble Ticket 01.02.....	533
Trouble Ticket 02.01.....	538
Trouble Ticket 03.01.....	541
Trouble Ticket 04.01.....	543
Trouble Ticket 05.01.....	546
Trouble Ticket 06.01.....	549

<i>Trouble Ticket 07.01</i>	551
<i>Trouble Ticket 08.01</i>	559
<i>Trouble Ticket 08.02</i>	562
<i>Trouble Ticket 08.03</i>	565
<i>Trouble Ticket 08.04</i>	568
<i>Trouble Ticket 08.05</i>	570
<i>Trouble Ticket 08.06</i>	572
<i>Trouble Ticket 08.07</i>	577
<i>Trouble Ticket 09.01 - Routing and Multicast</i>	581

LAB 14: TROUBLE TICKETS 590

TASK 1: PREPARE THE TROUBLESHOOTING SETUP	591
<i>Introduction</i>	591
TROUBLE TICKET SETUP 01	593
<i>Steps</i>	593
<i>Trouble Ticket 01.01</i>	594
<i>Trouble Ticket 01.02</i>	595
TROUBLE TICKET SETUP 02	596
<i>Steps</i>	596
<i>Trouble Ticket 02.01</i>	596
TROUBLE TICKET SETUP 03	597
<i>Steps</i>	597
<i>Trouble Ticket 03.01</i>	597
TROUBLE TICKET SETUP 04	598
<i>Steps</i>	598
<i>Trouble Ticket 04.01</i>	598
TROUBLE TICKET SETUP 05	599
<i>Steps</i>	599
<i>Trouble Ticket 05.01</i>	599
TROUBLE TICKET SETUP 06	600
<i>Steps</i>	600
<i>Trouble Ticket 06.01</i>	600
TROUBLE TICKET SETUP 07	601
<i>Steps</i>	601
<i>Trouble Ticket 07.01</i>	601
TROUBLE TICKET SETUP 08	602
<i>Steps</i>	602
<i>Trouble Ticket 08.01</i>	602
<i>Trouble Ticket 08.02</i>	602
<i>Trouble Ticket 08.03</i>	602
<i>Trouble Ticket 08.04</i>	603
<i>Trouble Ticket 08.05</i>	603
<i>Trouble Ticket 08.06</i>	603
<i>Trouble Ticket 08.07</i>	604
TROUBLE TICKET SETUP 09	605
<i>Steps</i>	605
<i>Trouble Ticket 09.01</i>	605

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.



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.

```
erase all zeroize
```

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.

```
Agg1# show version
```

```
-----  
ArubaOS-CX  
(c) Copyright 2017-2020 Hewlett Packard Enterprise Development LP  
-----
```

```
Version       : GL.10.05.0021  
Build Date    : 2020-10-29 10:43:38 PDT  
Build ID      : ArubaOS-CX:GL.10.05.0021:ef0fef7f4b9d:202010291658  
Build SHA     : ef0fef7f4b9d2732c70bed9fa25589c665319634  
Active Image  : primary
```

```
Service OS Version : GL.01.05.0003  
BIOS Version       : GL-01-0013
```

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.

```
Agg1# show image
```

```
-----  
ArubaOS-CX Primary Image  
-----
```

```
Version : GL.10.04.0030  
Size    : 415 MB  
Date    : 2020-01-28 00:14:09 PST  
SHA-256 : 3fa375918983bd020187db98a3ea94300787889bd023f09e5444c572b7cc4933  
-----
```

```
ArubaOS-CX Secondary Image  
-----
```

```
Version : GL.10.05.0021  
Size    : 425 MB  
Date    : 2020-10-29 10:43:38 PDT  
SHA-256 : 9a7e0739e9f0bc0bb46be3a64371baa2f8d6e3016e78a5ef7f0a14c7e78e1431  
-----
```

```
Default Image : secondary  
-----
```

```
Management Module 1/1 (Active)  
-----
```

```
Active Image      : secondary  
Service OS Version : GL.01.05.0003  
BIOS Version      : GL-01-0013  
-----
```

- 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.

```
Agg1# boot set-default ?
```

```
primary    Set the default boot image to primary for future reboots  
secondary  Set the default boot image to secondary for future reboots
```

- c. Verify the correct version has been set with the show image command and reboot the system.

```
Agg1# show image
```

```
Agg1# boot 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.

```
Agg1# show interface mgmt
Address Mode           : dhcp
Admin State           : up
Mac Address           : b8:d4:e7:41:38:01
IPv4 address/subnet-mask :
Default gateway IPv4   :
IPv6 address/prefix    :
IPv6 link local address/prefix: fe80::bad4:e7ff:fe41:3801/64
Default gateway IPv6   :
Primary Nameserver     :
Secondary Nameserver   :
```

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.

- 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.

```
Agg1# conf t
Agg1(config)# interface mgmt
Agg1(config)# ip static 10.251.x.2/24
Agg1(config)# default-gateway 10.251.x.254
Agg1(config)# end
Agg1# show interface mgmt

Address Mode           : static
Admin State            : up
Mac Address             : b8:d4:e7:41:38:01
IPv4 address/subnet-mask : 10.251.x.2/24
Default gateway IPv4    : 10.251.x.254
IPv6 address/prefix     :
IPv6 link local address/prefix: fe80::bad4:e7ff:fe41:3801/64
Default gateway IPv6    :
Primary Nameserver      :
Secondary Nameserver    :
```

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

```
Agg1# copy running-config checkpoint oobm
```

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

- Review the checkpoint list and verify that checkpoint **oobm** is listed.

```
Agg1# show checkpoint list
CPC20201116095626
oobm
```

- 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

- 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

```
Agg1# show interface brief
```

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'

Port Description	Native VLAN	Mode	Type	Enabled	Status	Reason	Speed (Mb/s)
1/1/1	--	routed	SFP+DAC1	no	down	Group speed mismatch	--

----- Output Omitted -----

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.

```
# ping 10.x.1.3 repetitions 2
PING 10.x.1.3 (10.x.1.3) 100(128) bytes of data.
108 bytes from 10.x.1.3: icmp_seq=1 ttl=64 time=0.125 ms
108 bytes from 10.x.1.3: icmp_seq=2 ttl=64 time=0.127 ms

--- 10.x.1.3 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1060ms
rtt min/avg/max/mdev = 0.125/0.126/0.127/0.001 ms

# ping 10.x.1.4 repetitions 2
PING 10.x.1.4 (10.x.1.4) 100(128) bytes of data.
108 bytes from 10.x.1.4: icmp_seq=1 ttl=64 time=0.150 ms
108 bytes from 10.x.1.4: icmp_seq=2 ttl=64 time=0.168 ms

--- 10.x.1.4 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1060ms
rtt min/avg/max/mdev = 0.150/0.159/0.168/0.009 ms

# ping 10.x.1.5 repetitions 2
PING 10.x.1.5 (10.x.1.5) 100(128) bytes of data.
108 bytes from 10.x.1.5: icmp_seq=1 ttl=64 time=0.155 ms
108 bytes from 10.x.1.5: icmp_seq=2 ttl=64 time=0.165 ms

--- 10.x.1.5 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1020ms
rtt min/avg/max/mdev = 0.155/0.160/0.165/0.005 ms
```

3. On Agg1 ping to the RouterA using the default routing table

```
# ping 10.255.101.11 repetitions 2
PING 10.255.101.11 (10.255.101.11) 100(128) bytes of data.
108 bytes from 10.255.101.11: icmp_seq=1 ttl=64 time=2.06 ms
108 bytes from 10.255.101.11: icmp_seq=2 ttl=64 time=1.98 ms
```

```
--- 10.255.101.11 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1000ms
rtt min/avg/max/mdev = 1.987/2.028/2.069/0.041 ms
```

4. On Agg1 ping to the ClearPass host (10.252.1.23) using VLAN1 as source IP address.

```
# ping 10.252.1.23 repetitions 2 source vlan1

PING 10.252.1.23 (10.252.1.23) from 10.x.1.2 : 100(128) bytes of data.
108 bytes from 10.252.1.23: icmp_seq=1 ttl=64 time=1.79 ms
108 bytes from 10.252.1.23: icmp_seq=2 ttl=64 time=1.78 ms

--- 10.255.101.11 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1001ms
rtt min/avg/max/mdev = 1.787/1.790/1.794/0.042 ms
```

OOBM Connection tests

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

```
# ping 10.251.x.200 repetitions 2 vrf mgmt

PING 10.251.x.200 (10.251.x.200) 100(128) bytes of data.
108 bytes from 10.251.x.200: icmp_seq=1 ttl=64 time=0.487 ms
108 bytes from 10.251.x.200: icmp_seq=2 ttl=64 time=0.551 ms

--- 10.251.x.200 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1025ms
rtt min/avg/max/mdev = 0.487/0.519/0.551/0.032 ms
```

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'.

```
# copy running-config tftp://10.251.x.91/asts-lab01-agg1.txt ?
```

```
cli    CLI format
json   JSON format
```

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

```
# copy running-config tftp://10.251.x.91/asts-lab01-agg1.txt cli vrf mgmt

% Total      % Received % Xferd  Average Speed   Time    Time       Time  Current
   Dload  Upload   Total     Spent    Left     Speed
100 1517      0      0 100 1517      0 59511 --:--:-- --:--:-- --:--:-- 59511
100 1517      0      0 100 1517      0 59338 --:--:-- --:--:-- --:--:-- 59338
Copying configuration: [Success]
```

9. Upload the running configuration as JSON file to PC1, use filename asts-lab01-agg1-json.txt

```
# copy running-config tftp://10.251.x.91/asts-lab01-agg1-json.txt json vrf mgmt

Copying configuration: [[]] % Total      % Received % Xferd  Average Speed   Time    Time
Time Current
   Dload  Upload   Total     Spent    Left     Speed
100 17324      0      0 100 17324      0 638k --:--:-- --:--:-- --:--:-- 638k
100 17324      0      0 100 17324      0 635k --:--:~ --:~:~ --:~:~ 635k
Copying configuration: [Success]
```

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'

```
# conf t
(config)# hostname test
test(config)#
```

15. To compare your current configuration with a saved checkpoint, use the checkpoint diff command.

```
test(config)# checkpoint diff running-config asts-lab01-done
--- /tmp/running-config1605784056185      2020-11-19 06:07:36.184469318 -0500
+++ /tmp/asts-lab01-done1605784056187      2020-11-19 06:07:36.186469318 -0500
@@ -1,7 +1,7 @@
!
!Version ArubaOS-CX GL.10.05.0021
!export-password: default
-hostname test
+hostname Tx-Agg1-lab01
 user admin group administrators password ciphertext
 AQBapUs4PRSyIHOf50An+ESWxxIICJf51f6uToSeMcr77M4/YgAAA0uwW3fRfUgFJ0pJJFqX1i0qQn0VwI08iMnEGqG4U
 owB4akn7CL60UsdibdQKFxfVyYedB+oDrKwb+8+0dGp4F8dGrH5tCJhpNcIthvzBQwR5gLXp1He/2SPKyQGFxvr3+Ct
 user labadmin group administrators
 clock timezone us/eastern
```

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 a 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.

```
# show event -r -n 10 | include config
```

```
2020-09-29T09:44:08.362524-04:00 TX-Agg1 hpe-restd[3038]: Event|4613|LOG_INFO|AMM|-|admin has written a new switch configuration to 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.

```
show checkpoint list
```

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.

```
boot system
y
```

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.

```
Agg1# show interface physical
```

Port		Link	Admin	Speed		Flow-Control		PoE	Power
Port	Type	Status	Config	Status	Config	Status	Config	(Watts)	
State Information									
1/1/1	SFP+DAC1	up	up	10G	n/a	off	off	--	
10G			sw1						
----- Output Omitted -----									
1/1/46	SFP28DAC0.65	up	up	25G	auto	off	off	--	
25G			agg2						

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.

```
Agg1# show system interface-group
```

Group	Speed	Member Ports	Mismatched Ports
1	10g	1/1/1-1/1/12	
2	25g	1/1/13-1/1/24	
3	25g	1/1/25-1/1/36	
4	25g	1/1/37-1/1/48	

3. Review the link-status

```
Agg1# show interface link-status
```

```

-----
Port      Type      Physical  Link
          Link State Transitions
-----
1/1/1     SFP+DAC1  up        3
----- Output Omitted -----

```

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.

```

Agg1# show interface transceiver detail
Transceiver in 1/1/1
Interface Name      : 1/1/1
Type                : SFP+DAC1
Connector Type      : Copper Pigtail
Transfer Distance   : 0.00km (SMF), 0m (OM1), 0m (OM2), 0m (OM3)
Diagnostic Support   : None
Product Number      : J9281D
Serial Number       : CN99KBZ7MJ
Part Number         : 8121-1300
----- Output Omitted -----

```

Interface statistics

5. Review the interface statistics.

```

Agg1# show interface statistics
-----
Interface          RX Bytes      RX Packets      RX Drops
TX Bytes           TX Packets     TX Drops        RX Broadcast    RX Multicast
TX Broadcast       TX Multicast
-----
1/1/1              1016211        6527            0
5085822            42092          0                3171            2881
6751               34801
1/1/2              0              0                0                0
0                  0              0                0                0
----- Output Omitted -----

```

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

```
Agg1# show interface statistics non-zero
```

Interface		RX Bytes		RX Packets		RX Drops
TX Bytes	TX Packets		TX Drops	RX Broadcast		RX Multicast
TX Broadcast	TX Multicast					
1/1/1		1021025		6551		0
5095495	42177		0	3183		2891
6763	34872					
1/1/7		101215720		1084325		0
103818736	1111901		0	1859		13
1884	9898					
1/1/46		5003357		40635		0
2183098	17180		0	4883		34864
5066	11203					

7. Clear the statistics and review the statistics again. It may take a few seconds before new statistics are available.

```
Agg1# clear interface statistics
```

```
Agg1# show interface statistics non-zero
```

Interface		RX Bytes		RX Packets		RX Drops
TX Bytes	TX Packets		TX Drops	RX Broadcast		RX Multicast
TX Broadcast	TX Multicast					

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.

```
Agg1# clear interface statistics global
```

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.

```
Agg1# copy support-files all tftp://10.251.x.91/agg1-support-all.tar.gz vrf mgmt.
The operation to copy all support files could take a while to complete.
Do you want to continue (y/n)? y
Collection in progress...

show tech output collection is in progress...
Done

Done
Attempting to copy...
  % Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
   % Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
  Dload  Upload  Total   Spent    Left     Speed
100 15.8M    0     0  100 15.8M      0  4537k   0:00:03  0:00:03 --:--:-- 4566k
100 15.8M    0     0  100 15.8M      0  4537k   0:00:03  0:00:03 --:--:-- 4537k
Done
```

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.

```
Agg1# conf t
Agg1(config)# interface lag 256
Agg1(config-lag-if)# lacp mode active
Agg1(config-lag-if)# no shut
Agg1(config)# exit
Agg1(config)# interface 1/1/46-1/1/47
Agg1(config-if-<1/1/46-1/1/47>)# lag 256
Agg1(config-if-<1/1/46-1/1/47>)# no shut
Agg1(config-if-<1/1/46-1/1/47>)# end
```

19. Review the accounting log for the last 10 entries. The last entries should show the LAG commands.

```
Agg1# show accounting log last 10
----
type=USYS_CONFIG msg=audit(Sep 29 2020 04:39:48.863:512) : msg='rec=ACCT_CMD op=stop
session=CONSOLE timezone=US/Eastern user=admin priv-lvl=15 auth-method=LOCAL auth-type=LOCAL
service=shell isconfig=yes data="interface lag 256" hostname=TX-Agg1 addr=0.0.0.0
res=success'
----- Output Omitted -----
----
type=USYS_CONFIG msg=audit(Sep 29 2020 04:50:35.648:536) : msg='rec=ACCT_CMD op=stop
session=SSH timezone=US/Eastern user=admin priv-lvl=15 auth-method=LOCAL auth-type=LOCAL
service=shell isconfig=no data="show interface brief" hostname=TX-Agg1 addr=10.251.x.91
res=success'
----- Output Omitted -----
```

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!

```
Agg1# show lacp interfaces
----- Output Omitted -----
```

Intf	Aggr Name	Port Id	Port Pri	State	System-ID	System Pri	Aggr Key	Forwarding State
1/1/46	lag256	47	1	ALFOE	b8:d4:e7:41:38:00	65534	256	lacp-block
1/1/47	lag256							down

```
----- Output Omitted -----
```

21. Review the status of the LAG 256 interface

```
Agg1# show int lag256
Aggregate lag256 is down
Admin state is up
State information : Disabled by LACP or LAG
Description :
MAC Address      : b8:d4:e7:41:38:00
Aggregated-interfaces : 1/1/46 1/1/47
Aggregation-key  : 256
Aggregate mode   : active
Speed           : 0 Mb/s
L3 Counters: Rx Disabled, Tx Disabled
qos trust none
```

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.

```
Agg1# show event -r -n 20
-----
Event logs from current boot
-----
2020-09-29T04:41:35.075569-04:00 TX-Agg1 lacpd[1862]: Event|1321|LOG_INFO|AMM|1/1|LAG 256
State change for interface 1/1/46: Actor state: ALFOE, Partner state PLFOEX
2020-09-29T04:41:35.075331-04:00 TX-Agg1 lacpd[1862]: Event|1311|LOG_WARN|AMM|1/1|Partner is
lost (timed out) for interface 1/1/46 LAG sport: 1. State: Expired State -> Defaulted State
----- Output Omitted -----
```

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.

```
Agg1# show event -d lacpd -r -n 10
```

```
-----
Event logs from current boot
-----
```

```
2020-09-29T04:41:35.075569-04:00 TX-Agg1 lacpd[1862]: Event|1321|LOG_INFO|AMM|1/1|LAG 256
State change for interface 1/1/46: Actor state: ALFOE, Partner state PLFOEX
2020-09-29T04:41:35.075331-04:00 TX-Agg1 lacpd[1862]: Event|1311|LOG_WARN|AMM|1/1|Partner is
lost (timed out) for interface 1/1/46 LAG sport: 1. State: Expired State -> Defaulted State
----- Output Omitted -----
```

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

```
Agg1# terminal-monitor
Terminal-monitor is enabled successfully
```

25. Enable and disable interface LAG256 again

```
Agg1# configure terminal
Agg1(config)# interface lag 256
Agg1(config-lag-if)# shut
Agg1(config-lag-if)# no shutdown
Agg1(config-lag-if)# end
```

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

A: No.

26. Review the terminal-monitor severity options

```
Agg1# show terminal-monitor

Terminal-monitor is enabled
-----
Notify      | Severity  | Filter
-----
all         | err       | Nil
-----
```

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

```
Agg1# terminal-monitor severity info
```

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.

```
Agg1# configure terminal
Agg1(config)# interface lag 256
Agg1(config-lag-if)# shutdown
Agg1(config-lag-if)# no shutdown
Agg1(config-lag-if)# end

2020-09-29T06:54:56.341566-0400 portd[1869] <INFO> Event|1703|LOG_INFO|AMM|1/1|Interface
lag256, configured administratively down
2020-09-29T06:54:56.345845-0400 ops-switchd[3336] <INFO> Event|1702|LOG_INFO|AMM|1/1|L3-
Interface lag256, deleted
2020-09-29T06:54:56.355218-0400 intfd[1857] <INFO> Event|404|LOG_INFO||Link status for
interface 1/1/46 is down
2020-09-29T06:54:59.341873-0400 portd[1869] <INFO> Event|1703|LOG_INFO|AMM|1/1|Interface
lag256, configured administratively up
----- Output Omitted -----
```

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

```
Agg1# terminal-monitor filter lacpd severity info
```

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.

```
Agg2# conf t
Agg2(config)# interface lag 256
Agg2(config-lag-if)# lacp mode active
Agg2(config-lag-if)# no shut
Agg2(config)# exit
Agg2(config)# interface 1/1/46-1/1/47
Agg2(config-if-1/1/46-1/1/47)# lag 256
Agg2(config-if-1/1/46-1/1/47)# no shut
Agg2(config-if-1/1/46-1/1/47)# end
Agg2# show lacp interfaces

----- Output Omitted -----
Actor details of all interfaces:
-----
```

Intf	Aggr Name	Port Id	Port Pri	State	System-ID	System Pri	Aggr Key	Forwarding State
1/1/46	lag256	47	1	ALFNCD	b8:d4:e7:41:68:00	65534	256	up
1/1/47	lag256	48	1	ALFNCD	b8:d4:e7:41:68:00	65534	256	up

```
----- Output Omitted -----
```


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

```
2020-09-29T07:00:13.800501-0400 lacpd[1862] <INFO> Event|1309|LOG_INFO|AMM|1/1|Partner is
detected for interface 1/1/46 LAG 256 : 65534,b8:d4:e7:41:68:00. Actor state: ALFO, partner
state ALFOE
2020-09-29T07:00:13.800667-0400 lacpd[1862] <INFO> Event|1321|LOG_INFO|AMM|1/1|LAG 256 State
change for interface 1/1/46: Actor state: ALFO, Partner state ALFOE
----- Output Omitted -----
2020-09-29T07:00:15.108373-0400 lacpd[1862] <INFO> Event|1309|LOG_INFO|AMM|1/1|Partner is
detected for interface 1/1/47 LAG 256 : 65534,00:00:00:00:00:00. Actor state: ALFOE, partner
state PLFO
2020-09-29T07:00:15.108789-0400 lacpd[1862] <INFO> Event|1321|LOG_INFO|AMM|1/1|LAG 256 State
change for interface 1/1/47: Actor state: ALFOE, Partner state PLFO
----- Output Omitted -----
```

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

```
Agg1# no terminal-monitor
```

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.

```
Sw1# configure terminal
Sw1(config)# interface 1/1/25-1/1/26
Sw1(config-if-<1/1/25-1/1/26>)# mtu 9198
Sw1(config-if-<1/1/25-1/1/26>)# exit
```

3. Configure SVI 1 on Sw1 with an IP MTU of 9198.

```
Sw1(config)# interface vlan1
Sw1(config-if-vlan)# ip mtu 9198
Sw1(config-if-vlan)# end
```

4. On Sw1, ping to 10.x.1.2 (the IP of Agg1) to confirm a normal sized packet works fine.

```
Sw1# ping 10.x.1.2 repetitions 2
PING 10.x.1.2 (10.x.1.2) 100(128) bytes of data.
108 bytes from 10.x.1.2: icmp_seq=1 ttl=64 time=0.143 ms
108 bytes from 10.x.1.2: icmp_seq=2 ttl=64 time=0.166 ms

--- 10.x.1.2 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1002ms
rtt min/avg/max/mdev = 0.143/0.154/0.166/0.016 ms
```

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.

```
Sw1# ping 10.x.1.2 repetitions 2 datagram-size 2000 do-not-fragment
PING 10.x.1.2 (10.x.1.2) 2000(2028) bytes of data.

--- 10.x.1.2 ping statistics ---
2 packets transmitted, 0 received, 100% packet loss, time 1010ms
```

6. On Agg1, attempt to detect the drop reason, review the interface 1/1/1 statistics.

```
Agg1# show interface 1/1/1 statistics
```

Interface		RX Bytes		RX Packets		RX Drops
TX Bytes	TX Packets		TX Drops	RX Broadcast		RX Multicast
TX Broadcast	TX Multicast					
1/1/1		902861		6838		0
387993	3481		0	790		5954
557	2823					

Q: Any dropped frames?

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

7. Review the available 'diag?' command options.

```
Agg1# diag ?
```

```
<cr>
```

8. Enable diagnostics mode.

```
Agg1# diag
```

9. Review the available 'diag?' command options again.

```
diag ?
```

```
audit-failure-notification  Configure audit failure notification
bgp                          IP information
crash-notification          Configure crash console log
erps                         ERPS MIB dump
evpn                         evpn information
fan                          Fan module diagnostic test
fruread                      Read FRU EEPROM
fruwrite                     Write FRU EEPROM
hardware-vtep-dump           Display Hardware VTEP DB information
interface                    Interface information
ledtest                      Diagnostic level of LED testing
ospf                         OSPF information
ospfv3                       Open Shortest Path First (OSPF) for IPv6
ovsdb-dump                   Display or copy compressed(tar.gz format) OVSDB
                             information
rib                          RIB routes
setkernelloglevel           Set kernel Log Level
show                         Show running system information
show-hw-ut-status            Display hardware status of unsupported
                             transceivers
showkernelloglevel           Displays Current Kernel Log Level
utilities                    Linux tool utilities
```

10. Explore the 'diag interface 1/1/1 statistics' command.

```
Agg1# diag interface 1/1/1 statistics
```

RX Statistic	Packets
Good (L2 + L3) unicast	571
Good IPv4 L3 unicast	447
Good IPv6 L3 unicast	0
Invalid VLAN drops	0
Multicast drops	0
Non forwarding ingress port state drops	0
ACL drops	0
IPinIP tunnel drops	0
802.1d frame size exceeded drops	4
802.1d frame forwarding filter discards	0
802.3 frame opcode not supported drops	0
802.3 frame size exceeded drops	4
802.3 non-integral octet length FCS check drops	0
802.3 integral octet length FCS check drops	0
802.3 internal MAC sublayer errors	0
802.3 LLC data octet errors	0
IPv4 header errors	0
IPv6 header errors	0
IPv4 L3 discards	0
IPv6 L3 discards	0
IPv6 address errors	0
IP multicast drops	0

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.

```
interface 1/1/1-1/1/48
mtu 9198
```

12. Reset the Sw1 SVI1 MTU to 1500.

```
Sw1# conf t
Sw1(config)# interface vlan1
Sw1(config-if-vlan)# ip mtu 1500
Sw1(config-if-vlan)# end
```

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

```
Agg2# conf t
Agg2(config)# vlan 11
Agg2(config-vlan-11)# exit
Agg2(config)# int vlan 11
Agg2(config-if-vlan)# ip address 10.x.11.3/24
Agg2(config-if-vlan)# end
```

3. Configure LAG256 as VLAN trunk and allow all VLANs.

```
Agg2# conf t
Agg2(config)# interface lag 256
Agg2(config-lag-if)# no routing
Agg2(config-lag-if)# vlan trunk allowed all
Agg2(config-lag-if)# end
```

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.

```
Agg2# ping 10.x.11.2 repetitions 1000
```

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.

```
Agg1# show int 1/1/46 extended
Agg1# show int 1/1/47 extended
```

Q: Are there any dropped frames?

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

```
# show interface 1/1/46 extended | in "RX Filtered"
RX Filtered                               9483

# show interface 1/1/46 extended | in "RX Filtered"
RX Filtered                               9494
```

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

```
Agg1# diag interface 1/1/46 statistics

-----
RX Statistic                               Packets
-----
Good (L2 + L3) unicast                     2782
Good IPv4 L3 unicast                       727
Good IPv6 L3 unicast                       0
Invalid VLAN drops                         2182
----- Output Omitted -----
```

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.

```
Agg1# conf t
Agg1(config)# vlan 11
Agg1(config-vlan-11)# exit
Agg1(config)# int vlan 11
Agg1(config-if-vlan)# ip address 10.x.11.2/24
Agg1(config-if-vlan)# exit
Agg1(config)# interface lag 256
Agg1(config-lag-if)# no routing
Agg1(config-lag-if)# vlan trunk allowed all
Agg1(config-lag-if)# end
```

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.

```
Sw1# show debug
-----
-----
module          sub_module          severity  vlan  port  ip
mac             instance            client
column
-----
-----
```

2. Review default debug destinations on Sw1.

```
Sw1# show debug destination
-----
show debug destination
-----
```

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

3. Clear the debug buffer on Sw1.

```
Sw1# clear debug buffer
```

4. Review the debug buffer on Sw1.

```
Sw1# show debug buffer
-----
show debug buffer
-----
No dlog messages logged in the system
```

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.

```
Sw1# debug ndm statemachine
Sw1# debug ndm unresolvedunicast
```

7. Clear the ARP table on Sw1.

```
Sw1# clear arp
```

8. Review the contents of the ARP table on Sw1.

```
Sw1# show arp
```

IPv4 Address	MAC	Port	Physical Port	State
10.x.1.2	b8:d4:e7:41:38:00	vlan1	1/1/25	reachable
Total Number Of ARP Entries Listed- 1.				

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.

```
Sw1# terminal-monitor severity debug filter 10.x.1.3
Terminal-monitor is enabled successfully
```

10. On Sw1, ping 10.x.1.3, then wait a few moments.

```
Sw1# ping 10.x.1.3 repetitions 2
PING 10.x.1.3 (10.x.1.3) 100(128) bytes of data.
108 bytes from 10.x.1.3: icmp_seq=1 ttl=64 time=0.143 ms
108 bytes from 10.x.1.3: icmp_seq=2 ttl=64 time=0.363 ms

--- 10.x.1.3 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1028ms
rtt min/avg/max/mdev = 0.143/0.253/0.363/0.110 ms
```

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:'

```
----- Output Omitted -----
2020-09-29T17:24:22.237140-0400 ndmd[3260] <DEBUG>
debug|LOG_DEBUG|MSTR|1|NDM|NDM_STATEMACHINE|[nd_nbr_mac_update(585)] [VRF:default | NBR:
10.x.1.3 | PORT: vlan1] Updated the MAC from old 00:00:00:00:00:00 to a new MAC :
b8:d4:e7:41:68:00
2020-09-29T17:24:22.237319-0400 ndmd[3260] <INFO>
debug|LOG_INFO|MSTR|1|NDM|NDM_STATEMACHINE|[nd_nbr_fsm_set_nbr_reachable(3047)]
[NBR:10.x.1.3][PORT:vlan1] Neighbor state set to ND_NBR_STATE_REACHABLE
```

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.

```
Sw1# terminal-monitor severity debug filter 10.x.1.30
Terminal-monitor is enabled successfully
```

13. On Sw1, ping 10.x.1.30 a single time (repetition 1)

```
Sw1# ping 10.x.1.30 repetitions 1
PING 10.x.1.30 (10.x.1.30) 100(128) bytes of data.

--- 10.x.1.30 ping statistics ---
1 packets transmitted, 0 received, 100% packet loss, time 0ms
```

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.

```
----- Output Omitted -----
2020-09-29T17:18:59.481937-0400 ndmd[3260] <INFO>
debug|LOG_INFO|MSTR|1|NDM|NDM_STATEMACHINE|[nd_print_nbr_fsm_trace_dlog(131)] [NBR:10.x.1.30]
Neighbor current_state: probe [event: probeTimerExpire]
2020-09-29T17:18:59.482148-0400 ndmd[3260] <INFO>
debug|LOG_INFO|MSTR|1|NDM|NDM_STATEMACHINE|[nd_print_nbr_fsm_trace_dlog(139)] [NBR:10.x.1.30]
|00:00:00:00:00:00|Flag 0X0|probe count- 3|mac_deleted:false | Reprobing:false
2020-09-29T17:18:59.482410-0400 ndmd[3260] <INFO>
debug|LOG_INFO|MSTR|1|NDM|NDM_STATEMACHINE|[nd_nbr_fsm_is_max_probe_count_reached(2157)]
[NBR:10.x.1.30][PORT:vlan1] Max probe count reached
2020-09-29T17:18:59.482625-0400 ndmd[3260] <INFO>
debug|LOG_INFO|MSTR|1|NDM|NDM_STATEMACHINE|[nd_nbr_fsm_is_max_probe_count_reached(2157)]
[NBR:10.x.1.30][PORT:vlan1] Max probe count reached
2020-09-29T17:18:59.482806-0400 ndmd[3260] <INFO>
debug|LOG_INFO|MSTR|1|NDM|NDM_STATEMACHINE|[nd_nbr_fsm_is_max_probe_count_reached(2157)]
[NBR:10.x.1.30][PORT:vlan1] Max probe count reached
```

```
2020-09-29T17:18:59.482986-0400 ndmd[3260] <INFO>
debug|LOG_INFO|MSTR|1|NDM|NDM_STATEMACHINE|[nd_nbr_is_incomplete(2734)]
[NBR:10.x.1.30][PORT:vlan1] Neighbor is INCOMPLETE
```

15. Disable all debugging on Sw1 with the 'no debug all' command.

```
Sw1# no debug all
```

16. Disable the terminal-monitor on Sw1.

```
Sw1# no terminal-monitor
```

17. Review the last 40 lines of the debug buffer for the 'ndm' module on Sw1. This output will also show the NDM entries.

```
Sw1# show debug buffer module ndm lines 40
----- Output Omitted -----
2020-09-
29:17:22:58.776755|ndmd|LOG_INFO|MSTR|1|NDM|NDM_STATEMACHINE|[nd_print_nbr_fsm_trace_dlog(139
)] [NBR:10.x.1.2] |b8:d4:e7:41:38:00|Flag 0X8|probe count- 1|mac_deleted:false |
Reprobing:true
2020-09-
29:17:22:58.776774|ndmd|LOG_INFO|MSTR|1|NDM|NDM_STATEMACHINE|[nd_nbr_fsm_set_nbr_reachable(30
47)] [NBR:10.x.1.2][PORT:vlan1] Neighbor state set to ND_NBR_STATE_REACHABLE
2020-09-
29:17:22:58.776798|ndmd|LOG_INFO|MSTR|1|NDM|NDM_STATEMACHINE|[nd_nbr_fsm_set_nbr_reachable(30
47)] [NBR:10.x.1.2][PORT:vlan1] Neighbor state set to ND_NBR_STATE_REACHABLE
2020-09-29:17:22:58.776887|ndmd|LOG_DEBUG|MSTR|1|NDM|NDM_UNRESOLVEDUNICAST|Dropped 3 log
messages in last 1 seconds (most recently, 1 seconds ago) due to excessive rate
2020-09-
29:17:22:58.776905|ndmd|LOG_DEBUG|MSTR|1|NDM|NDM_UNRESOLVEDUNICAST|[nd_pd_nl_recv_msgs(283)]
[VRF:default] ND kernel msg recvd. sock = 20
2020-09-29:17:22:58.776954|ndmd|LOG_DEBUG|MSTR|1|NDM|NDM_UNRESOLVEDUNICAST|Dropped 1 log
messages in last 1 seconds (most recently, 1 seconds ago) due to excessive rate
2020-09-
29:17:22:58.776969|ndmd|LOG_DEBUG|MSTR|1|NDM|NDM_UNRESOLVEDUNICAST|[nd_rx_receive_kernel_ndm_
message_decode(251)] [NDM-DECODE]Ignoring ipv4 NL MSG 128 for port vlan1 in vrf default
2020-09-
29:17:23:04.181235|ndmd|LOG_DEBUG|MSTR|1|NDM|NDM_UNRESOLVEDUNICAST|[nd_pd_nl_recv_msgs(283)]
[VRF:default] ND kernel msg recvd. sock = 20
2020-09-
29:17:23:04.181289|ndmd|LOG_DEBUG|MSTR|1|NDM|NDM_UNRESOLVEDUNICAST|[nd_rx_receive_kernel_ndm_
message_decode(259)] [NDM-DECODE]Incomplete ipv4 neighbor event received for port vlan1 in
vrf default
2020-09-
29:17:23:04.181313|ndmd|LOG_DEBUG|MSTR|1|NDM|NDM_UNRESOLVEDUNICAST|[nd_nl_msg_translate_to_nb
r_data(122)] [NBR:default|VRF:10.x.1.3] [NL-TRANSLATE] Translated nl_nbr IP address to
nbr_data.
```

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

```
Sw1# show mirror
No mirror sessions exist
```

3. Configure local mirror session 1

```
Sw1# conf t
Sw1(config)# mirror session 1
```

4. Add port 1/1/25 as the source interface and monitor both receive and transmit traffic

```
Sw1(config-mirror-1)# source interface 1/1/25 both
```

5. Add port 1/1/1 as the destination interface

```
Sw1(config-mirror-1)# destination interface 1/1/1
```

6. Enable the mirror

```
Sw1(config-mirror-1)# enable
Sw1(config-mirror-1)# end
```

7. Review the mirror

```
Sw1# show mirror 1
Mirror Session: 1
Admin Status: enable
Operation Status: enabled
Source: vlan rx none
Source: vlan tx none
Source: interface 1/1/25 both
Destination: interface 1/1/1
Output Packets: 31
```

Output Bytes: 5642

8. Review the mirror details

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

A: Depends on your setup

9. Disable the mirror

```
Sw1# conf t
Sw1(config)# mirror session 1
Sw1(config-mirror-1)# disable
```

10. Remove the destination interface from this mirror

```
Sw1(config-mirror-1)# no destination interface
Sw1(config-mirror-1)# show run current

mirror session 1
    source interface 1/1/25 both
Sw1(config-mirror-1)# end
```

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.

```
Sw1# ping 10.x.1.91 repetitions 2

PING 10.x.1.91 (10.x.1.91) 100(128) bytes of data.
108 bytes from 10.x.1.91: icmp_seq=1 ttl=128 time=0.603 ms
108 bytes from 10.x.1.91: icmp_seq=2 ttl=128 time=0.543 ms

--- 10.x.1.91 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1019ms
rtt min/avg/max/mdev = 0.543/0.573/0.603/0.030 ms
```

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.

```
Sw1# conf t
Sw1(config)# mirror session 1
Sw1(config-mirror-1)# destination tunnel 10.x.1.91 source 10.x.1.4
```

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

13. Enable the mirror.

```
Sw1(config-mirror-1)# enable
Sw1(config-mirror-1)# end
```

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)

```
Sw1# ping 10.x.1.2 repetitions 2

PING 10.x.1.2 (10.x.1.2) 100(128) bytes of data.
108 bytes from 10.x.1.2: icmp_seq=1 ttl=64 time=0.167 ms
108 bytes from 10.x.1.2: icmp_seq=2 ttl=64 time=0.151 ms

--- 10.x.1.2 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1000ms
rtt min/avg/max/mdev = 0.151/0.159/0.167/0.008 ms
```

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	25.998666	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 JL635A 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 JL635A 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 JL668A 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

```
Sw1# conf t
Sw1(config)# mirror session 1
Sw1(config-mirror-1)# disable
```

24. Remove the tunnel destination

```
Sw1(config-mirror-1)# no destination tunnel
Sw1(config-mirror-1)# end
```

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

```
Sw1# conf t
Sw1(config)# mirror session 1
Sw1(config-mirror-1)# destination cpu
```

26. Enable the mirror session

```
Sw1(config-mirror-1)# enable
Sw1(config-mirror-1)# end
```

27. Launch the diagnostic utility “tshark” without any options on Sw1. This enables the live view.

```
Sw1# diagnostics
Sw1# diag utilities tshark

Inspecting traffic mirrored to the CPU until Ctrl-C is entered.
```

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.

```
Agg1# ping 10.x.1.91 repetitions 2

PING 10.x.1.91 (10.x.1.91) 100(128) bytes of data.
108 bytes from 10.x.1.91: icmp_seq=1 ttl=128 time=13.9 ms
108 bytes from 10.x.1.91: icmp_seq=2 ttl=128 time=0.506 ms

--- 10.x.1.91 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1001ms
rtt min/avg/max/mdev = 0.506/7.217/13.929/6.712 ms
```

30. In the Sw1 session, packets should be shown in the session

```
----- Output Omitted -----
```



```

Frame 14: 142 bytes on wire (1136 bits), 142 bytes captured (1136 bits) on interface 0
  Interface id: 0 (MirrorRxNet)
    Interface name: MirrorRxNet
    Encapsulation type: Ethernet (1)
    Arrival Time: Sep 30, 2020 07:24:13.107448511 EDT
    [Time shift for this packet: 0.000000000 seconds]
    Epoch Time: 1601465053.107448511 seconds
    [Time delta from previous captured frame: 0.009739428 seconds]
    [Time delta from previous displayed frame: 0.009739428 seconds]
    [Time since reference or first frame: 11.996298611 seconds]
    Frame Number: 14
    Frame Length: 142 bytes (1136 bits)
    Capture Length: 142 bytes (1136 bits)
    [Frame is marked: False]
    [Frame is ignored: False]
    [Protocols in frame: eth:ethertype:ip:icmp:data]
  Ethernet II, Src: b8:d4:e7:41:38:00, Dst: 00:50:56:b1:14:aa
    Destination: 00:50:56:b1:14:aa
      Address: 00:50:56:b1:14:aa
        .... 0. .... = LG bit: Globally unique address (factory default)
        .... 0. .... = IG bit: Individual address (unicast)
    Source: b8:d4:e7:41:38:00
      Address: b8:d4:e7:41:38:00
        .... 0. .... = LG bit: Globally unique address (factory default)
        .... 0. .... = IG bit: Individual address (unicast)
    Type: IPv4 (0x0800)
  Internet Protocol Version 4, Src: 10.x.1.2, Dst: 10.x.1.91
    0100 .... = Version: 4
    .... 0101 = Header Length: 20 bytes (5)
    Differentiated Services Field: 0x00 (DSCP: CS0, ECN: Not-ECT)
      0000 00.. = Differentiated Services Codepoint: Default (0)
      .... ..00 = Explicit Congestion Notification: Not ECN-Capable Transport (0)
    Total Length: 128
    Identification: 0x6f10 (28432)
    Flags: 0x4000, Don't fragment
      0... .... = Reserved bit: Not set
      .1.. .... = Don't fragment: Set
      ..0. .... = More fragments: Not set
      ...0 0000 0000 0000 = Fragment offset: 0
    Time to live: 64
    Protocol: ICMP (1)
    Header checksum: 0xb4f8 [validation disabled]
    [Header checksum status: Unverified]
    Source: 10.x.1.2
    Destination: 10.x.1.91
  Internet Control Message Protocol
    Type: 8 (Echo (ping) request)
  ----- Output Omitted -----

```

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.

```
Sw1# diag utilities tshark delete-file
```

```
Capture file does not exist.
```

34. Start the tshark utility with the file option on Sw1 and leave it running.

```
Sw1# diag utilities tshark file
```

35. On Agg1, ping to PC1. This traffic should go through Sw1 via port 1/1/25.

```
Agg1# ping 10.x.1.91 repetitions 2
```

```
PING 10.x.1.91 (10.x.1.91) 100(128) bytes of data.  
108 bytes from 10.x.1.91: icmp_seq=1 ttl=128 time=13.9 ms  
108 bytes from 10.x.1.91: icmp_seq=2 ttl=128 time=0.506 ms
```

```
--- 10.x.1.91 ping statistics ---  
2 packets transmitted, 2 received, 0% packet loss, time 1001ms  
rtt min/avg/max/mdev = 0.506/7.217/13.929/6.712 ms
```

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 OOB IP and VRF can be used for the copy operation to PC1.

```
Sw1# copy tshark-pcap tftp://10.x.1.91/lab02-sw1.pcap
```

% Total	% Received	% Xferd	Average Speed	Time	Time	Time	Current
			Dload Upload	Total	Spent	Left	Speed
100	1460	0	0	100	1460	0	44508
100	1460	0	0	100	1460	0	43850

Copied successfully.

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

```
Sw1# conf t
Sw1(config)# mirror session 1
Sw1(config-mirror-1)# disable
```

40. Remove the mirror source from the mirror session 1

```
Sw1(config-mirror-1)# no source interface 1/1/25
Sw1(config-mirror-1)# end
```

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

```
Sw1# conf t
Sw1(config)# class ip icmp
Sw1(config-class-ip)# match icmp any any
Sw1(config-class-ip)# exit
```

42. Define a new policy named 'mirror'. Add the 'icmp' class and assign action to mirror to the previously defined mirror session.

```
Sw1(config)# policy mirror
Sw1(config-policy)# class ip icmp action mirror 1
Sw1(config-policy)# exit
```

43. Enable the mirror session

```
Sw1(config)# mirror session 1
Sw1(config-mirror-1)# enable
Sw1(config-mirror-1)# exit
```

44. Apply the policy 'mirror' to the interface 1/1/25 to both in and out directions.

```
Sw1(config)# interface 1/1/25
Sw1(config-if)# apply policy mirror in
Sw1(config-if)# apply policy mirror out
Sw1(config-if)# exit
```

45. Open a second SSH session to switch Sw1.

46. In the second SSH session, start a live Tshark session.

```
Sw1# diagnostics
Sw1# diag utilities tshark
```

Inspecting traffic mirrored to the CPU until Ctrl-C is entered.

47. On Agg1, ping 1 time to PC1 (ping 10.x.1.91 repetitions 1).

```
Agg1# ping 10.x.1.91 repetitions 1
PING 10.x.1.91 (10.x.1.91) 100(128) bytes of data.
108 bytes from 10.x.1.91: icmp_seq=1 ttl=128 time=0.628 ms

--- 10.x.1.91 ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 0.628/0.628/0.628/0.000 ms
```

48. This ICMP echo request and reply should be shown in the live Tshark session.

```
----- Output Omitted -----
Source: 10.x.1.2
Destination: 10.x.1.91
Internet Control Message Protocol
Type: 8 (Echo (ping) request)
----- Output Omitted -----
Source: 10.x.1.91
Destination: 10.x.1.2
Internet Control Message Protocol
Type: 0 (Echo (ping) reply)
----- Output Omitted -----
```

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.

```
Sw1# conf t
Sw1(config)# mirror session 1
Sw1(config-mirror-1)# disable
Sw1(config-mirror-1)# exit
```

52. On Sw1, remove the policy 'mirror'

```
Sw1(config)# no policy mirror
Sw1(config)# no class ip icmp
Sw1(config)# end
```

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.

```
Sw1# conf t
Sw1(config)# radius-server host 10.252.1.23 key plaintext aruba123 vrf mgmt
Sw1(config)# radius-server tracking user-name test password plaintext test
```

54. Open a second SSH session to switch Sw1

55. Run tcpdump in one session. Port 1812 is used for RADIUS communication.

```
Sw1# diagnostics
Sw1# diag utilities tcpdump host-port 1812 vrf mgmt count 4
```

56. In the other SSH session, enable RADIUS tracking.

```
Sw1(config)# radius-server host 10.252.1.23 vrf mgmt tracking enable
```

57. Use the disable/enable feature of the tracking command to trigger a new RADIUS request if you want to repeat this test

```
Sw1(config)# radius-server host 10.252.1.23 vrf mgmt tracking disable
Sw1(config)# radius-server host 10.252.1.23 vrf mgmt tracking enable
```

58. Review TCPDUMP output

```
tcpdump: listening on any, link-type LINUX_SLL (Linux cooked), capture size 262144 bytes
1 08:15:41.495881 IP (tos 0x0, ttl 64, id 2812, offset 0, flags [DF], proto UDP (17),
length 105)
TX-Sw1.47498 > 10.252.1.23.radius: RADIUS, length: 77
Access-Request (1), id: 0x00, Authenticator: de887a6028218f5911f71718f9b4c613
User-Name Attribute (1), length: 6, Value: test
User-Password Attribute (2), length: 18, Value:
Message-Authenticator Attribute (80), length: 18, Value: .,8.+..d._..Qs...
NAS-Identifier Attribute (32), length: 9, Value: TX-Sw1
NAS-IP-Address Attribute (4), length: 6, Value: TX-Sw1
2 08:15:45.602964 IP (tos 0x0, ttl 64, id 3832, offset 0, flags [DF], proto UDP (17),
length 105)
TX-Sw1.47498 > 10.252.1.23.radius: RADIUS, length: 77
Access-Request (1), id: 0x00, Authenticator: de887a6028218f5911f71718f9b4c613
User-Name Attribute (1), length: 6, Value: test
User-Password Attribute (2), length: 18, Value:
Message-Authenticator Attribute (80), length: 18, Value: .,8.+..d._..Qs...
NAS-Identifier Attribute (32), length: 9, Value: TX-Sw1
NAS-IP-Address Attribute (4), length: 6, Value: TX-Sw1
3 08:15:45.604243 IP (tos 0x0, ttl 62, id 61279, offset 0, flags [DF], proto UDP (17),
length 48)
```

```
10.252.1.23.radius > TX-Sw1.47498: RADIUS, length: 20
    Access-Reject (3), id: 0x00, Authenticator: b2f00eaf2453790da6af4da33feebcb9
^C
3 packets captured
3 packets received by filter
0 packets dropped by kernelStop the TCPDUMP using CTRL-C
```

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

59. Disable the RADIUS tracking.

```
Sw1(config)# radius-server host 10.252.1.23 tracking disable
```

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

```
Agg1# show https-server
HTTPS Server Configuration
-----
VRF                : mgmt
REST Access Mode   : read-write
```

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




- 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.



- 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.

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



Time	User	Category	Subcategory	Message
09/30/20 10:...	system	PLAN	PLAN_CREATE	Plan 'Device Import' added 1 devices for hardware addition. Status is now DEPLOY_SUCCESS.
09/30/20 10:...	system	PLAN	DEVICE_ADD	Device '10.251.12.6' added to plan 'Device Import' for hardware addition.
09/30/20 10:...	system	PLAN	FW_CHANG...	Plan 'Initial firmware' detected firmware change for device '10.251.12.6'. New version is 'DeviceMaintenanceService.DeviceFirmwar...
09/30/20 10:...	system	PLAN	FW_CHANG...	Device '10.251.12.6' detected firmware change for plan 'Initial firmware'. New version is 'DeviceMaintenanceService.DeviceFirmwar...
09/30/20 10:...	system	DEVICE	DEVICE_ADD	IP '115.1.1.0.6' is invalid for discovery.
09/30/20 10:...	admin	DEVICE	DEVICE_ADD	Discovery of device '10.251.12.6' is complete.
09/30/20 10:...	system	DEVICE	DEVICE_ADD	Discovery classification failed for device '10.12.1.3' since the device is unreachable.
09/30/20 10:...	system	DEVICE	DEVICE_ADD	Discovery classification failed for device '10.12.1.4' since the device is unreachable.
09/30/20 10:...	system	DEVICE	DEVICE_ADD	Discovery of device '10.12.1.3' has begun.
09/30/20 10:...	system	DEVICE	DEVICE_ADD	Device '10.12.1.3' added.
09/30/20 10:...	system	DEVICE	DEVICE_ADD	Discovery of device '10.12.1.4' has begun.
09/30/20 10:...	system	DEVICE	DEVICE_ADD	Device '10.12.1.4' added.
09/30/20 10:...	system	PLAN	CONFIG_CH...	Plan 'Startup-config change' detected startup configuration change for device '10.251.12.2'. Status is now DEPLOY_SUCCESS.
09/30/20 10:...	system	PLAN	PLAN_CREATE	Plan 'Initial config' imported devices: 10.251.12.2. Status is now DEPLOY_SUCCESS.

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.

- On Agg1, review the LLDP neighbor details for port 1/1/46. Note the current Neighbor Management-address.

```
Agg1# show lldp neighbor-info 1/1/46
Port : 1/1/46
Neighbor Entries : 1
Neighbor Entries Deleted : 0
Neighbor Entries Dropped : 0
Neighbor Entries Aged-Out : 0
Neighbor Chassis-Name : TX-Agg2
Neighbor Chassis-Description : Aruba JL635A GL.10.05.0001
Neighbor Chassis-ID : b8:d4:e7:41:68:00
Neighbor Management-Address : 10.x.1.3
Chassis Capabilities Available : Bridge, Router
Chassis Capabilities Enabled : Bridge, Router
Neighbor Port-ID : 1/1/46
Neighbor Port-Desc : 1/1/46
Neighbor Port VLAN ID : 1
TTL : 120
```

- On Agg2, adjust this by manually assigning the OOBM IP address as the IPv4 management IP in LLDP.

```
Agg2# conf t
Agg2(config)# lldp management-ipv4 10.251.x.3
```

10. On Agg1, apply the LLDP OOBM management IP configuration.

```
Agg1(config)# lldp management-ipv4 10.251.x.2
```

11. On Agg1, confirm the LLDP neighbor details have been updated with the correct management IP address.

```
Agg1# show lldp neighbor-info 1/1/46

Port                               : 1/1/46
Neighbor Entries                   : 1
Neighbor Entries Deleted           : 0
Neighbor Entries Dropped           : 0
Neighbor Entries Aged-Out          : 0
Neighbor Chassis-Name              : TX-Agg2
Neighbor Chassis-Description       : Aruba JL635A GL.10.05.0001
Neighbor Chassis-ID                : b8:d4:e7:41:68:00
Neighbor Management-Address        : 10.251.x.3
Chassis Capabilities Available     : Bridge, Router
Chassis Capabilities Enabled       : Bridge, Router
Neighbor Port-ID                   : 1/1/46
Neighbor Port-Desc                  : 1/1/46
Neighbor Port VLAN ID              : 1
TTL                                : 120
```

12. In NetEdit, confirm that the Agg2 has now been automatically discovered. It can take some time to complete the discovery.

Name	Address	Managed	Status	NAE	MAC	Serial	Current Firmware	Manufact...	Model	Running Config Modified	Cur
T12-Agg1	10.251.12.2	✓	✓	N	b8d4e7-4...	TW04KM0...	GL.10.05.0001	Aruba	8325	09/30/20 10:09:43	1
T12-Agg2	10.251.12.3	✓	✓	N	b8d4e7-4...	TW04KM0...	GL.10.05.0001	Aruba	8325	09/30/20 10:28:00	1
icx-T12-m...	10.251.12.6		✓		204c03-b...	CNJJKLB030	8.5.0.5	Aruba	A9004-US		


13. Correct the LLDP Management IP on Sw1 and Sw2.

```
Sw1# conf t
Sw1(config)# lldp management-ipv4 10.251.x.4
```

```
Sw2# conf t
Sw2(config)# lldp management-ipv4 10.251.x.5
```

14. Verify all 4 switches are now listed in the device list. This may take a few minutes.


Diagnostics page



The screenshot shows the Aruba NetEdit interface with the 'Devices' tab selected. A table lists 5 devices with columns for Name, Address, Managed, Status, NAE, MAC, Serial, Current Firmware, Manufact..., Model, Running Config Modified, and Cur. The devices are T12-Agg1, T12-Agg2, T12-Sw1, T12-Sw2, and lcx-T12-m....


Name	Address	Managed	Status	NAE	MAC	Serial	Current Firmware	Manufact...	Model	Running Config Modified	Cur
T12-Agg1	10.251.12.2	✓	✓	N	b8d4e7-4...	TW04KM0...	GL-10.05.0001	Aruba	8325	09/30/20 10:09:43	1
T12-Agg2	10.251.12.3	✓	✓	N	b8d4e7-4...	TW04KM0...	GL-10.05.0001	Aruba	8325	09/30/20 10:28:00	1
T12-Sw1	10.251.12.4	✓	✓	N	883a30-a...	SG9BKN7...	FL-10.05.0001	Aruba	6300	09/30/20 10:38:04	1
T12-Sw2	10.251.12.5	✓	✓	N	883a30-a...	SG9BKN7...	FL-10.05.0001	Aruba	6300	09/30/20 10:38:05	1
lcx-T12-m...	10.251.12.6		✓		204c03-b...	CNJJKLB030	8.5.0.5	Aruba	A9004-US		

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.



The screenshot shows the 'Discover Devices' dialog box in the Aruba NetEdit interface. The dialog has a 'Subnet' section with a dropdown menu showing '10.0.0.0/8' and a 'Credentials' section with a dropdown menu showing 'asts'. Below these sections, there is a 'Seed Addresses' section with a checkbox labeled 'Seed Addresses' and a text input field containing '10.251.12.254'. The dialog also includes a 'PREVIOUS' button and a 'NEXT' button.

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'



Time	User	Category	Subcategory	Message
09/30/20 10:...	admin	PLAN	PLAN_CREATE	Plan 'Configuration change requested by admin at 2020.09.30 10.50.19 UTC' added devices: ['10.251...
09/30/20 10:...	admin	PLAN	PLAN_CREATE	Device '10.251.12.2' added to plan 'Configuration change requested by admin at 2020.09.30 10.50.1...
09/30/20 10:...	admin	PLAN	PLAN_CREATE	Plan 'Configuration change requested by admin at 2020.09.30 10.50.19 UTC' created.
09/30/20 10:...	admin	PLAN	ATTRIBUTE_...	Plan 'Configuration change requested by admin at 2020.09.30 10.50.19 UTC' attribute 'State' created with value 'Draft'.
09/30/20 10:...	admin	PLAN	ATTRIBUTE_...	Plan 'Configuration change requested by admin at 2020.09.30 10.50.19 UTC' attribute 'Approved-By' created with value '...
09/30/20 10:...	admin	PLAN	ATTRIBUTE_...	Plan 'Configuration change requested by admin at 2020.09.30 10.50.19 UTC' attribute 'Change-ID' created with value '...
09/30/20 10:...	admin	DEVICE	DEVICE_ADD	Discovery of device 'null' is complete.
09/30/20 10:...	admin	DEVICE	DEVICE_DEL...	Deleted device 10.251.12.254.
09/30/20 10:...	admin	DEVICE	DEVICE_ADD	Discovery of device 'null' is complete.
09/30/20 10:...	admin	DEVICE	DEVICE_DEL...	Deleted device 10.12.1.3.

21. On Agg1, enable terminal-monitor with severity debug by CLI if using ssh to connect the device.

```
Agg1# terminal-monitor severity debug
```

Terminal-monitor is enabled successfully

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

```
Agg1# debug rest all
```

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.

```
2020-09-30T11:02:25.340773-0400 hpe-restd[3038] <DEBUG> debug|LOG_DEBUG|cookie not found
2020-09-30T11:02:25.353986-0400 hpe-restd[3038] <INFO> Event|4602|LOG_INFO|AMM|-
|Authentication succeeded for user admin in session D1_gSLZwB0ET2Q7sVVB_rA==
2020-09-30T11:02:25.354311-0400 hpe-restd[3038] <INFO> Event|4604|LOG_INFO|AMM|-|Session
started for user admin, session D1_gSLZwB0ET2Q7sVVB_rA==
2020-09-30T11:02:25.354643-0400 hpe-restd[3038] <INFO> debug|LOG_INFO|Authenticated user
admin with privilege level 15
2020-09-30T11:02:25.363318-0400 hpe-restd[3038] <INFO> Event|4607|LOG_INFO|AMM|-
|Authorization succeeded for user admin, for resource RoleDefault, with action POST
2020-09-30T11:02:26.389257-0400 hpe-restd[3038] <DEBUG> debug|LOG_DEBUG|Authorization
succeeded for user admin, for resource RoleDefault, with action GET
```



```
2020-09-30T11:02:27.408441-0400 hpe-restd[3038] <DEBUG> debug|LOG_DEBUG|Authorization
succeeded for user admin, for resource RoleDefault, with action GET
2020-09-30T11:02:28.429883-0400 hpe-restd[3038] <DEBUG> debug|LOG_DEBUG|Authorization
succeeded for user admin, for resource RoleDefault, with action GET
2020-09-30T11:02:28.489068-0400 hpe-restd[3038] <INFO> Event|4607|LOG_INFO|AMM|-
|Authorization succeeded for user admin, for resource RoleDefault, with action DELETE
2020-09-30T11:02:33.492005-0400 hpe-restd[3038] <INFO> Event|4608|LOG_INFO|AMM|-
|Authorization allowed for user admin, for resource SessionMgmt, with action POST
```

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.

```
2020-09-30T11:22:25.012747-0400 hpe-restd[3038] <DEBUG> debug|LOG_DEBUG|cookie not found
2020-09-30T11:22:25.022336-0400 hpe-restd[3038] <INFO> Event|4602|LOG_INFO|AMM|-
|Authentication succeeded for user admin in session 5NmzKXL-fMVZ3AxsR1DyWw==
2020-09-30T11:22:25.022629-0400 hpe-restd[3038] <INFO> Event|4604|LOG_INFO|AMM|-|Session
started for user admin, session 5NmzKXL-fMVZ3AxsR1DyWw==
2020-09-30T11:22:25.022921-0400 hpe-restd[3038] <INFO> debug|LOG_INFO|Authenticated user
admin with privilege level 15
2020-09-30T11:22:25.025096-0400 hpe-restd[3038] <DEBUG> debug|LOG_DEBUG|Authorization
succeeded for user admin, for resource ConfigFile, with action GET
2020-09-30T11:22:25.111468-0400 hpe-restd[3038] <DEBUG> debug|LOG_DEBUG|Authorization
succeeded for user admin, for resource RoleDefault, with action GET
2020-09-30T11:22:25.407680-0400 hpe-restd[3038] <INFO> Event|4607|LOG_INFO|AMM|-
|Authorization succeeded for user admin, for resource ConfigFile, with action PUT
2020-09-30T11:22:25.633923-0400 hpe-restd[3038] <INFO> Event|4613|LOG_INFO|AMM|-|admin has
written a new switch configuration to running-config
2020-09-30T11:22:32.585933-0400 hpe-restd[3038] <INFO> Event|4608|LOG_INFO|AMM|-
|Authorization allowed for user admin, for resource SessionMgmt, with action POST
2020-09-30T11:22:32.586564-0400 hpe-restd[3038] <INFO> Event|4605|LOG_INFO|AMM|-|Session
ended for user admin, session 5NmzKXL-fMVZ3AxsR1DyWw==
```

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

```
Agg1# show debug
```

```
-----
module          sub_module      severity  vlan  port  ip
mac             instance        vrf
column          client
-----
```

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.

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'.

Submit
[View Response](#)

Curl

```
curl -X GET --header 'Accept: application/json' 'https://10.251.12.4/rest/v1/system/interfaces/1%2F1%2F25/lldp_neighbors'
```

Request URL

```
https://10.251.12.4/rest/v1/system/interfaces/1%2F1%2F25/lldp_neighbors
```

Response Body

```
[
  "/rest/v1/system/interfaces/1%2F1%2F25/lldp_neighbors/b8%3Ad4%3Ae7%3A41%3A38%3A00/1%2F1%2F1"
]
```

Response Code

```
200
```

Response Headers

```
{
  "connection": "keep-alive",
  "content-length": "94",
  "content-security-policy": "script-src 'self' 'unsafe-inline'; object-src 'none'; font-src *; media-src 'none'; form-action 'self'",
  "content-type": "application/json; charset=utf-8",
  "date": "Wed, 30 Sep 2020 15:37:05 GMT",
  "etag": "bb7a782d6dda296bd15838d5872442fd",
  "server": "nginx",
  "strict-transport-security": "max-age=31536000; includeSubdomains;",
  "x-content-type-options": "nosniff",
  "x-frame-options": "SAMEORIGIN",
  "x-xss-protection": "1; mode=block"
}
```

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

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

Response Body

```
[
  "/rest/v1/system/interfaces/1%2F1%2F1/lldp_neighbors/00%3A24%3Aa8%3Ad9%3A1f%3A80/12",
  "/rest/v1/system/interfaces/1%2F1%2F25/lldp_neighbors/b8%3Ad4%3Ae7%3A41%3A38%3A00/1%2F1%2F1"
]
```

Response Code

```
200
```

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

```
Agg1# show lldp neighbor-info
```

```
LLDP Neighbor Information
```

```
=====
```

```
Total Neighbor Entries      : 3
Total Neighbor Entries Deleted : 4
Total Neighbor Entries Dropped : 0
Total Neighbor Entries Aged-Out : 4
```

LOCAL-PORT	CHASSIS-ID	PORT-ID	PORT-DESC	TTL	SYS-NAME
1/1/1	00:24:a8:d9:1f:80	12	12	120	P15-PC1-
Fanout					
1/1/25	b8:d4:e7:41:38:00	1/1/1	1/1/1	120	TX-Agg1
mgmt	00:23:89:d9:c4:8b	GigabitEther...	TX-6300-A-00BM	120	P15-00BM-
Fanout					

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

```
Sw1# https-server session close all
```

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
```

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.

```
Sw1# debug nae script
```

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.

Aruba Analytics > Scripts 10.251.12.4(T12-Sw1) admin

+ Create Agent

Script: LLDP_Renaming

Agent Name: lldp

Parameters:

Type	Name	Description	More Info	Value
STRING	change_current	Change existing description	Default: Yes	
STRING	include_lag	Include LAGs	Default: Yes	
STRING	interfaces	Monitored interfaces	Default: all	

☐ Save running config to startup **CREATE** **CANCEL**

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.

```
Sw1# show debug buffer module nae | in rename
2020-09-30:13:02:19.512554|hpe-policyd|LOG_DEBUG|This interface has to be renamed - 1/1/25
2020-09-30:13:02:19.512588|hpe-policyd|LOG_DEBUG|Interface is going to be renamed
2020-09-30:13:02:19.512672|hpe-policyd|LOG_DEBUG|This interface has to be renamed - 1/1/1
2020-09-30:13:02:19.512713|hpe-policyd|LOG_DEBUG|Interface is going to be renamed
2020-09-30:13:02:19.513020|hpe-policyd|LOG_DEBUG|This interface has to be renamed - 1/1/3
2020-09-30:13:02:19.513052|hpe-policyd|LOG_DEBUG|This interface cannot be renamed - Not
Enough Information
```

12. Use 'show interface brief' to verify that the port descriptions have been updated based on the LLDP peer name.

```
Sw1# show interface brief | exclude down
```

Port Description	Native VLAN	Mode	Type	Enabled	Status	Reason	Speed (Mb/s)
1/1/1 to P15-PC1-Fa...	1	access	1GbT	yes	up		100 Link
1/1/3	1	access	1GbT	yes	up		1000 PC3
1/1/25 to TX-Agg1 (...)	1	access	SFP+DAC1	yes	up		10000 Link

13. Disable all debugging.

```
Sw1# no debug all
```

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.

```
Agg1# show interface brief | exclude down
```

Port Description	Native VLAN	Mode	Type	Enabled	Status	Reason	Speed (Mb/s)
1/1/1 to TX-Sw1 (1...	1	access	SFP+DAC1	yes	up		10000 Link
1/1/7	--	routed	SFP-BT	yes	up		1000 --
1/1/46 to TX-Agg2 (...)	1	trunk	SFP28DAC0.65	yes	up		25000 Link
1/1/47 to TX-Agg2 (...)	1	trunk	SFP28DAC0.65	yes	up		25000 Link

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.

```
Sw1# clear nae-data
```

```
This action will delete NAE data monitored so far.  
Do you want to continue (y/n)? y
```

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.

```
Agg1# show snmp vrf
```

```
SNMP enabled VRF
```

```
-----
```

3. Enable snmp-server on Agg1 for mgmt VRF and the default VRF.

```
Agg1# conf t
```

```
Agg1(config)# snmp-server vrf default
```

```
Agg1(config)# snmp-server vrf mgmt
```

4. Verify the active VRFs for SNMP again, both default and mgmt should be active.

```
Agg1# show snmp vrf
```

```
SNMP enabled VRF
```

```
-----
```

```
default
```

```
mgmt
```

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

```
Agg1# show snmp community
```

```
-----  
SNMP communities  
-----  
public
```

6. Attempt to remove the public community

```
Agg1(config)# no snmp-server community public  
The community is not configured
```

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'

```
Agg1(config)# snmp-server community aruba123
```

8. Review the SNMP communities

```
Agg1# show snmp community  
-----  
SNMP communities  
-----  
aruba123
```

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.

```
Agg1(config)# snmpv3 user amp auth sha auth-pass plaintext aruba123 priv aes priv-pass  
plaintext aruba123
```

10. Define an SNMPv3 context 'aruba' and link it to the default VRF and community 'aruba123'

```
Agg1(config)# snmpv3 context aruba vrf default community aruba123
```

11. Link the SNMPv3 account to the SNMPv3 context.

```
Agg1(config)# snmpv3 user amp context aruba
```

Troubleshooting SNMP packets

12. Use `tcpdump` to review the SNMP packets and community strings. Make sure to run `TCPDUMP` in the correct VRF context.

```
Agg1# diagnostics
Agg1# diag utilities tcpdump host-port 161 vrf mgmt count 2
```

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.

```
tcpdump: listening on any, link-type LINUX_SLL (Linux cooked), capture size 262144 bytes
 1 13:36:00.806108 IP (tos 0x0, ttl 128, id 19925, offset 0, flags [none], proto UDP (17), length 75)
    10.251.x.91.50055 > TX-Agg1.snmp: { SNMPv2c C="badcommunity" { GetRequest(26) R=14244 system.sysUpTime.0 } }
 2 13:36:01.811774 IP (tos 0x0, ttl 128, id 19927, offset 0, flags [none], proto UDP (17), length 75)
    10.251.x.91.50055 > TX-Agg1.snmp: { SNMPv2c C="badcommunity" { GetRequest(26) R=14244 system.sysUpTime.0 } }
2 packets captured
2 packets received by filter
0 packets dropped by kernel
```

14. Stop the `TCPDUMP`.

SNMP Events

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

```
Agg1# show event -r -d snmpd_wrapper

-----
Event logs from current boot
-----
2020-09-30T13:25:33.387958-04:00 TX-Agg1 snmpd_wrapper[8369]: Event|7101|LOG_INFO|AMM|-|Snmp
agent is up and running in namespace VRF_1
2020-09-30T13:25:29.381891-04:00 TX-Agg1 snmpd_wrapper[8369]: Event|7101|LOG_INFO|AMM|-|Snmp
agent is up and running in namespace swns
```

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.

```
Agg1# diag utilities snmpget oid system.sysUpTime.0 community aruba123  
DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (98545) 0:16:25.45
```

18. Optional: Save a new checkpoint on Agg1, Agg2, Sw1 and Sw2

```
copy run checkpoint asts-lab03-<your-name>
```

You have completed this Lab!

Aruba Training-Confidential

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

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

```
Agg1# configure terminal
Agg1(config)# vrf KA
Agg1(config-vrf)# interface 1/1/45
Agg1(config-if)# no shutdown
Agg1(config-if)# vrf attach KA
Agg1(config-if)# description VSX keepalive
Agg1(config-if)# ip address 192.168.0.0/31
Agg1(config-if)# exit
Agg1(config)# vsx
Agg1(config-vsx)# role primary
Agg1(config-vsx)# inter-switch-link lag 256
Agg1(config-vsx)# keepalive peer 192.168.0.1 source 192.168.0.0 vrf KA
Agg1(config-vsx)# end
```

3. Configure Agg2 for VSX

```
Agg2# configure terminal
Agg2(config)# vrf KA
Agg2(config-vrf)# interface 1/1/45
Agg2(config-if)# no shutdown
Agg2(config-if)# vrf attach KA
Agg2(config-if)# description VSX keepalive
Agg2(config-if)# ip address 192.168.0.1/31
Agg2(config-if)# exit
Agg2(config)# vsx
Agg2(config-vsx)# role secondary
Agg2(config-vsx)# inter-switch-link lag 256
Agg2(config-vsx)# keepalive peer 192.168.0.0 source 192.168.0.1 vrf KA
```

```
Agg2(config-vsx)# end
```

4. Verify that the VSX is operational on Agg1 switch

```
Agg1# show vsx status
```

VSX Operational State

```
-----
ISL channel           : In-Sync
ISL mgmt channel      : operational
Config Sync Status    : in-sync
NAE                   : peer_reachable
HTTPS Server          : peer_reachable
```

Attribute	Local	Peer
ISL link	lag256	lag256
ISL version	2	2
System MAC	b8:d4:e7:41:38:00	b8:d4:e7:41:68:00
Platform	8325	8325
Software Version	GL.10.05.0021	GL.10.05.0021
Device Role	primary	secondary

5. Review the VLAN status of the ISL port (LAG256)

```
Agg1# show vlan port lag256
```

VLAN	Name	Mode	Mapping
1	DEFAULT_VLAN_1	native-tagged	port
11	VLAN11	trunk	port

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

A: The mode is native-tagged.

6. Review the running configuration of LAG256

```
Agg1# show running interface lag256
```

```
interface lag 256
  no shutdown
  description LAG 256 - Link to TX-Agg2
  no routing
  vlan trunk native 1 tag
  vlan trunk allowed all
  lacp mode active
  exit
```

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

```
Agg1# configure terminal
Agg1(config)# vsx
Agg1(config-vsx)# vsx-sync dhcp-relay lldp mclag-interfaces
Agg1(config-vsx)# vsx-sync qos-global stp-global vsx-global
Agg1(config-vsx)# exit
```

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

```
Agg1(config)# interface vlan 1
Agg1(config-if-vlan)# l3-counters
Agg1(config-if-vlan)# vsx-sync active-gateways policies
Agg1(config-if-vlan)# active-gateway ip mac 02:00:00:00:00:01
Agg1(config-if-vlan)# active-gateway ip 10.x.1.1
Agg1(config-if-vlan)# exit
```

VSX SVI Active Gateway - SVI 12

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

```
Agg1(config)# vlan 12
Agg1(config-vlan-12)# vsx-sync
Agg1(config-vlan-12)# exit
```

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

```
Agg1(config)# interface vlan 12
Agg1(config-if-vlan)# ip address 10.x.12.2/24
Agg1(config-if-vlan)# l3-counters
Agg1(config-if-vlan)# vsx-sync active-gateways policies
Agg1(config-if-vlan)# active-gateway ip mac 02:00:00:00:00:01
Agg1(config-if-vlan)# active-gateway ip 10.x.12.1
```



```
Agg1(config-if-vlan)# exit
```

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

```
Agg1# show run int vlan12
```

```
interface vlan12
  vsx-sync active-gateways policies
  ip address 10.x.12.2/24
  active-gateway ip mac 02:00:00:00:00:01
  active-gateway ip 10.x.12.1
  l3-counters
  exit
```

```
Agg1# show run int vlan12 vsx-peer
```

Unknown interface vlan12.

```
Agg1# show vlan 12
```

VLAN	Name	Status	Reason	Type	Interfaces
12	VLAN12	up	ok	static	lag256

```
Agg1# show vlan 12 vsx-peer
```

VLAN	Name	Status	Reason	Type	Interfaces
12	VLAN12	up	ok	static	lag256

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

```
Agg1# show vsx status config-sync
```

```
Admin state           : Enabled
Operational State     : Operational
Error State           : Missing reference error
Recommended remediation : A. Execute 'show running-config vsx-sync peer-diff' to determine
                           which lines did not sync correctly on secondary
                           B. Identify the configuration that is missing from secondary and
                           manually fix via CLI.
Current time          : Thu Oct 1 13:24:30 2020
Last sync time        : Thu Oct 1 12:36:08 2020
```

13. Review the differences in the running config

```
Agg1# show running-config vsx-sync peer-diff
```

```
--- /tmp/running-config-vsx.2f9f      2020-10-01 13:25:51.755712370 -0400
+++ /tmp/peer-running-config-vsx.2f9f  2020-10-01 13:25:51.746712472 -0400
@@ -1,16 +1,13 @@
!
```

```
!Version ArubaOS-CX GL.10.05.0001
!export-password: default
+lldp management-ipv4-address 10.251.x.3
vlan 12
    vsx-sync
interface vlan 1
    vsx-sync active-gateways policies
    active-gateway ip mac 02:00:00:00:00:01
    active-gateway ip 10.x.1.1
-interface vlan 12
-    vsx-sync active-gateways policies
-    active-gateway ip mac 02:00:00:00:00:01
-    active-gateway ip 10.x.12.1
```

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.

```
Agg2# conf t
Agg2(config)# int vlan 12
Agg2(config-if-vlan)# ip address 10.x.12.3/24
```

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

```
Agg1# show vsx status config-sync
```

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.

```
# ping 10.252.1.21 repetitions 2 source vlan1
PING 10.252.1.21 (10.252.1.21) from 10.x.1.2 : 100(128) bytes of data.
108 bytes from 10.252.1.21: icmp_seq=1 ttl=125 time=3.42 ms
108 bytes from 10.252.1.21: icmp_seq=2 ttl=125 time=2.97 ms

--- 10.252.1.21 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1001ms
rtt min/avg/max/mdev = 2.978/3.199/3.420/0.221 ms
```

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

```
Agg1# conf t
Agg1(config)# interface vlan12
Agg1(config-if-vlan)# ip helper-address 10.252.1.21
Agg1(config-if-vlan)# exit
```

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.

```
Agg1# show interface brief | exclude down
```

Port Description	Native VLAN	Mode	Type	Enabled	Status	Reason	Speed (Mb/s)	
1/1/1 to TX-Sw1 (1...	1	access	SFP+DAC1	yes	up		10000	Link
1/1/7	--	routed	SFP-BT	yes	up		1000	--
1/1/45	--	routed	SFP28DAC0.65	yes	up		25000	
Keepalive Link port								
1/1/46 Inter-Switch L...	1	trunk	SFP28DAC0.65	yes	up		25000	VSX
1/1/47 Inter-Switch L...	1	trunk	SFP28DAC0.65	yes	up		25000	VSX

----- Output Omitted -----

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

```
Agg1# conf t
Agg1(config)# interface lag 5 multi-chassis
Agg1(config-lag-if)# VLAN trunk allowed 1,11,12
Agg1(config-lag-if)# no shutdown
Agg1(config-lag-if)# interface 1/1/5
Agg1(config-if)# lag 5
Agg1(config-if)# no shutdown
Agg1(config-if)# exit
```

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

```
Agg2# conf t
Agg2(config)# interface lag 5 multi-chassis
Agg2(config-lag-if)# show run current
Agg2(config-lag-if)# no shutdown
```

```
Agg2(config-lag-if)# exit
```

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

```
Agg2(config)# interface 1/1/5
Agg2(config-if)# lag 5
Agg2(config-if)# no shutdown
Agg2(config-if)# exit
```

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

```
Agg1# show lacp interfaces
```

```
----- Output Omitted -----
Actor details of all interfaces:
```

Intf	Aggr Name	Port Id	Port Pri	State	System-ID	System Pri	Aggr Key	Forwarding State
1/1/5	lag5(mc)	5	1	ALFNCD	b8:d4:e7:41:38:00	65534	5	up
1/1/46	lag256	47	1	ALFNCD	b8:d4:e7:41:38:00	65534	256	up
1/1/47	lag256	48	1	ALFNCD	b8:d4:e7:41:38:00	65534	256	up

```
----- Output Omitted -----
```

```
Agg1# show lacp interfaces vsx-peer
```

```
----- Output Omitted -----
Actor details of all interfaces:
```

Intf	Aggr Name	Port Id	Port Pri	State	System-ID	System Pri	Aggr Key	Forwarding State
1/1/5	lag5(mc)	1005	1	ALFNCD	b8:d4:e7:41:38:00	65534	5	up
1/1/46	lag256	47	1	ALFNCD	b8:d4:e7:41:68:00	65534	256	up
1/1/47	lag256	48	1	ALFNCD	b8:d4:e7:41:68:00	65534	256	up

```
----- Output Omitted -----
```

```
Agg1# show lacp interfaces multi-chassis
```

```
----- Output Omitted -----
Actor details of all interfaces:
```

Intf	Aggregate name	Port id	Port Priority	State	System-ID	System Priority	Aggr Key
1/1/5	lag5(mc)	5	1	ALFNCD	b8:d4:e7:41:38:00	65534	5

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

Intf	Aggregate name	Partner Port-id	Port Priority	State	System-ID	System Priority	Aggr Key
1/1/5	lag5(mc)	2	255	ASFNCD	20:4c:03:b1:ce:ba	32768	2

----- Output Omitted -----

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

```
Agg1# conf t
Agg1(config)# interface lag 2 multi-chassis
Agg1(config-lag-if)# description Sw2
Agg1(config-lag-if)# vlan trunk allowed 1,11-12
Agg1(config-lag-if)# no shutdown
Agg1(config-lag-if)# interface 1/1/2
Agg1(config-if)# description Sw2
Agg1(config-if)# no shutdown
Agg1(config-if)# lag 2
Agg1(config-if)# end
```

24. Verify the configuration on Agg1.

```
Agg1# show lacp interfaces
----- Output Omitted -----
Actor details of all interfaces:
-----
Intf      Aggr    Port  Port  State  System-ID          System Aggr  Forwarding
         Name    Id    Pri           Pri    Key    State
-----
1/1/2     lag2(mc)
1/1/5     lag5(mc)   5    1    ALFNCD b8:d4:e7:41:38:00 65534  5    up
1/1/46    lag256    47    1    ALFNCD b8:d4:e7:41:38:00 65534 256  up
1/1/47    lag256    48    1    ALFNCD b8:d4:e7:41:38:00 65534 256  up
----- Output Omitted -----
```

25. Review the 'show interface brief' output.

```
Agg1# show interface brief
-----
Port      Native  Mode   Type           Enabled Status  Reason           Speed
Description
          VLAN                                (Mb/s)
-----
```

1/1/1	1	access	SFP+DAC1	yes	up		10000	Link
to TX-Sw1 (1...								
1/1/2	1	trunk	SFP+DAC1	yes	down	Waiting for link	--	Sw2
----- Output Omitted -----								

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.

```
Sw2# conf t
Sw2(config)# vlan 11-12
Sw2(config-vlan-<11-12>)# exit
Sw2(config)# interface lag 255
Sw2(config-lag-if)# vlan trunk allowed 1,11,12
Sw2(config-lag-if)# no shutdown
Sw2(config-lag-if)# interface 1/1/25-1/1/26
Sw2(config-if-<1/1/25-1/1/26>)# lag 255
Sw2(config-if-<1/1/25-1/1/26>)# no shut
Sw2(config-if-<1/1/25-1/1/26>)# end
```

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

```
Agg2# conf t
Agg2(config)# interface lag 2 multi-chassis
Agg2(config-lag-if)# no shutdown
Agg2(config-lag-if)# exit
Agg2(config)# interface 1/1/2
Agg2(config-if)# lag 2
Agg2(config-if)# no shutdown
Agg2(config-if)# exit
```

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

```
Sw2# show lacp interfaces
----- Output Omitted -----
Actor details of all interfaces:
-----
```

Intf	Aggr Name	Port Id	Port Pri	State	System-ID	System Pri	Aggr Key	Forwarding State
1/1/25	lag255							up
1/1/26	lag255							up

```
----- Output Omitted -----
```

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.

```
Agg1# show interface brief
```

Port Description	Native VLAN	Mode	Type	Enabled	Status	Reason	Speed (Mb/s)	
1/1/1 to TX-Sw1 (1...	1	access	SFP+DAC1	yes	up		10000	Link
1/1/2 to TX-Sw2 (1...	1	trunk	SFP+DAC1	yes	up		10000	Link
----- Output Omitted -----								

```
Agg1# show lacp interfaces
```

----- Output Omitted -----

Actor details of all interfaces:

Intf	Aggr Name	Port Id	Port Pri	State	System-ID	System Pri	Aggr Key	Forwarding State
1/1/2	lag2(mc)	2	1	ALFOE	b8:d4:e7:41:38:00	65534	2	lacp-block
1/1/5	lag5(mc)	5	1	ALFNCD	b8:d4:e7:41:38:00	65534	5	up
1/1/46	lag256	47	1	ALFNCD	b8:d4:e7:41:38:00	65534	256	up
1/1/47	lag256	48	1	ALFNCD	b8:d4:e7:41:38:00	65534	256	up
----- Output Omitted -----								

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

```
Agg1# diag-dump lag basic
```

----- Output Omitted -----

LACP PDUs counters:

LAG lag2:

Configured interfaces:

Interface: 1/1/2

```
lacp_pdus_sent: 20
marker_response_pdus_sent: 0
lacp_pdus_received: 0
marker_pdus_received: 0
----- Output Omitted -----
```

Adjust the LAG to LACP

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

```
Sw2# conf t
Sw2(config)# interface lag 255
Sw2(config-lag-if)# lacp mode active
Sw2(config-lag-if)# end
```

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

```
Agg1# show lacp interfaces
----- Output Omitted -----
Actor details of all interfaces:
-----
```

Intf	Aggr Name	Port Id	Port Pri	State	System-ID	System Pri	Aggr Key	Forwarding State
1/1/2	lag2(mc)	2	1	ALFNCD	b8:d4:e7:41:38:00	65534	2	up
1/1/5	lag5(mc)	5	1	ALFNCD	b8:d4:e7:41:38:00	65534	5	up
1/1/46	lag256	47	1	ALFNCD	b8:d4:e7:41:38:00	65534	256	up
1/1/47	lag256	48	1	ALFNCD	b8:d4:e7:41:38:00	65534	256	up

```
----- Output Omitted -----
```

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

```
Agg1# diag-dump lag basic
----- Output Omitted -----
LACP PDUs counters:
LAG lag2:
Configured interfaces:
Interface: 1/1/2
lacp_pdus_sent: 34
marker_response_pdus_sent: 0
lacp_pdus_received: 8
marker_pdus_received: 0
----- Output Omitted -----
```

Client PC Validation

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

```
Sw2# conf t
Sw2(config)# interface 1/1/4
Sw2(config-if)# vlan access 12
Sw2(config-if)# no shut
Sw2(config-if)# end
```


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

```
C:\Users\student>ping 10.252.1.23 -n 1

Pinging 10.252.1.23 with 32 bytes of data:
Reply from 10.252.1.23: bytes=32 time=3ms TTL=60

Ping statistics for 10.252.1.23:
    Packets: Sent = 1, Received = 1, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 3ms, Maximum = 3ms, Average = 3ms

C:\Users\student>ping 10.252.1.24 -n 1

Pinging 10.252.1.24 with 32 bytes of data:
Reply from 10.252.1.24: bytes=32 time=3ms TTL=60

Ping statistics for 10.252.1.24:
    Packets: Sent = 1, Received = 1, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 3ms, Maximum = 3ms, Average = 3ms
```

VSX DHCP Relay Services

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

```
Agg1# show dhcp-relay
----- Output Omitted -----

DHCP Relay Statistics:

Valid Requests Dropped Requests Valid Responses Dropped Responses
-----
2              0              2              0

DHCP Relay Option 82 Statistics:

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

Agg1# show dhcp-relay vsx-peer
----- Output Omitted -----

DHCP Relay Statistics:

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

DHCP Relay Option 82 Statistics:

Valid Requests	Dropped Requests	Valid Responses	Dropped Responses
0	0	0	0

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.

```
# debug dhcprelay all
# terminal-monitor severity debug
```

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

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

39. Review the debug output on Agg1 and Agg2.

```
2020-10-01T16:12:52.410178-0400 hpe-relay[1819] <DEBUG> debug|LOG_DEBUG|AMM|-
|DHCPRELAY|DHCPRELAY|DHCP BOOTREQUEST Packet received on interface vlan12 from client MAC
005056-b19530.
2020-10-01T16:12:52.410397-0400 hpe-relay[1819] <INFO> debug|LOG_INFO|AMM|-
|DHCPRELAY|DHCPRELAY|Packet send to Server(10.252.1.21) successfully via interface vlan12.
Client MAC is 005056-b19530.
2020-10-01T16:12:52.413637-0400 hpe-relay[1819] <DEBUG> debug|LOG_DEBUG|AMM|-
|DHCPRELAY|DHCPRELAY|DHCP BOOTREPLY Packet received on interface 1/1/7 from Server
10.252.1.21 to client MAC 005056-b19530.
2020-10-01T16:12:52.413753-0400 hpe-relay[1819] <INFO> debug|LOG_INFO|AMM|-
|DHCPRELAY|DHCPRELAY|Packet send to client successfully via interface vlan12. Client MAC is
005056-b19530. & Client VRF Table ID:0
2020-10-01T16:12:52.414157-0400 hpe-relay[1819] <DEBUG> debug|LOG_DEBUG|AMM|-
|DHCPRELAY|DHCPRELAY|DHCP BOOTREQUEST Packet received on interface vlan12 from client MAC
005056-b19530.
2020-10-01T16:12:52.414310-0400 hpe-relay[1819] <INFO> debug|LOG_INFO|AMM|-
|DHCPRELAY|DHCPRELAY|Packet send to Server(10.252.1.21) successfully via interface vlan12.
Client MAC is 005056-b19530.
2020-10-01T16:12:52.416803-0400 hpe-relay[1819] <DEBUG> debug|LOG_DEBUG|AMM|-
|DHCPRELAY|DHCPRELAY|DHCP BOOTREPLY Packet received on interface 1/1/7 from Server
10.252.1.21 to client MAC 005056-b19530.
2020-10-01T16:12:52.416925-0400 hpe-relay[1819] <INFO> debug|LOG_INFO|AMM|-
|DHCPRELAY|DHCPRELAY|Packet send to client successfully via interface vlan12. Client MAC is
005056-b19530. & Client VRF Table ID:0
```

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.

```
Agg1# conf t
Agg1(config)# interface vlan12
Agg1(config-if-vlan)# shut
```

Warning: active gateway is configured on interface vlan12.
Disabling the interface may result in traffic loss.

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.

```
Agg1(config-if-vlan)# no shut
Agg1(config-if-vlan)# end
```

45. Disable all debugging on both Agg1 and Agg2.

```
# no debug all
```

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.

```
C:\Users\student> ping 10.252.1.23 -n 1

Pinging 10.252.1.23 with 32 bytes of data:
Reply from 10.252.1.23: bytes=32 time=3ms TTL=60

Ping statistics for 10.252.1.23:
    Packets: Sent = 1, Received = 1, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 3ms, Maximum = 3ms, Average = 3ms

C:\Users\student> ping 10.252.1.24 -n 1

Pinging 10.252.1.24 with 32 bytes of data:
Reply from 10.252.1.24: bytes=32 time=3ms TTL=60

Ping statistics for 10.252.1.24:
    Packets: Sent = 1, Received = 1, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 3ms, Maximum = 3ms, Average = 3ms
```

48. On Agg1, shutdown the Layer3 VLAN12 SVI.

```
Agg1# conf t
Agg1(config)# interface vlan12
Agg1(config-if-vlan)# shut

Warning: active gateway is configured on interface vlan12.
Disabling the interface may result in traffic loss.
```

49. On PC4, attempt both pings again.

```
C:\Users\student> ping 10.252.1.23 -n 1

Pinging 10.252.1.23 with 32 bytes of data:
Request timed out.

Ping statistics for 10.252.1.23:
    Packets: Sent = 1, Received = 0, Lost = 1 (100% loss),

C:\Users\student> ping 10.252.1.24 -n 1

Pinging 10.252.1.24 with 32 bytes of data:
Request timed out.

Ping statistics for 10.252.1.24:
```

Packets: Sent = 1, Received = 0, Lost = 1 (100% loss),

C:\Users\student>

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.

```
Agg1(config-if-vlan)# no shut
Agg1(config-if-vlan)# end
```

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 VSW 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.

```
interface vlan 11
ip address 10.x.11.1/24
active-gateway ip mac 02:00:00:00:00:01
active-gateway ip 10.x.11.1
```

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.

```
Agg1# conf t
Agg1(config)# interface vlan 11
Agg1(config-if-vlan)# ip address 10.x.11.1/24
Agg1(config-if-vlan)# l3-counters
Agg1(config-if-vlan)# active-gateway ip mac 02:00:00:00:00:01
```

```
Agg1(config-if-vlan)# active-gateway ip 10.x.11.1
```

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

```
Agg1(config-if-vlan)# vsx-sync active-gateways policies
```

54. On SVI 11, configure the DHCP Relay

```
Agg1(config-if-vlan)# ip helper-address 10.252.1.21
Agg1(config-if-vlan)# end
```

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

```
Agg2# conf t
Agg2(config)# interface vlan 11
Agg2(config-if-vlan)# ip address 10.x.11.1/24
Active Gateway is using current configured subnet.
```

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:

```
Agg2(config-if-vlan)# no ip address 10.x.11.3/24
Active Gateway is using current configured subnet.
```

- 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.

```
Agg1# show run int vlan11 vsx-peer
interface vlan11
  vsx-sync active-gateways policies
  ip address 10.x.11.1/24
  active-gateway ip mac 02:00:00:00:00:01
  active-gateway ip 10.x.11.1
```

```
ip helper-address 10.252.1.21
exit
```

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.

```
Sw1# conf t
Sw1(config)# vlan 11
Sw1(config-vlan-11)# interface 1/1/27
Sw1(config-if)# vlan trunk allowed 1,11
Sw1(config-if)# no shut
Sw1(config-if)# interface 1/1/3
Sw1(config-if)# vlan access 11
Sw1(config-if)# exit
```

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

```
Sw1# show vlan port 1/1/27
```

VLAN	Name	Mode	Mapping
1	DEFAULT_VLAN_1	native-untagged	port
11	VLAN11	trunk	port

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.

```
Sw2# conf t
Sw2(config)# vlan 11
Sw2(config-vlan-11)# interface 1/1/27
Sw2(config-if)# vlan trunk allowed 1,11
Sw2(config-if)# no shut
Sw2(config-if)# exit
```

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?

Sw1# show interface brief

Port Description	Native VLAN	Mode	Type	Enabled	Status	Reason	Speed (Mb/s)	
1/1/1	1	access	1GbT	yes	up		100	Link
to P15-PC1-Fa...								
1/1/2	1	access	1GbT	no	down	Administratively down	--	--
1/1/3	1	access	1GbT	yes	up		1000	PC3
----- Output Omitted -----								
1/1/25	1	access	SFP+DAC1	yes	up		10000	Link
to TX-Agg1 (...)								
1/1/26	1	access	SFP+DAC1	no	down	Administratively down	--	--
1/1/27	1	trunk	SFP+DAC1	yes	up		10000	Link
to TX-Sw2 (1...								
1/1/28	1	access	SFP+DAC1	no	down	Administratively down	--	--

Is there an LLDP neighbor?

Sw1# show lldp neighbor-info

LLDP Neighbor Information

=====

Total Neighbor Entries : 3
Total Neighbor Entries Deleted : 7
Total Neighbor Entries Dropped : 0
Total Neighbor Entries Aged-Out : 3

LOCAL-PORT	CHASSIS-ID	PORT-ID	PORT-DESC	TTL	SYS-NAME
1/1/1	00:24:a8:d9:1f:80	12	12	120	P15-PC1-
Fanout					
1/1/25	b8:d4:e7:41:38:00	1/1/1	1/1/1	120	TX-Agg1
1/1/27	88:3a:30:a4:1f:40	1/1/27	1/1/27	120	TX-Sw2
mgmt	00:23:89:d9:c4:8b	GigabitEther...	TX-6300-A-00BM	120	P15-00BM-
Fanout					

Does the VLAN contain the correct ports?

Sw1# show vlan 11

VLAN	Name	Status	Reason	Type	Interfaces
11	VLAN11	up	ok	static	1/1/3,1/1/27

Is the VLAN tagged or untagged status correct?

Sw1# show vlan port 1/1/27

VLAN	Name	Mode	Mapping

1	DEFAULT_VLAN_1	native-untagged port
11	VLAN11	trunk port

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)

```
Sw1# show mac-address-table vlan 11
```

Is there any Layer2 protocol that is impacting the forwarding?

```
Sw1# show spanning-tree
```

```
----- Output Omitted -----
1/1/1      Designated      Forwarding 200000      128      P2P      147088
0          0              0
1/1/2      Disabled        Blocking 20000      128      P2P      0
0          0              0
1/1/3      Designated      Forwarding 20000      128      P2P      147086
0          0              0
----- Output Omitted -----
1/1/25     Root              Forwarding 2000      128      P2P      Bound      8561
126541    3                14
1/1/26     Disabled        Blocking 20000      128      P2P      0
0          0              0
1/1/27     Alternate       Blocking 2000      128      P2P      Bound      4
12154     0                9
1/1/28     Disabled        Blocking 20000      128      P2P      4
11973     2                10
```

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.

```
Sw1(config)# int 1/1/25,1/1/26
Sw1(config)# shutdown
```

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.

```
Sw1# show mac-address-table vlan 11
MAC age-time      : 300 seconds
```

Number of MAC addresses : 3

MAC Address	VLAN	Type	Port
b8:d4:e7:41:38:00	11	dynamic	1/1/27
00:50:56:b1:04:25	11	dynamic	1/1/3
02:00:00:00:00:01	11	dynamic	1/1/27

Aruba Training-Confidential

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.

```
Agg1# show vsx status
VSX Operational State
-----
ISL channel           : In-Sync
ISL mgmt channel      : operational
Config Sync Status    : in-sync
NAE                   : peer_reachable
HTTPS Server          : peer_reachable

Attribute             Local                Peer
-----
ISL link               lag256                lag256
ISL version            2                    2
System MAC             b8:d4:e7:41:38:00    b8:d4:e7:41:68:00
Platform               8325                 8325
Software Version       GL.10.05.0001         GL.10.05.0001
Device Role            primary               secondary
```

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.

```
# show lacp interfaces
# repeat delay 5
```

4. On Agg1, apply a user-defined VSX system MAC.

```
Agg1# conf t
Agg1(config)# vsx
Agg1(config-vsx)# system-mac 02:01:00:00:00:00
```

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.

Partner details of all interfaces:

Intf	Aggr Name	Port Id	Port Pri	State	System-ID	System Pri	Aggr Key
1/1/25	lag255	2	1	ALFNCD	02:01:00:00:00:00	65534	2
1/1/26	lag255	1002	1	ALFNCD	02:01:00:00:00:00	65534	2

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.

```
# conf t
(config)# int vlan1,vlan11,vlan12
(config-if-vlan)# ip mtu 9198
(config-if-vlan)# end
```

9. Configure all ports **on both Sw1 and Sw2** to support jumbo frames.

```
# conf t
(config)# int 1/1/1-1/1/28
(config-if-<1/1/1-1/1/28>)# mtu 9198
(config-if-<1/1/1-1/1/28>)# end
```

Uplink delay-timer

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

```
Agg1# show vsx status linkup-delay
Configured linkup delay-timer           : 180 seconds
Initial sync status                     : Completed
Delay timer status                      : Completed
Linkup Delay time left                  :
Interfaces that will be brought up after delay timer expires :
Interfaces that are excluded from 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.

```
Agg2# show mac-address-table
Agg2# show mac-address-table vsx-peer
```

12. There is a convenient command that groups these 2 commands:

```
Agg2# show vsx mac-address-table
```

MAC Address	VLAN	Type	LocalPort	PeerPort
88:3a:30:a4:1f:40	1	dynamic	lag2	lag2
00:50:56:b1:04:25	11	dynamic	lag2	lag2
88:3a:30:a4:1f:4f	1	dynamic	lag2	lag2

00:50:56:b1:14:aa	1	dynamic	lag2	lag2
20:4c:03:b1:ce:bd	1	dynamic	lag5	lag5
88:3a:30:a5:52:80	1	dynamic	lag2	lag2
00:50:56:b1:95:30	12	dynamic	lag2	lag2
88:3a:30:a4:1f:50	1	dynamic	lag2	lag2
20:4c:03:b1:ce:bc	1	dynamic	lag5	lag5
b8:d4:e7:41:38:00	12	dynamic	lag256	
b8:d4:e7:41:38:00	1	dynamic	lag256	
b8:d4:e7:41:38:00	11	dynamic	lag256	
b8:d4:e7:41:68:00	12	dynamic		lag256
b8:d4:e7:41:68:00	11	dynamic		lag256
b8:d4:e7:41:68:00	1	dynamic		lag256

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.

```
Agg2# show vsx mac-address-table
Agg2# show vsx status
Agg2# show vsx status linkup-delay
Agg2# repeat id 1,2,3 delay 5
```

14. On Agg1, disable the LAG 256 of the ISL.

```
Agg1# conf t
Agg1(config)# interface lag 256
Agg1(config-lag-if)# shutdown
```

15. On Agg1, enable the LAG 256 again.

```
Agg1(config-lag-if)# no shutdown
```

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).

```
Agg1(config-lag-if)# shutdown
Agg1(config-lag-if)# 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).

```
Agg2# show interface brief
Agg2# repeat

*****
Iteration : 11 Command : show interface brief
*****
1/1/1    --      routed SFP+DAC1    no      down    Administratively down  --    --
1/1/2    1      trunk  SFP+DAC1    yes     down    Disabled by VSX       --    Link
to TX-Sw2 (1...
1/1/5    1      trunk  SFP-BT      yes     down    Disabled by VSX       --    Link
to icx-TX-mc...
1/1/7    --      routed SFP-BT      no      down    Administratively down  --    --
----- Output Omitted -----
1/1/45    --      routed SFP28DAC0.65  yes     up       25000
Keepalive Link port
1/1/46    1      trunk  SFP28DAC0.65  yes     up       25000  VSX
Inter-Switch L...
1/1/47    1      trunk  SFP28DAC0.65  yes     up       25000  VSX
Inter-Switch L...
----- Output Omitted -----
vlan1     --      --                yes     down    Disabled by VSX       --    --
vlan11    --      --                yes     down    Disabled by VSX       --    --
vlan12    --      --                yes     down    Disabled by VSX       --    --
lag2      1      trunk  --          yes     down    --                   auto
lag5      1      trunk  --          yes     down    --                   auto
lag256    1      trunk  --          yes     up      --                   25000
```

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.

```
Agg2# conf t
Agg2(config)# vsx
Agg2(config-vsx)# linkup-delay-timer exclude lag 5
Agg2(config-vsx)# end
```

21. Verify the configuration on Agg2.

```
Agg2# show vsx status linkup-delay
Configured linkup delay-timer           : 180 seconds
Initial sync status                     : Completed
Delay timer status                      : Completed
Linkup Delay time left                  :
Interfaces that will be brought up after delay timer expires :
Interfaces that are excluded from delay timer : lag5
```

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

```
Agg2(config-lag-if)# shutdown
Agg2(config-lag-if)# no shutdown
```

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

```
Agg2# show interface brief
```

Port	Native	Mode	Type	Enabled	Status	Reason	Speed
Description	VLAN						(Mb/s)
1/1/1	--	routed	SFP+DAC1	no	down	Administratively down	--
1/1/2	1	trunk	SFP+DAC1	yes	down	Disabled by VSX	--
to TX-Sw2 (1...							
1/1/3	--	routed	--	no	down	No XCVR installed	--
1/1/4	--	routed	--	no	down	No XCVR installed	--
1/1/5	1	trunk	SFP-BT	yes	up		1000
to icx-TX-mc...							


```

----- Output Omitted -----
1/1/45  --      routed SFP28DAC0.65  yes    up                25000
Keepalive Link port
1/1/46  1      trunk  SFP28DAC0.65  yes    up                25000  VSX
Inter-Switch L...
1/1/47  1      trunk  SFP28DAC0.65  yes    up                25000  VSX
Inter-Switch L...
----- Output Omitted -----
vlan1   --      --              yes    down    Disabled by VSX    --      --
vlan11  --      --              yes    down    Disabled by VSX    --      --
vlan12  --      --              yes    down    Disabled by VSX    --      --
lag2    1      trunk  --      yes    down    --              auto
lag5    1      trunk  --      yes    up      --              1000
lag256  1      trunk  --      yes    up      --              50000

```

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.

```

Agg1(config-lag-if)# interface lag 2
Agg1(config-lag-if)# no vlan trunk allowed 1

```

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

```

Agg1(config-lag-if)# interface lag 256

```

```
Agg1(config-lag-if)# shutdown
Agg1(config-lag-if)# no shutdown
```

26. On Agg2, review the status.

```
Agg2# show interface brief
```

----- Output Omitted -----

vlan1	--	--	yes	up		--	--
vlan11	--	--	yes	down	Disabled by VSX	--	--
vlan12	--	--	yes	down	Disabled by VSX	--	--
lag2	1	trunk	--	yes	down	--	auto
lag5	1	trunk	--	yes	blocked	--	auto
lag256	1	trunk	--	yes	up	--	50000

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

```
Agg1(config-lag-if)# interface lag 2
Agg1(config-lag-if)# vlan trunk allowed 1
Agg1(config-lag-if)# end
```

VSX LAG Hashing algorithm

Review the VSX LAG hashing algorithms

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

```
Agg1# show lacp aggregates lag2
Aggregate name   : lag2 (multi-chassis)
Interfaces       : 1/1/2
Peer interfaces  : 1/1/2
Heartbeat rate   : Slow
Hash             : l3-src-dst
Aggregate mode   : Active
```

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.

```
Agg1# conf t
Agg1(config)# interface lag 2
Agg1(config-lag-if)# hash ?

 12-src-dst  Base the hash on 12-src-dst (Default:13-src-dst)
 13-src-dst  Base the hash on 13-src-dst (Default)
 14-src-dst  Base the hash on 14-src-dst (Default:13-src-dst)
```

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.**

```
# show spanning-tree
Spanning-tree is disabled
```

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

```
Agg1# conf t
Agg1(config)# spanning-tree priority 1
```

3. Configure STP priority 1 on Agg2.

```
Agg2# conf t
Agg2(config)# spanning-tree priority 1
```

4. On both Agg1 and Agg2, enable STP.

```
Agg1(config)# spanning-tree
Agg1(config)# end
```

```
Agg2(config)# spanning-tree
Agg2(config)# end
```

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.

```
Agg1# show spanning-tree
Spanning tree status      : Enabled Protocol: MSTP
```

```
MST0
  Root ID    Priority    : 4096
             MAC-Address: 02:01:00:00:00:00
             This bridge is the root
----- Output Omitted -----
```

```
Agg2# show spanning-tree

Spanning tree status      : Enabled Protocol: MSTP

MST0
  Root ID    Priority    : 4096
             MAC-Address: 02:01:00:00:00:00
             This bridge is the root
----- Output Omitted -----
```

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 BPDUs-Tx statistics for LAG2.

```
Agg1# show spanning-tree

----- Output Omitted -----
```

Port	Role	State	Cost	Priority	Type	BPDUs-Tx
BPDUs-Rx	TCN-Tx	TCN-Rx				
1/1/1	Disabled	Blocking	20000	128	P2P	0
0	0	0				
lag2	Designated	Forwarding	2000	64	P2P	179
2	2	2				
lag5	Designated	Forwarding	20000	64	P2P	178
0	0	0				
lag256	Designated	Forwarding	1	64	P2P	179
177	2	0				

```
Number of topology changes      : 2
Last topology change occurred  : 355 seconds ago
```

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

```
Agg1# show spanning-tree vsx
```

----- Output Omitted -----							
Port	Role	State	Cost	Priority	Type	BPDU-Tx	
BPDU-Rx	TCN-Tx	TCN-Rx					
lag2	Designated	Forwarding	2000	64	P2P	0	
0	0	0					
lag5	Designated	Forwarding	20000	64	P2P	0	
0	0	0					
lag256	Designated	Forwarding	1	64	P2P	229	
228	0	2					
Number of topology changes : 2							
Last topology change occurred : 451 seconds ago							

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

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

```
Agg1# conf t
Agg1(config)# int lag 2
Agg1(config-lag-if)# spanning-tree root-guard
Agg1(config-lag-if)# end
```

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

- Configure Sw2 with STP priority 0

```
Sw2# conf t
Sw2(config)# spanning-tree priority 0
```

- On Agg1, verify the local STP status and the VSX peer STP status.

```
Agg1# show spanning-tree
Spanning tree status      : Enabled Protocol: MSTP

MST0
  Root ID    Priority    : 4096
             MAC-Address: 02:01:00:00:00:00
             This bridge is the root
             Hello time(in seconds):2  Max Age(in seconds):20
```

Forward Delay(in seconds):15

Bridge ID Priority : 4096
MAC-Address: 02:01:00:00:00:00
Hello time(in seconds):2 Max Age(in seconds):20
Forward Delay(in seconds):15

Port BPDU-Rx	Role TCN-Tx	State TCN-Rx	Cost	Priority	Type	BPDU-Tx
1/1/1 0	Disabled 0	Blocking 0	20000	128	P2P	0
lag2	Alternate	Root-Inc	2000	64	P2P Bound	352
47	4	8				
lag5	Designated	Forwarding	20000	64	P2P	386
0	0	0				
lag256	Designated	Forwarding	1	64	P2P	387
385	4	0				

Number of topology changes : 4
Last topology change occurred : 24 seconds ago

Agg1# show spanning-tree vsx-peer

Spanning tree status : Enabled Protocol: MSTP

MST0

Root ID Priority : 4096
MAC-Address: 02:01:00:00:00:00
This bridge is the root
Hello time(in seconds):2 Max Age(in seconds):20
Forward Delay(in seconds):15

Bridge ID Priority : 4096
MAC-Address: 02:01:00:00:00:00
Hello time(in seconds):2 Max Age(in seconds):20
Forward Delay(in seconds):15

Port BPDU-Rx	Role TCN-Tx	State TCN-Rx	Cost	Priority	Type	BPDU-Tx
lag2	Alternate	Root-Inc	2000	64	P2P Bound	0
0	0	0				
lag5	Designated	Forwarding	20000	64	P2P	0
0	0	0				
lag256	Designated	Forwarding	1	64	P2P	419
419	0	4				

Number of topology changes : 4

Last topology change occurred : 85 seconds ago

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.

```
Agg1# show span inconsistent-ports
```

Instance ID	Blocked Port	Reason
0	lag2	Root Guard

Q: What is the reason?

A: Root Guard.

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

```
Agg1# show events -r -n 20 -d hpe-mstpd
```

```
2020-10-02T05:00:08.111434-04:00 TX-Agg1 hpe-mstpd[1847]:
Event|2011|LOG_INFO|AMM|1/1|Topology Change received on port lag2 for CIST from source:
88:3a:30:a4:1f:4f
2020-10-02T05:00:07.056544-04:00 TX-Agg1 hpe-mstpd[1847]:
Event|2011|LOG_INFO|AMM|1/1|Topology Change received on port lag2 for CIST from source:
88:3a:30:a4:1f:4f
2020-10-02T05:00:05.216785-04:00 TX-Agg1 hpe-mstpd[1847]: Event|2014|LOG_INFO|AMM|1/1|Port
lag2 blocked on CIST
```

Restore the configuration

13. On Sw2, set the default STP priority.

```
Sw2(config)# no span priority
```

14. On Sw2, verify you can now ping 10.x.1.2 and 10.x.1.3 again.

```
Sw2# ping 10.x.1.2 repetitions 1
```

```
PING 10.x.1.2 (10.x.1.2) 100(128) bytes of data.
108 bytes from 10.x.1.2: icmp_seq=1 ttl=64 time=0.166 ms

--- 10.x.1.2 ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 0.166/0.166/0.166/0.000 ms

# ping 10.x.1.3 repetitions 1
```



```
PING 10.x.1.3 (10.x.1.3) 100(128) bytes of data.
108 bytes from 10.x.1.3: icmp_seq=1 ttl=64 time=0.136 ms

--- 10.x.1.3 ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 0.136/0.136/0.136/0.000 ms
```

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.

```
Agg1# show vsx config-consistency
Agg1# repeat
```

16. On Agg2, change the STP mode rpvst.

```
Agg2# conf t
Agg2(config)# span 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

```
Sw2# show spanning-tree

Spanning tree status      : Enabled Protocol: MSTP

MST0
  Root ID    Priority    : 4096
             MAC-Address: 02:01:00:00:00:00
             Hello time(in seconds):2  Max Age(in seconds):20
             Forward Delay(in seconds):15

  Bridge ID  Priority    : 32768
             MAC-Address: 88:3a:30:a4:1f:40
             Hello time(in seconds):2  Max Age(in seconds):20
```

Forward Delay(in seconds):15

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

```
Sw2# conf t
Sw2(config)# interface lag 255
Sw2(config-lag-if)# spanning-tree loop-guard
Sw2(config-lag-if)# end
```

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.

```
Sw2# terminal-monitor severity info
Terminal-monitor is enabled successfully
```

22. On Agg1, enter the LAG2 context and apply the BPDU-filter.

```
Agg1# conf t
Agg1(config)# interface lag 2
Agg1(config-lag-if)# spanning-tree bpdu-filter
Disable routing on the interface to configure spanning-tree options.
```

23. On Sw2, review the events that were reported.

```
2020-10-02T05:32:09.111357-0400 hpe-mstpd[3493] <INFO> Event|2008|LOG_INFO|MSTR|1|CIST
starved for a BPDU Rx on port lag255 from 4096:020100-000000
2020-10-02T05:32:09.111593-0400 hpe-mstpd[3493] <INFO> Event|2009|LOG_INFO|MSTR|1|BPDU loss-
port lag255 moved to inconsistent state for CIST
2020-10-02T05:32:09.111987-0400 hpe-mstpd[3493] <INFO> Event|2006|LOG_INFO|MSTR|1|CST - Root
changed from 4096: 02:01:00:00:00:00 to 32768: 88:3a:30:a4:1f:40
2020-10-02T05:32:09.114197-0400 hpe-mstpd[3493] <INFO> Event|2014|LOG_INFO|MSTR|1|Port lag255
blocked on CIST
```

24. On Sw2, review the spanning-tree state of the port LAG255.

```
Sw2# show spanning-tree
```

```
----- Output Omitted -----
```

Port	Role	Loop-Inc	Cost	Priority	Port	Cost
lag255	Alternate	2000	64	P2P	24120	
1461	22	10				

```
Number of topology changes : 36
```

```
Last topology change occurred : 891 seconds ago
```

25. On Sw2, review the STP inconsistent ports

```
Sw2# show spanning-tree inconsistent-ports
```

Instance ID	Blocked Port	Reason
0	lag255	Loop Guard

26. After you have confirmed the feature, restore the configuration by removing the BPDU-guard on the Agg1 LAG2 port.

```
Agg1(config-lag-if)# interface lag 2
```

```
Agg1(config-lag-if)# no span bpdu-filter
```

```
Agg1(config-lag-if)# end
```

27. On Sw2, disable terminal monitor and verify that port lag255 is now forwarding again.

```
Sw2# no terminal-monitor
```

```
Sw2# show spanning-tree inconsistent-ports
```

Instance ID	Blocked Port	Reason

```
Sw2# show spanning-tree | include lag
```

Port	Role	Forwarding	Cost	Priority	Port	Cost
lag255	Root	2000	64	P2P	Bound	24122
1477	24	10				

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

```
Sw2# show span | i 1/1/4
Sw2# repeat
```

29. In the first SSH connection, bounce the port 1/1/4 (shutdown > no shutdown).

```
Sw2# conf t
Sw2(config)# interface 1/1/4
Sw2(config-if)# shut
Sw2(config-if)# no shut
```

30. Follow the state in the second SSH connection

```
Iteration : 17 Command : show span | i 1/1/4
*****
1/1/4      Designated      Forwarding 20000      128      P2P      167575
0          0              0
*****
Iteration : 18 Command : show span | i 1/1/4
*****
1/1/4      Disabled        Blocking  20000      128      P2P      167575
0          0              0
*****
Iteration : 19 Command : show span | i 1/1/4
*****
1/1/4      Disabled        Blocking  20000      128      P2P      167575
0          0              0
*****
Iteration : 20 Command : show span | i 1/1/4
*****
1/1/4      Designated      Blocking  20000      128      P2P      167576
0          0              0
*****
Iteration : 21 Command : show span | i 1/1/4
*****
1/1/4      Designated      Forwarding 20000      128      P2P      167576
0          0              0
```

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.

```
Agg1# debug mstp event
Agg1# terminal-monitor severity debug
```

Verify TCN generation for access ports

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

```
Sw2(config)# interface 1/1/4
Sw2(config-if)# shut
Sw2(config-if)# no shut
```

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.

```
Sw2(config-if)# int 1/1/27
Sw2(config-if)# shut
Sw2(config-if)# no shut
Sw2(config-if)# end
```

35. On Agg1, review the messages.

```
2020-10-02T05:46:24.119108-0400 hpe-mstpd[1847] <INFO> Event|2011|LOG_INFO|AMM|1/1|Topology
Change received on port lag2 for CIST from source: 88:3a:30:a4:1f:4f
2020-10-02T05:46:24.119295-0400 hpe-mstpd[1847] <DEBUG> debug|LOG_DEBUG|AMM|-
|MSTP|MSTP_EVENT|CIST -Topology Change received on port lag2 from Source MAC 883a30-a41f40
2020-10-02T05:46:26.112328-0400 hpe-mstpd[1847] <INFO> Event|2011|LOG_INFO|AMM|1/1|Topology
Change received on port lag2 for CIST from source: 88:3a:30:a4:1f:4f
2020-10-02T05:46:26.112513-0400 hpe-mstpd[1847] <DEBUG> debug|LOG_DEBUG|AMM|-
|MSTP|MSTP_EVENT|CIST -Topology Change received on port lag2 from Source MAC 883a30-a41f40
```

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.

```
Agg1# show spanning-tree
----- Output Omitted -----
Number of topology changes      : 9
Last topology change occurred   : 71 seconds ago
```

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.

```
Agg1# no debug all
Agg1# no terminal-monitor
```

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.

```
Sw2# terminal-monitor severity info
```

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

```
Sw2# conf t
Sw2(config)# span bpdu-guard timeout 30
```

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.

```
Sw2(config)# int 1/1/4,1/1/27,1/1/28
Sw2(config-if-<1/1/4,1/1/27,1/1/28>)# spanning-tree bpdu-guard
```

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

```
Sw2(config-if-<1/1/4,1/1/27,1/1/28>)# interface 1/1/27
Sw2(config-if)# shut
Sw2(config-if)# no shut
```

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.

```
2020-10-02T05:53:04.294765-0400 intf[d]756] <INFO> Event|404|LOG_INFO|||Link status for
interface 1/1/27 is down
2020-10-02T05:53:08.084503-0400 lldpd[3325] <INFO> Event|110|LOG_INFO|MSTR|1|Configured LLDP
reinit-delay to 2
2020-10-02T05:53:08.108870-0400 intf[d]756] <INFO> Event|403|LOG_INFO|||Link status for
interface 1/1/27 is up
2020-10-02T05:53:08.133465-0400 hpe-mstpd[3493] <INFO> Event|2012|LOG_INFO|MSTR|1|CIST -
Topology Change generated on port 1/1/27 going in to forwarding
2020-10-02T05:53:08.190601-0400 hpe-mstpd[3493] <WARN> Event|2007|LOG_WARN|MSTR|1|Port 1/1/27
disabled - BPDU received on protected port
2020-10-02T05:53:08.190795-0400 hpe-mstpd[3493] <ERR>
debug|LOG_ERR|MSTR|1|MSTP|MSTP_PACKET|port 1/1/27 disabled - BPDU received on protected port.
2020-10-02T05:53:08.217302-0400 intf[d]756] <INFO> Event|404|LOG_INFO|||Link status for
interface 1/1/27 is down
2020-10-02T05:53:40.303972-0400 lldpd[3325] <INFO> Event|110|LOG_INFO|MSTR|1|Configured LLDP
reinit-delay to 2
2020-10-02T05:53:40.333604-0400 intf[d]756] <INFO> Event|403|LOG_INFO|||Link status for
interface 1/1/27 is up
2020-10-02T05:53:40.358953-0400 hpe-mstpd[3493] <INFO> Event|2012|LOG_INFO|MSTR|1|CIST -
Topology Change generated on port 1/1/27 going in to forwarding
2020-10-02T05:53:40.414038-0400 hpe-mstpd[3493] <WARN> Event|2007|LOG_WARN|MSTR|1|Port 1/1/27
disabled - BPDU received on protected port
2020-10-02T05:53:40.414252-0400 hpe-mstpd[3493] <ERR>
debug|LOG_ERR|MSTR|1|MSTP|MSTP_PACKET|port 1/1/27 disabled - BPDU received on protected port.
2020-10-02T05:53:40.446001-0400 intf[d]756] <INFO> Event|404|LOG_INFO|||Link status for
interface 1/1/27 is down
```

44. On Sw2, review the spanning tree port states

```
Sw2# show spanning-tree
1/1/27      Disabled      Bpdu-Error 20000      128      P2P      33026
54          28          11

Sw2# show spanning-tree inconsistent
Instance ID  Blocked Port  Reason
-----
0            1/1/27      BPDU Guard
```

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

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.

```
Sw1(config)# interface 1/1/1,1/1/3
Sw1(config-if-<1/1/1,1/1/3>)# shutdown
Sw1(config-if-<1/1/1,1/1/3>)# exit
```

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

```
Sw1# show interface brief
```

Port	Native VLAN	Mode	Type	Enabled	Status	Reason	Speed (Mb/s)
1/1/1	1	access	1GbT	no	down	Administratively down	--
1/1/2	1	access	1GbT	no	down	Administratively down	--
1/1/3	11	access	1GbT	no	down	Administratively down	--

----- Output Omitted -----

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.

```
Sw1(config)# interface 1/1/27,1/1/28
Sw1(config-if-<1/1/27,1/1/28>)# spanning-tree bpdu-filter
```


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.

```
Sw1(config-if)# int 1/1/28
Sw1(config-if)# no shut
Sw1(config-if)# end
```

5. On Sw2, review the current (default) loop protect interface configuration and status. Loop protect is not enabled by default.

```
Sw2# show loop-protect
Status and Counters - Loop Protection Information

Transmit Interval      : 5 (sec)
Port Re-enable Timer   : 30 (sec)
Loop Detected Trap     : Disabled

Sw2# show loop-protect loop-detected
Status and Counters - Loop Protection Information

Transmit Interval      : 5 (sec)
Port Re-enable Timer   : 30 (sec)
Loop Detected Trap     : Disabled
```

Enable the loop

6. On Sw2, enable port 1/1/28, wait about 10 seconds.

```
Sw2# conf t
Sw2(config)# interface 1/1/28
Sw2(config-if)# no shut
Sw2(config-if)# end
```

7. On Sw2, verify the port status of the ports 1/1/27 and 1/1/28, both should be 'up'

```
Sw2# show interface brief
----- Output Omitted -----
1/1/25  1      trunk  SFP+DAC1    yes    up      10000  --
1/1/26  1      trunk  SFP+DAC1    yes    up      10000  --
1/1/27  1      trunk  SFP+DAC1    yes    up      10000  --
1/1/28  1      access SFP+DAC1    yes    up      10000  --
vlan1   --      --      --          yes    up      --      --
lag255  1      trunk  --          yes    up      20000  --
```

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.

```
Sw2# clear interface statistics global

Warning: clearing statistics globally will be reflected in all CLI
sessions, any agents running in the analytics engine, and any external
systems monitoring switch statistics.
```

Continue (y/n)? y

Sw2# show interface statistics non-zero

Interface		RX Bytes		RX Packets		RX Drops	
TX Bytes	TX Packets		TX Drops	RX Broadcast		RX Multicast	
TX Broadcast	TX Multicast						
1/1/4		0		0		0	
1670	15	0		0		0	
3	12						
1/1/25	- lag255	2801		26		0	
15109528191	236086375		0		6		20
236086373	2						
1/1/26	- lag255	1065		10		0	
9305482943	145398169		0		2		8
145398167	2						
1/1/27		12095262271		188988472		0	
12098026832	z 189031658		0		188988471		1
189031646	12						
1/1/28		12186036607		190406821		0	
12183264264	190363493		0		190406821		1
190363481	12						

Loop Protection configuration

9. On Sw1, configure the global loop-protect re-enable timer to 30 seconds.

```
Sw1# conf t
Sw1(config)# loop-protect re-enable-timer 30
```

10. On Sw1, enable loop-protect on 1/1/27 and 1/1/28.

```
Sw1(config)# interface 1/1/27,1/1/28
Sw1(config-if-1/1/27,1/1/28)# loop-protect
```

11. Observe how the ports are disabled by loop-protect.

```
Sw1# show loop-protect

Status and Counters - Loop Protection Information

Transmit Interval      : 5 (sec)
Port Re-enable Timer   : 30 (sec)
Loop Detected Trap     : Disabled

Interface 1/1/27
Loop-protect enabled   : Yes
Action on loop detection : TX disable
Loop detected count    : 2
Loop detected          : Yes
Detected on VLAN       : 1
Detected at            : 2020-10-02T06:39:03
```

```
Interface status          : down

Interface 1/1/28
Loop-protect enabled     : Yes
Action on loop detection : TX disable
Loop detected count      : 0
Loop detected            : No
Interface status         : up
```

Sw1# show loop-protect loop-detected

Status and Counters - Loop Protection Information

```
Transmit Interval       : 5 (sec)
Port Re-enable Timer    : 30 (sec)
Loop Detected Trap      : Disabled
```

```
Interface 1/1/27
Loop-protect enabled     : Yes
Action on loop detection : TX disable
Loop detected count      : 3
Loop detected            : Yes
  Detected on VLAN       : 1
  Detected at            : 2020-10-02T06:39:43
Interface status         : down
```

12. Review the last 20 events on Sw1, filter for the hpe-lpd daemon.

show event -r -d hpe-lpd -n 20

```
2020-10-02T06:40:23.604626-04:00 TX-Sw1 hpe-lpd[3333]: Event|2808|LOG_INFO|MSTR|1|Ports TX
1/1/27 and RX 1/1/28 are involved during TX port disabling
2020-10-02T06:40:23.604336-04:00 TX-Sw1 hpe-lpd[3333]: Event|2801|LOG_WARN|MSTR|1|Port 1/1/27
is disabled by Loop-protection after loop detection on VLAN 1
2020-10-02T06:40:14.329706-04:00 TX-Sw1 hpe-lpd[3333]: Event|2804|LOG_INFO|MSTR|1|Port 1/1/27
enabled after disable time expired
2020-10-02T06:39:43.604460-04:00 TX-Sw1 hpe-lpd[3333]: Event|2808|LOG_INFO|MSTR|1|Ports TX
1/1/27 and RX 1/1/28 are involved during TX port disabling
2020-10-02T06:39:43.604146-04:00 TX-Sw1 hpe-lpd[3333]: Event|2801|LOG_WARN|MSTR|1|Port 1/1/27
is disabled by Loop-protection after loop detection on VLAN 1
2020-10-02T06:39:34.276414-04:00 TX-Sw1 hpe-lpd[3333]: Event|2804|LOG_INFO|MSTR|1|Port 1/1/27
enabled after disable time expired
```

13. Use 'show interface brief', check the 'Reason' description.

Sw1# show interface brief

```
----- Output Omitted -----
1/1/25    1      access SFP+DAC1      no      down      Administratively down  --      Link
to TX-Agg1 (...
1/1/26    1      access SFP+DAC1      no      down      Administratively down  --      --
1/1/27    1      trunk  SFP+DAC1      yes     down      Network loop detected  --      Link
to TX-Sw2 (1...
```

1/1/28	1	access	SFP+DAC1	yes	up		10000	Link
to TX-Sw2 (1...								
vlan1	--	--		yes	up	--	--	--

Restore the configuration

14. On Sw2, disable port 1/1/28

```
Sw2(config)# interface 1/1/28
Sw2(config-if)# shutdown
```

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.

```
Sw1# show interface brief
----- Output Omitted -----
1/1/25    1      access SFP+DAC1      no      down      Administratively down  --      Link
to TX-Agg1 (...
1/1/26    1      access SFP+DAC1      no      down      Administratively down  --      --
1/1/27    1      trunk  SFP+DAC1      yes     up        10000    Link
to TX-Sw2 (1...
1/1/28    1      access SFP+DAC1      yes     down      Waiting for link      --      Link
to TX-Sw2 (1...
vlan1     --      --                yes     up        --      --
```

16. On Sw1, enable port 1/1/1 (PC1) and 1/1/3 (PC3).

```
Sw1(config)# interface 1/1/1,1/1/3
Sw1(config-if-<1/1/1,1/1/3>)# no shutdown
```

17. Optional: On Agg1, Agg2, Sw1 and Sw2, you can make a new checkpoint named 'asts-lab04-<your-name>'.

```
# write memory
Copying configuration: [Success]
# copy running-config checkpoint asts-lab04-<your-name>
Copying configuration: [Success]
```

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.

```
Agg1(config)# router ospf 1
Agg1(config-ospf-1)# router-id 10.x.0.2
Agg1(config-ospf-1)# area 0.0.0.0
Agg1(config-ospf-1)# exit
Agg1(config)#
```

3. On Agg2, configure OSPF area 0. Use 10.x.0.3 as the router-id.

```
Agg2(config)# router ospf 1
Agg2(config-ospf-1)# router-id 10.x.0.3
Agg2(config-ospf-1)# area 0.0.0.0
Agg2(config-ospf-1)# exit
```

```
Agg2(config)#
```

4. On Agg1, configure and enable the loopback 0 and interface VLAN 1 for OSPF.

```
Agg1(config)# interface loopback 0
Agg1(config-loopback-if)# ip address 10.x.0.2/32
Agg1(config-loopback-if)# ip ospf 1 area 0
Agg1(config-loopback-if)# interface vlan 1
Agg1(config-if-vlan)# ip ospf 1 area 0
Agg1(config-if-vlan)# exit
```

5. On Agg2, configure and enable the loopback 0 and interface VLAN 1 for OSPF.

```
Agg2(config)# interface loopback 0
Agg2(config-loopback-if)# ip address 10.x.0.3/32
Agg2(config-loopback-if)# ip ospf 1 area 0
Agg2(config-loopback-if)# interface vlan 1
Agg2(config-if-vlan)# ip ospf 1 area 0
Agg2(config-if-vlan)# exit
```

6. On both Agg1 and Agg2, enable the SVI 11 and SVI 12 for OSPF and configure them as passive interfaces.

```
(config)# interface vlan 11
(config-if-vlan)# ip ospf 1 area 0
(config-if-vlan)# ip ospf passive
(config-if-vlan)# interface vlan 12
(config-if-vlan)# ip ospf 1 area 0
(config-if-vlan)# ip ospf passive
(config-if-vlan)# exit
```

7. Wait about 30 seconds for the adjacencies to establish, then verify the OSPF LSDB

```
Agg1# show ip ospf lsdb
OSPF Router with ID (10.x.0.2) (Process ID 1 VRF default)
=====
Router Link State Advertisements (Area 0.0.0.0)
-----
LSID          ADV Router    Age           Seq#           Checksum      Link Count
-----
10.x.0.2      10.x.0.2      105           0x80000004    0x0000b580    4
10.x.0.3      10.x.0.3      103           0x80000004    0x0000e150    4

Network Link State Advertisements (Area 0.0.0.0)
-----
LSID          ADV Router    Age           Seq#           Checksum
-----
10.x.1.2      10.x.0.2      364           0x80000001    0x0000688f
```

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.


Aruba Training-Confidential

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.

```
Agg1(config)# vlan 41
Agg1(config-vlan-41)# int 1/1/1
Agg1(config-if)# no routing
Agg1(config-if)# vlan trunk allow 41
Agg1(config-if)# exit
```

2. On Agg1, assign SVI 41 IP 10.x.41.2/24, enable it for OSPF area 0.

```
Agg1(config)# interface vlan 41
Agg1(config-if-vlan)# ip ad 10.x.41.2/24
Agg1(config-if-vlan)# ip ospf 1 area 0
Agg1(config-if-vlan)# exit
```

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.

```
Agg2(config)# vlan 42
Agg2(config-vlan-42)# int 1/1/1
Agg2(config-if)# no routing
Agg2(config-if)# no shutdown
Agg2(config-if)# vlan trunk allow 42
Agg2(config-if)# exit
```

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

```
Agg2(config)# interface vlan 42
Agg2(config-if-vlan)# ip ad 10.x.42.3/24
Agg2(config-if-vlan)# ip ospf 1 area 0
Agg2(config-if-vlan)# exit
```

Sw1

5. On Sw1, define VLAN 41, make 1/1/25 VLAN trunk and allow VLAN 41.

```
Sw1(config)# config t
Sw1(config)# vlan 41
Sw1(config-vlan-41)# int 1/1/25
Sw1(config-if)# vlan trunk allow 41
Sw1(config-if)# exit
```

6. On Sw1, assign SVI 41 IP 10.x.41.4/24.

```
Sw1(config)# interface vlan 41
Sw1(config-if-vlan)# ip ad 10.x.41.4/24
Sw1(config-if-vlan)# exit
```

7. On Sw1, define VLAN 42, make 1/1/26 VLAN trunk and allow VLAN 42.

```
Sw1(config)# vlan 42
Sw1(config-vlan-42)# int 1/1/26
Sw1(config-if)# vlan trunk allow 42
```

```
Sw1(config-if)# exit
```

8. On Sw1, assign SVI 42 IP 10.x.42.4/24.

```
Sw1(config)# interface vlan 42
Sw1(config-if-vlan)# ip ad 10.x.42.4/24
Sw1(config-if-vlan)# exit
```

9. On Sw1, define loopback 0 with IP 10.x.0.4.

```
Sw1(config)# interface loopback 0
Sw1(config-loopback-if)# ip address 10.x.0.4/32
Sw1(config-loopback-if)# exit
```

10. On Sw1, configure OSPF with router-id 10.x.0.4 and define area 0.

```
Sw1(config)# router ospf 1
Sw1(config-ospf-1)# router-id 10.x.0.4
Sw1(config-ospf-1)# area 0
Sw1(config-ospf-1)# exit
```

11. On Sw1, enable interfaces SVI41, SVI42 and the Loopback0 for OSPF.

```
Sw1(config)# interface loopback 0
Sw1(config-loopback-if)# ip ospf 1 area 0
Sw1(config-loopback-if)# interface vlan 41
Sw1(config-if-vlan)# ip ospf 1 area 0
Sw1(config-if-vlan)# interface vlan 42
Sw1(config-if-vlan)# ip ospf 1 area 0
Sw1(config-if-vlan)# exit
```

12. On Sw1, enable interfaces 1/1/25 and 1/1/26.

```
Sw1(config)# interface 1/1/25-1/1/26
Sw1(config-if-<1/1/25-1/1/26>)# no shutdown
Sw1(config-if-<1/1/25-1/1/26>)#
```

13. On Sw1, disable SVI 1.

```
Sw1(config-if-<1/1/25-1/1/26>)# interface vlan 1
Sw1(config-if-vlan)# shutdown
Sw1(config-if-vlan)# exit
```

14. On Sw1, verify the OSPF adjacencies to Agg1 and Agg2.

```
Sw1# show ip ospf neighbors
OSPF Process ID 1 VRF default
=====
```

Total Number of Neighbors: 1

Neighbor ID	Priority	State	Nbr Address	Interface
10.x.0.2	1	FULL/DR	10.x.41.2	vlan41

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'?

Sw1# show interface brief

Port Description	Native VLAN	Mode	Type	Enabled	Status	Reason	Speed (Mb/s)	
1/1/1 to P15-PC1-Fa...	1	access	1GbT	yes	up		100	Link
1/1/2	1	access	1GbT	no	down	Administratively down	--	--
1/1/3	11	access	1GbT	yes	up		1000	PC3
----- Output Omitted -----								
1/1/25 to TX-Agg1 (...)	1	trunk	SFP+DAC1	yes	up		10000	Link
1/1/26 to TX-Agg2 (...)	1	trunk	SFP+DAC1	yes	up		10000	Link
1/1/27 to TX-Sw2 (1...	1	trunk	SFP+DAC1	yes	up		10000	Link
1/1/28 to TX-Sw2 (1...	1	access	SFP+DAC1	yes	down	Waiting for link	--	Link
loopback0	--	routed	--	yes	up		--	--
vlan1	--	--	--	no	down	Administratively down	--	--
vlan41	--	--	--	yes	up		--	--
vlan42	--	--	--	yes	down		--	--

Are there any LLDP neighbors?

Sw1# show lldp neighbor

LLDP Neighbor Information
=====

Total Neighbor Entries : 5
Total Neighbor Entries Deleted : 3
Total Neighbor Entries Dropped : 0
Total Neighbor Entries Aged-Out : 3

LOCAL-PORT	CHASSIS-ID	PORT-ID	PORT-DESC	TTL	SYS-NAME
1/1/1 Fanout	00:24:a8:d9:1f:80	12	12	120	P15-PC1-

1/1/25	b8:d4:e7:41:38:00	1/1/1	1/1/1	120	TX-Agg1
1/1/26	b8:d4:e7:41:68:00	1/1/1	1/1/1	120	TX-Agg2
1/1/27	88:3a:30:a4:1f:40	1/1/27	1/1/27	120	TX-Sw2
mgmt	00:23:89:d9:c4:8b	GigabitEther...	TX-6300-A-00BM	120	P15-00BM-
Fanout					

Are the ports in the VLAN correctly assigned?

Sw1# show vlan 41

VLAN	Name	Status	Reason	Type	Interfaces
41	VLAN41	up	ok	static	1/1/25

show vlan 42

VLAN	Name	Status	Reason	Type	Interfaces
42	VLAN42	down	no_member_forwarding	static	1/1/26
42	VLAN42	down	no_member_forwarding	static	1/1/26

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.

Sw1# show spanning-tree

Spanning tree status : Enabled Protocol: MSTP

MST0

Root ID Priority : 4096
MAC-Address: 02:01:00:00:00:00
Hello time(in seconds):2 Max Age(in seconds):20
Forward Delay(in seconds):15

Bridge ID Priority : 32768
MAC-Address: 88:3a:30:a5:52:80
Hello time(in seconds):2 Max Age(in seconds):20
Forward Delay(in seconds):15

Port	Role	State	Cost	Priority	Type	BPDU-Tx
BPDU-Rx	TCN-Tx	TCN-Rx				


```

1/1/1      Designated      Forwarding 200000      128      P2P      173467
0          0              0
1/1/2      Disabled        Blocking  20000      128      P2P      0
0          0              0
1/1/3      Designated      Forwarding 20000      128      P2P      173466
0          0              0
----- Output Omitted -----
1/1/25     Root            Forwarding 2000      128      P2P      Bound      3
516        2              4
1/1/26     Alternate        Blocking  2000      128      P2P      Bound      2
437        0              2
1/1/27     Designated      Forwarding 2000      128      P2P      18
4          6              0
1/1/28     Disabled        Blocking  20000      128      P2P      0
0          0              0

Number of topology changes      : 5
Last topology change occurred   : 1027 seconds ago

```

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.

```
Sw1# show vlan port 1/1/25
```

```

-----
VLAN  Name                      Mode           Mapping
-----
41    VLAN41                     trunk          port

```

```
Sw1# show vlan port 1/1/26
```

```

-----
VLAN  Name                      Mode           Mapping
-----
42    VLAN42                     trunk          port

```

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

```

Sw1(config)# interface 1/1/25,1/1/26
Sw1(config-if-<1/1/25,1/1/26>)# spanning-tree bpdu-filter
Sw1(config-if-<1/1/25,1/1/26>)# exit

```

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

```

Sw1# show ip ospf neighbors

OSPF Process ID 1 VRF default
=====

```

Total Number of Neighbors: 2

Neighbor ID	Priority	State	Nbr Address	Interface
10.x.0.2	1	FULL/DR	10.x.41.2	vlan41
10.x.0.3	1	FULL/DR	10.x.42.3	vlan42

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.

```
Sw1# debug ospfv2 all
```

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

```
Sw1# terminal-monitor severity debug
```

Terminal Monitor is supported in SSH session only

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

```
Sw1(config)# interface 1/1/25
Sw1(config-if)# shutdown

2020-10-06T07:26:45.082884-0400 intfcd[736] <INFO> Event|404|LOG_INFO||Link status for
interface 1/1/25 is down
2020-10-06T07:26:45.154132-0400 hpe-routing[3471] <INFO>
debug|LOG_INFO|MSTR|1|OSPFV2|OSPFV2|OSPF 268698624 i/f idx 0X10000029 rtr ID 10.x.0.2 IP
addr 10.x.41.2 neighbor FSM has processed an input.
2020-10-06T07:26:45.154368-0400 hpe-routing[3471] <INFO>
debug|LOG_INFO|MSTR|1|OSPFV2|OSPFV2|FSM input = QOAM_NBR_KILL_NBR (12)
2020-10-06T07:26:45.154573-0400 hpe-routing[3471] <INFO>
debug|LOG_INFO|MSTR|1|OSPFV2|OSPFV2|Old FSM state = AMB_OSPF_NBR FULL (8)
2020-10-06T07:26:45.154766-0400 hpe-routing[3471] <INFO>
debug|LOG_INFO|MSTR|1|OSPFV2|OSPFV2|New FSM state = AMB_OSPF_NBR DOWN (1)
```

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

```
Sw1(config-if)# no shutdown

2020-10-06T07:33:29.257068-0400 hpe-routing[3471] <INFO>
debug|LOG_INFO|MSTR|1|OSPFV2|OSPFV2|Old FSM state = AMB_OSPF_NBR DOWN (1)
2020-10-06T07:33:29.259021-0400 hpe-routing[3471] <INFO>
debug|LOG_INFO|MSTR|1|OSPFV2|OSPFV2|Old FSM state = AMB_OSPF_NBR INIT (3)
----- Output Omitted -----
2020-10-06T07:33:29.270580-0400 hpe-routing[3471] <INFO>
debug|LOG_INFO|MSTR|1|OSPFV2|OSPFV2|Old FSM state = AMB_OSPF_NBR TWO_WAY (4)
----- Output Omitted -----
2020-10-06T07:33:29.285418-0400 hpe-routing[3471] <INFO>
debug|LOG_INFO|MSTR|1|OSPFV2|OSPFV2|Old FSM state = AMB_OSPF_NBR EXCHANGE_START (5)
```

```
----- Output Omitted -----
2020-10-06T07:33:29.288222-0400 hpe-routing[3471] <INFO>
debug|LOG_INFO|MSTR|1|OSPFV2|OSPFV2|Old FSM state = AMB_OSPF_NBR_EXCHANGE (6)
----- Output Omitted -----
2020-10-06T07:33:29.288407-0400 hpe-routing[3471] <INFO>
debug|LOG_INFO|MSTR|1|OSPFV2|OSPFV2|New FSM state = AMB_OSPF_NBR_FULL (8)
```

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

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

23. Disable the terminal monitor.

```
Sw1# no terminal-monitor
```

Invalid Timers

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

```
Agg1(config)# interface vlan 41
Agg1(config-if-vlan)# ip ospf hello-interval 5
```

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

```
Agg1(config-if-vlan)# ip ospf dead-interval 15
```

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

```
Agg1# show ip ospf neighbors
OSPF Process ID 1 VRF default
=====
Total Number of Neighbors: 1

Neighbor ID      Priority   State          Nbr Address      Interface
-----
10.x.0.3         1         FULL/BDR       10.x.1.3         vlan1
```

OSPFv2 Debugging

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

```
Sw1# do clear debug buffer
Sw1# show debug buffer reverse

-----
show debug buffer
-----
2020-10-06:07:54:19.279361|hpe-routing|LOG_ERR|MSTR|1|OSPFV2|OSPFV2|(OSPFv2 process Id = 1,
VRF Name = default)
2020-10-06:07:54:19.279347|hpe-routing|LOG_ERR|MSTR|1|OSPFV2|OSPFV2|Neighboring Hello
Interval = 5
2020-10-06:07:54:19.279331|hpe-routing|LOG_ERR|MSTR|1|OSPFV2|OSPFV2|My Hello Interval = 10
```

```
2020-10-06:07:54:19.279309|hpe-routing|LOG_ERR|MSTR|1|OSPFV2|OSPFV2|OSPF 268698624 Hello
packet with mismatched hello interval received from router 10.x.0.2.
```

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.

```
Sw1(config)# interface vlan 41
Sw1(config-if-vlan)# ip ospf hello-interval 5
```

Q: Did the adjacency come online now?

A: No, so something else must be wrong.

```
2020-10-06:09:20:14.445296|hpe-routing|LOG_ERR|MSTR|1|OSPFV2|OSPFV2|My Dead Interval = 40
2020-10-06:09:20:14.445311|hpe-routing|LOG_ERR|MSTR|1|OSPFV2|OSPFV2|Neighboring Dead Interval
= 15
2020-10-06:09:20:14.445325|hpe-routing|LOG_ERR|MSTR|1|OSPFV2|OSPFV2|(OSPFv2 process Id = 1,
VRF Name = default)
2020-10-06:09:20:24.444992|hpe-routing|LOG_ERR|MSTR|1|OSPFV2|OSPFV2|OSPF 268698624 Hello
packet with mismatched dead interval received from router 10.x.0.2.
```

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

```
Sw1(config-if-vlan)# ip ospf dead-interval 15
```

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

```
Sw1# show ip ospf statistics interface vlan41
----- Output Omitted -----
Reason                               Packets Dropped
-----
Invalid type                         0
Invalid length                       0
Invalid checksum                     0
Invalid version                      0
Bad or unknown source                1
Area mismatch                        0
Self-originated                      0
Duplicate router ID                  0
Interface standby                    0
Total Hello packets dropped          211
Network Mask mismatch                0
Hello interval mismatch              185
Dead interval mismatch                26
Options mismatch                     0
```

```

MTU mismatch          0
Neighbor ignored      0
Authentication errors  0
Type mismatch         0
Authentication failures 0
Wrong protocol        0
Resource failures     0
Bad LSA length        0
Bad DD packets        0
----- Output Omitted -----

```

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.

```

Agg1(config)# interface vlan 41
Agg1(config-if-vlan)# ip ospf authentication simple-text
Agg1(config-if-vlan)# ip ospf authentication-key plaintext aruba123

```

32. On Sw1, review the OSPF interface statistics

```
Sw1# show ip ospf statistics interface vlan41
```

Reason	Packets Dropped
Invalid type	0
Invalid length	0
Invalid checksum	0
Invalid version	0
Bad or unknown source	1
Area mismatch	0
Self-originated	0
Duplicate router ID	0
Interface standby	0
Total Hello packets dropped	211
Network Mask mismatch	0
Hello interval mismatch	185
Dead interval mismatch	26
Options mismatch	0
MTU mismatch	0
Neighbor ignored	0
Authentication errors	8
Type mismatch	8
Authentication failures	0
Wrong protocol	0
Resource failures	0
Bad LSA length	0
Bad DD packets	0

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

```
Sw1# clear debug buffer
Sw1# show debug buffer

-----
show debug buffer
-----
2020-10-06:09:58:09.514562|hpe-routing|LOG_WARN|MSTR|1|OSPFV2|OSPFV2|(OSPFv2 process Id = 1,
VRF Name = default)
2020-10-06:09:58:14.512376|hpe-routing|LOG_WARN|MSTR|1|OSPFV2|OSPFV2|OSPF 268698624 Packet
received with unexpected authentication type 1.
2020-10-06:09:58:14.512990|hpe-routing|LOG_WARN|MSTR|1|OSPFV2|OSPFV2|Expected authentication
type = 0.
```

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.

```
Sw1# diagnostics
Sw1# diag utilities tcpdump proto 89 count 4 vrf default

tcpdump: listening on any, link-type LINUX_SLL (Linux cooked), capture size 262144 bytes
1 10:00:25.032835 ethertype IPv4, IP (tos 0xc0, ttl 1, id 58268, offset 0, flags [none],
proto OSPF (89), length 68)
  10.x.1.2 > ospf-all.mcast.net: OSPFv2, Hello, length 48
    Router-ID 10.x.0.2, Backbone Area, Authentication Type: none (0)
    Options [External]
    Hello Timer 10s, Dead Timer 40s, Mask 255.255.255.0, Priority 1
    Designated Router 10.x.1.2, Backup Designated Router 10.x.1.3
    Neighbor List:
      10.x.0.3
2 10:00:27.559870 ethertype IPv4, IP (tos 0xc0, ttl 1, id 26597, offset 0, flags [none],
proto OSPF (89), length 68)
  10.x.1.3 > ospf-all.mcast.net: OSPFv2, Hello, length 48
    Router-ID 10.x.0.3, Backbone Area, Authentication Type: none (0)
    Options [External]
    Hello Timer 10s, Dead Timer 40s, Mask 255.255.255.0, Priority 1
    Designated Router 10.x.1.2, Backup Designated Router 10.x.1.3
    Neighbor List:
      10.x.0.2
3 10:00:28.567285 IP (tos 0xc0, ttl 1, id 36168, offset 0, flags [none], proto OSPF
(89), length 64)
  TX-Sw1 > ospf-all.mcast.net: OSPFv2, Hello, length 44
    Router-ID TX-Sw1, Backbone Area, Authentication Type: none (0)
    Options [External]
    Hello Timer 5s, Dead Timer 15s, Mask 255.255.255.0, Priority 1
    Designated Router TX-Sw1
4 10:00:29.517020 ethertype IPv4, IP (tos 0xc0, ttl 1, id 34853, offset 0, flags [none],
proto OSPF (89), length 64)
  10.x.41.2 > ospf-all.mcast.net: OSPFv2, Hello, length 44
```

```
Router-ID 10.x.0.2, Backbone Area, Authentication Type: simple (1)
Simple text password: aruba123
Options [External]
  Hello Timer 5s, Dead Timer 15s, Mask 255.255.255.0, Priority 1
  Designated Router 10.x.41.2
4 packets captured
5 packets received by filter
0 packets dropped by kernel
```

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.

```
Agg1(config)# interface vlan 41
Agg1(config-if-vlan)# ip ospf authentication null
Agg1(config-if-vlan)# exit
```

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.

```
Agg1(config)# router ospf 1
Agg1(config-ospf-1)# router-id 10.x.0.4
```

37. Confirm to reset the OSPF process.

```
OSPFv2 protocol will be reset.
Do you want to continue (y/n)? y
```

38. On Sw1, review the OSPF adjacencies.

```
Sw1# show ip ospf neighbors
OSPF Process ID 1 VRF default
=====
Total Number of Neighbors: 1
```

Neighbor ID	Priority	State	Nbr Address	Interface
10.x.0.3	1	FULL/DR	10.x.42.3	vlan42

39. Review the debug buffer. You should see an OSPF ERR message about the neighbor router ID.

```
Sw1# clear debug buffer
Sw1# show debug buffer

-----
show debug buffer
-----
2020-10-06:10:08:17.339143|hpe-routing|LOG_ERR|MSTR|1|OSPFV2|OSPFV2|(OSPFv2 process Id = 1,
VRF Name = default)
2020-10-06:10:08:22.339567|hpe-routing|LOG_ERR|MSTR|1|OSPFV2|OSPFV2|OSPF 268698624 IP addr
10.x.41.4 i/f idx 0X00000000 Hello packet dropped because source router ID matches local
router ID.
2020-10-06:10:08:22.340007|hpe-routing|LOG_ERR|MSTR|1|OSPFV2|OSPFV2|Source IP address =
10.x.41.2
2020-10-06:10:08:22.340276|hpe-routing|LOG_ERR|MSTR|1|OSPFV2|OSPFV2|(OSPFv2 process Id = 1,
VRF Name = default)
```

40. Review the OSPF interface statistics, notice the 'Duplicate router ID' counter.

```
Sw1# show ip ospf statistics interface vlan41
```

```
----- Output Omitted -----
Reason                      Packets Dropped
-----
Invalid type                 0
Invalid length               0
Invalid checksum             0
Invalid version              0
Bad or unknown source        1
Area mismatch                0
Self-originated              0
Duplicate router ID          71
Interface standby            0
----- Output Omitted -----
```

41. On Agg1, correct the router-id again to 10.x.0.2.

```
Agg1(config-ospf-1)# router-id 10.x.0.2
```

```
OSPFv2 protocol will be reset.
Do you want to continue (y/n)? y
```

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.

```
Sw1# conf t
Sw1(config)# int vlan 41
Sw1(config-if-vlan)# ip ospf network point-to-point
```



```
Sw1(config-if-vlan)# int vlan 42
Sw1(config-if-vlan)# ip ospf network point-to-point
Sw1(config-if-vlan)# exit
```

43. On Agg1, configure SVI 41 as OSPF point to point.

```
Agg1(config)# interface vlan 41
Agg1(config-if-vlan)# ip ospf network point-to-point
Agg1(config-if-vlan)# exit
```

44. On Agg2, configure SVI 42 as OSPF point to point.

```
Agg2(config)# interface vlan 42
Agg2(config-if-vlan)# ip ospf network point-to-point
Agg2(config-if-vlan)# exit
```

Test the point to point

45. On Agg1, disable port 1/1/1 (connected to Sw1, transporting SVI 41)

```
Agg1(config)# interface 1/1/1
Agg1(config-if)# shutdown
```

46. On Sw1, enable OSPF debugging and terminal-monitor for INFO

```
Sw1# debug ospfv2 event
Sw1# terminal-monitor severity info filter hpe-routing
```

47. On Agg1, enable the port 1/1/1.

```
Agg1(config-if)# no shutdown
```

48. On Sw1, review the debug output.

```
2020-10-06T10:49:12.572515-0400 hpe-routing[3471] <INFO>
debug|LOG_INFO|MSTR|1|OSPFV2|OSPFV2|Old FSM state = AMB_OSPF_IF_DOWN (1)
2020-10-06T10:49:12.572712-0400 hpe-routing[3471] <INFO>
debug|LOG_INFO|MSTR|1|OSPFV2|OSPFV2|New FSM state = AMB_OSPF_IF_POINT_TO_POINT (4)
2020-10-06T10:49:17.076813-0400 hpe-routing[3471] <INFO>
debug|LOG_INFO|MSTR|1|OSPFV2|OSPFV2|Old FSM state = AMB_OSPF_NBR_DOWN (1)
2020-10-06T10:49:17.077016-0400 hpe-routing[3471] <INFO>
debug|LOG_INFO|MSTR|1|OSPFV2|OSPFV2|New FSM state = AMB_OSPF_NBR_INIT (3)
2020-10-06T10:49:17.079686-0400 hpe-routing[3471] <INFO>
debug|LOG_INFO|MSTR|1|OSPFV2|OSPFV2|Old FSM state = AMB_OSPF_NBR_TWO_WAY (4)
2020-10-06T10:49:17.079876-0400 hpe-routing[3471] <INFO>
debug|LOG_INFO|MSTR|1|OSPFV2|OSPFV2|New FSM state = AMB_OSPF_NBR_EXCHANGE_START (5)
```

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.

```
Agg1# show lacp interfaces
State abbreviations :
A - Active           P - Passive           F - Aggregable I - Individual
S - Short-timeout    L - Long-timeout  N - InSync       O - OutofSync
C - Collecting       D - Distributing
X - State m/c expired      E - Default neighbor state

Actor details of all interfaces:
-----
Intf      Aggr   Port  Port  State  System-ID          System Aggr  Forwarding
      Name   Id    Pri             Pri      Key   State
-----
1/1/2     lag2(mc)  2     1     ALFNCD 02:01:00:00:00:00 65534  2     up
1/1/5     lag5(mc)  5     1     ALFNCD 02:01:00:00:00:00 65534  5     up
1/1/46    lag256    47    1     ALFNCD b8:d4:e7:41:38:00 65534  256   up
1/1/47    lag256    48    1     ALFNCD b8:d4:e7:41:38:00 65534  256   up
```

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

```
Agg1(config)# vlan 40
Agg1(config-vlan-40)# int vlan 40
Agg1(config-if-vlan)# ip address 10.x.40.2/24
Agg1(config-if-vlan)# ip ospf 1 area 0
Agg1(config-if-vlan)# exit
Agg1(config-if-vlan)# interface lag 5
Agg1(config-lag-if)# vlan trunk allowed 40
Agg1(config-lag-if)# exit
```

```
Agg2(config)# vlan 40
Agg2(config-vlan-40)# int vlan 40
Agg2(config-if-vlan)# ip address 10.x.40.3/24
Agg2(config-if-vlan)# ip ospf 1 area 0
Agg2(config-if-vlan)# exit
Agg2(config-if-vlan)# interface lag 5
Agg2(config-lag-if)# vlan trunk allowed 40
Agg2(config-lag-if)# exit
```

- 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.

```
Agg1# show ip ospf neighbors
OSPF Process ID 1 VRF default
=====
Total Number of Neighbors: 4
Neighbor ID      Priority  State                Nbr Address      Interface
-----
10.x.0.3         1        FULL/DR              10.x.1.3         vlan1
10.x.0.6         1        FULL/DROther        10.x.40.6        vlan40
10.x.0.3         1        FULL/DR              10.x.40.3        vlan40
10.x.0.4         n/a      FULL                 10.x.41.4        vlan41
```

- 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.

```
Aruba_gw# show ip ospf neighbor
OSPF Neighbor Table
-----
Neighbor ID  Pri  State      Address      Interface
-----
10.x.0.2     1    FULL/BDR   10.x.40.2    Vlan 40
10.x.0.3     1    FULL/DR    10.x.40.3    Vlan 40
```

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.

```
Sw1(config)# interface loopback 0
Sw1(config-loopback-if)# ip address 10.x.0.7/32 secondary
Sw1(config-loopback-if)# exit
```

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.

```
Sw1# ping 10.x.40.6 source 10.x.0.4 rep 1
Sw1# ping 10.x.40.6 source 10.x.0.7 rep 1
```

7. On the Aruba Gateway, review the ARP table so you know the MAC of Agg1 and Agg2 on VLAN40.

```
Aruba_gw# show arp
```

Protocol	Address	Hardware Address	Interface
Internet	10.x.40.3	b8:d4:e7:41:68:00	vlan40
Internet	10.251.x.200	00:50:56:b1:8a:3c	vlan2
Internet	10.251.x.91	00:50:56:b1:2d:b9	vlan2
Internet	10.x.40.2	b8:d4:e7:41:38:00	vlan40

8. On Agg1, define a class ip 'icmp' that matches any ICMP traffic.

```
Agg1(config)# class ip icmp
Agg1(config-class-ip)# vsx-sync
Agg1(config-class-ip)# match icmp any any count
Agg1(config-class-ip)# exit
```

9. Define a policy 'mirror' that applies action 'mirror 1' to the class 'icmp'

```
Agg1(config)# policy mirror
Agg1(config-policy)# vsx-sync
```

```
Agg1(config-policy)# class ip icmp action mirror 1
Agg1(config-policy)# exit
```

10. Apply the policy 'mirror' to LAG5 (Aruba GW) and LAG256 (VSX ISL).

```
Agg1(config)# int lag 5
Agg1(config-lag-if)# apply policy mirror in
Agg1(config-lag-if)# int lag 256
Agg1(config-lag-if)# apply policy mirror in
Agg1(config-lag-if)# exit
```

11. Define mirror 1 with destination CPU, enable the mirror.

```
Agg1(config)# mirror session 1
Agg1(config-mirror-1)# destination cpu
Agg1(config-mirror-1)# enable
Agg1(config-mirror-1)# exit
```

12. On Agg2, repeat the mirror definition and verify the other settings have been synced.

```
Agg2(config)# mirror session 1
Agg2(config-mirror-1)# destination cpu
Agg2(config-mirror-1)# enable
Agg2(config-mirror-1)# exit
```

```
Agg2# show policy
```

```

      Name
      Additional Policy Parameters
Sequence Comment
      Class Type
              action
-----
      mirror
      VSX synchronization: Enabled
10      icmp  ipv4
              mirror 1
```

```
Agg2# show run interface lag 5
```

```
interface lag 5 multi-chassis
no shutdown
no routing
vlan trunk native 1
vlan trunk allowed 1,11-12,40
lacp mode active
apply policy mirror in
exit
```

13. Start 'diag utilities tshark' on both Agg1 and Agg2.

```
(config)# end
```

```
# diagnostics
# diag utilities tshark
```

Testing

14. On Sw1, send a single ping to 10.x.40.6 using source IP 10.x.41.4.

```
Sw1# ping 10.x.40.6 source 10.x.41.4 rep 1
```

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.

```
Sw1# ping 10.x.40.6 source 10.x.42.4 rep 1
```

17. Take note of the source/destination IP combination ping that showed up on both switches.

```
Ethernet II, Src: 20:4c:03:b1:ce:ba, Dst: b8:d4:e7:41:38:00
Ethernet II, Src: 20:4c:03:b1:ce:ba, Dst: b8:d4:e7:41:38:00
```

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.

```
# show vsx active-forwarding

No VSX active-forwarding enabled interfaces found.
```

20. On Agg1, disable IP ICMP redirects and enable 'vsx active-forwarding' on the VLAN 40 SVI.

```
Agg1# conf t
Agg1(config)# no ip icmp redirect
Agg1(config)# interface vlan 40
Agg1(config-if-vlan)# vsx active-forwarding
Agg1(config-if-vlan)# exit
```

21. Review the VSX Active-forwarding state on Agg1.

```
Agg1# show vsx active-forwarding interface vlan40

Interface vlan40 has VSX active-forwarding enabled.
No peer info available for interface vlan40.
```


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.

```
Agg2# show vsx active-forwarding
No VSX active-forwarding enabled interfaces found.
```

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.

```
Agg2# conf t
Agg2(config)# no ip icmp redirect
Agg2(config)# interface vlan 40
Agg2(config-if-vlan)# vsx active-forwarding
Agg2(config-if-vlan)# exit
```

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

```
# show vsx active-forwarding interface vlan40
Interface vlan40 has VSX active-forwarding enabled.
Interface vlan40 Peer Data:
Peer MAC: b8:d4:e7:41:68:00
Peer primary IPv4 Address: 10.x.40.3
```

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.


```
(config)# interface lag 256
(config-lag-if)# no apply policy mirror in
(config-lag-if)# int lag 5
(config-lag-if)# no apply policy mirror in
(config-lag-if)# exit
```

30. On Sw1, remove the secondary IP 10.x.0.7 from the Loopback0 interface.

```
Sw1# conf t
Sw1(config)# interface loopback 0
Sw1(config-loopback-if)# no ip address 10.x.0.7/32 secondary
Sw1(config-loopback-if)# exit
```

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.

```
Agg1# conf t
Agg1(config)# keychain asts
Agg1(config-keychain)# key 1
Agg1(config-keychain-key)# key-string plaintext aruba123
Agg1(config-keychain-key)# end
```

2. Enter the SVI 41 context on Agg1, link the 'asts' keychain to VLAN 41 SVI using the command 'ip ospf keychain asts'

```
Agg1(config)# interface vlan41
Agg1(config-if-vlan)# ip ospf keychain asts
Agg1(config-if-vlan)# exit
```

3. Explore the 'ip ospf authentication' command options, however, do **not** enable keychain authentication yet.

```
Agg1(config-if-vlan)# ip ospf authentication ?
keychain          Configure keychain authentication
message-digest    Configure message-digest authentication
null              Configure null authentication
simple-text        Configure simple-text authentication
```

Q: What are the authentication options?

A:	null simple-text message-digest keychain	default> no authentication clear-text password single MD5 key with a single key-id 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.

```
Sw1(config)# keychain asts
Sw1(config-keychain)# key 1
Sw1(config-keychain-key)# key-string plaintext aruba123
Sw1(config-keychain-key)# end
Sw1# conf t
Sw1(config)# interface vlan41
Sw1(config-if-vlan)# ip ospf keychain asts
Sw1(config-if-vlan)# exit
```

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).

```
(config-if-vlan)# ip ospf authentication keychain
(config-if-vlan)# exit
```

6. On Agg1, verify that the adjacency is formed correctly with the new key without resetting the adjacency.

```
Agg1# show ip ospf neighbors 10.x.0.4
Neighbor 10.x.0.4, interface address 10.x.41.4
-----

Process ID 1 VRF default, in area 0.0.0.0 via interface vlan41
Neighbor priority is n/a, State is FULL
Options is 0x42
Dead timer due in 00:00:13
Time since last state change 05h:02m:47s
```

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'

```
Sw1(config)# keychain asts
Sw1(config-keychain)# key 2
Sw1(config-keychain-key)# key-string plaintext arubaRocks
Sw1(config-keychain-key)# end
```

8. On Agg1, wait about 15 seconds, check the OSPF neighbors.

```
Agg1# show ip ospf neighbors
```

```
OSPF Process ID 1 VRF default
=====
```

```
Total Number of Neighbors: 2
```

Neighbor ID	Priority	State	Nbr Address	Interface
10.x.0.2	n/a	INIT	10.x.41.2	vlan41
10.x.0.3	n/a	FULL	10.x.42.3	vlan42

Q: What do you observe?

A: The adjacency with Sw1 was lost.

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

```
Agg1# show ip ospf statistics interface vlan41
```

```
----- Output Omitted -----
```

Reason	Packets Dropped

```
----- Output Omitted -----
```

MTU mismatch	0

Neighbor ignored	0

Authentication errors	35

Type mismatch	1

Authentication failures	34

Wrong protocol	0

```
----- Output Omitted -----
```

Q: What do you observe?

A: The authentication failures counter is increasing.

10. On Agg1, enable ospf all debugging.

```
Agg1# debug ospfv2 all
```

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

```
Agg1# terminal-monitor severity debug filter hpe-routing
```

```
Terminal-monitor is enabled successfully
```

```
2020-10-06T16:35:53.144439-0400 hpe-routing[8390] <ERR> debug|LOG_ERR|AMM|-
|OSPFV2|OSPFV2|OSPF 268698624 Packet cannot be authenticated.
```

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.

```
Sw1# show keychain asts
Keychain Name : asts
Number of Keys      : 2
Active Send Key ID  : 2
Active Recv Key IDs : 1,2
----- Output Omitted -----
```

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

```
Agg1# no debug all
```

TCPDUMP on Agg1

14. On Agg1, enable diagnostics mode

```
Agg1# diagnostics
```

15. Use TCPDUMP of the OSPF protocol (protocol 89).

```
Agg1# diag utilities tcpdump proto 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).

```
8 16:42:53.149766 ethertype IPv4, IP (tos 0xc0, ttl 1, id 12688, offset 0, flags [none],
proto OSPF (89), length 84)
  10.x.41.4 > 224.0.0.5: OSPFv2, Hello, length 48
    Router-ID 10.x.0.4, Backbone Area, Authentication Type: MD5 (2)
    Key-ID: 2, Auth-Length: 16, Crypto Sequence Number: 0x00008a74
    Options [External]
      Hello Timer 5s, Dead Timer 15s, Mask 255.255.255.0, Priority 1
      Neighbor List:
        TX-Agg1
```

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.

```
Sw1# show keychain asts
Keychain Name : asts
Number of Keys      : 2
Active Send Key ID  : 2
Active Recv Key IDs : 1,2

Key ID : 1
Key string      : flXXyBPHucV4ScfFigKSUFWcNJavUClNBQenFaJwwRo=
Send Key Validity : 19:00:00 12/31/2019 to Infinite
Recv Key Validity : 19:00:00 12/31/2019 to Infinite
Key ID : 2
Key string      : bP7CWG8jbdP270Szr+t3sKHHmQ0N1kPn0cIvArFLnM4=
Send Key Validity : 19:00:00 12/31/2019 to Infinite
Recv Key Validity : 19:00:00 12/31/2019 to Infinite
```

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.

```
Sw1# conf t
Sw1(config)# keychain asts
Sw1(config-keychain)# no key 2
Sw1(config-keychain)# end
```

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

```
Sw1# show ip ospf neighbors
OSPF Process ID 1 VRF default
=====

Total Number of Neighbors: 2
```

Neighbor ID	Priority	State	Nbr Address	Interface
10.x.0.2	n/a	FULL	10.x.41.2	vlan41
10.x.0.3	n/a	FULL	10.x.42.3	vlan42

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

```
Agg1(config)# keychain asts
```

```
Agg1(config-keychain)# key 2
```

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

```
Agg1(config-keychain-key)# send-lifetime start-time 00:00:00 1/1/2030 duration infinite
```

22. Configure the key-string 'arubaRocks'

```
Agg1(config-keychain-key)# key-string plaintext arubaRocks
```

23. Review the keychain status.

```
Agg1# show keychain asts
```

```
Keychain Name : asts
Number of Keys : 2
Active Send Key ID : 1
Active Recv Key IDs : 1,2

Key ID : 1
Key string : r2aKp7v21MqZmd2KTCB+iVWcNJavUC1NBQenFaJwwRo=
Send Key Validity : 19:00:00 12/31/2019 to Infinite
Recv Key Validity : 19:00:00 12/31/2019 to Infinite

Key ID : 2
Key string : Uxy+4kunStWFG3Jo0BIAWqHHmQ0N1kPn0cIvArFLnM4=
Send Key Validity : 00:00:00 01/01/2030 to Infinite
Recv Key Validity : 19:00:00 12/31/2019 to Infinite
```

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.

```
Sw1# show ip ospf neighbors
```

```
OSPF Process ID 1 VRF default
```

```
=====
```

```
Total Number of Neighbors: 2
```

Neighbor ID	Priority	State	Nbr Address	Interface
10.x.0.2	n/a	FULL	10.x.41.2	vlan41
10.x.0.3	n/a	FULL	10.x.42.3	vlan42

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

```
Sw1# conf t
```

```
Sw1(config)# keychain asts
```

```
Sw1(config-keychain)# key 2
Sw1(config-keychain-key)# key-string plaintext arubaRocks
Sw1(config-keychain-key)# end
```

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

```
Sw1# show keychain asts
Keychain Name : asts
Number of Keys      : 2
Active Send Key ID  : 2
Active Recv Key IDs : 1,2
----- Output Omitted -----
```

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.

```
Agg1(config-keychain-key)# send-lifetime start-time 00:00:00 1/1/2020 duration infinite
```

28. On Agg1, review the keychain status.

```
Agg1# show keychain asts
Keychain Name : asts
Number of Keys      : 2
Active Send Key ID  : 2
Active Recv Key IDs : 1,2
```

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.

```
# keychain asts
(config-keychain)# no key 1
(config-keychain)# show ip ospf neighbors

OSPF Process ID 1 VRF default
=====
Total Number of Neighbors: 2

Neighbor ID      Priority  State              Nbr Address        Interface
-----
```


10.x.0.2	n/a	FULL	10.x.41.2	vlan41
10.x.0.3	n/a	FULL	10.x.42.3	vlan42

Aruba Training-Confidential

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.

```
Sw1# show ip route 10.x.0.6
Displaying ipv4 routes selected for forwarding

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

10.x.0.6/32, vrf default
    via 10.x.41.2, [110/201], ospf
    via 10.x.42.3, [110/201], ospf
```

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.

```
Agg2(config)# router ospf 1
Agg2(config-ospf-1)# disable
```

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

```
Agg2(config-ospf-1)# max-metric router-lsa on-startup 60
```

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

```
Sw1# show ip route 10.x.0.6
Sw1# repeat delay 5
```

35. On Agg2, enable OSPF again.

```
Agg2(config-ospf-1)# enable
```

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

```
Agg2# show ip ospf routes
```

```
Codes: i - Intra-area route, I - Inter-area route  
       E1 - External type-1, E2 - External type-2
```

```
OSPF Process ID 1 VRF default, Routing Table
```

```
-----
```

```
Total Number of Routes : 12
```

```
10.x.0.2/32      (i) area: 0.0.0.0  
    via 10.x.1.2 interface vlan1, cost 65535 distance 110  
10.x.0.2/32      (i) area: 0.0.0.0  
    via 10.x.40.2 interface vlan40, cost 65535 distance 110  
10.x.0.4/32      (i) area: 0.0.0.0  
    via 10.x.42.4 interface vlan42, cost 65535 distance 110  
10.x.0.6/32      (i) area: 0.0.0.0  
    via 10.x.40.6 interface vlan40, cost 65536 distance 110
```

```
----- Output Omitted -----
```

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.

```
*****  
Iteration : 11 Command : show ip route 10.x.0.6  
*****
```

```
Displaying ipv4 routes selected for forwarding
```

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

```
10.x.0.6/32, vrf default  
    via 10.x.41.2, [110/201], ospf
```

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.

```
Agg2# show ip ospf routes
```

```
Codes: i - Intra-area route, I - Inter-area route  
       E1 - External type-1, E2 - External type-2
```

```
OSPF Process ID 1 VRF default, Routing Table
```

```
-----
```

Total Number of Routes : 12

```
10.x.0.2/32      (i) area: 0.0.0.0
  via 10.x.1.2 interface vlan1, cost 100 distance 110
10.x.0.2/32      (i) area: 0.0.0.0
  via 10.x.40.2 interface vlan40, cost 100 distance 110
10.x.0.4/32      (i) area: 0.0.0.0
  via 10.x.42.4 interface vlan42, cost 100 distance 110
10.x.0.6/32      (i) area: 0.0.0.0
  via 10.x.40.6 interface vlan40, cost 101 distance 110
----- Output Omitted -----
```

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.

```
Sw1# conf t
Sw1(config)# vlan 46
Sw1(config-vlan-46)# vlan 47
Sw1(config-vlan-47)# int 1/1/27
Sw1(config-if)# vlan trunk allowed 46,47
Sw1(config-if)# exit
```

2. 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

```
Sw1(config)# int vlan 46
Sw1(config-if-vlan)# ip address 10.x.46.4/24
Sw1(config-if-vlan)# l3-counters
Sw1(config-if-vlan)# ip ospf 1 area 0
Sw1(config-if-vlan)# ip ospf cost 10
Sw1(config-if-vlan)# ip ospf network point-to-point
```

```
Sw1(config-if-vlan)# int vlan 47
Sw1(config-if-vlan)# ip address 10.x.47.4/24
Sw1(config-if-vlan)# l3-counters
Sw1(config-if-vlan)# ip ospf 1 area 0
Sw1(config-if-vlan)# ip ospf cost 100
Sw1(config-if-vlan)# ip ospf network point-to-point
Sw1(config-if-vlan)# exit
```

3. On Sw2, define VLAN 46 and VLAN47, enable them on the port 1/1/27.

```
Sw2# conf t
Sw2(config)# vlan 46
Sw2(config-vlan-46)# vlan 47
Sw2(config-vlan-47)# int 1/1/27
Sw2(config-if)# vlan trunk allowed 46,47
Sw2(config-if)# exit
```

4. On Sw2, define VRF 'branch', enable OSPF process 1 for VRF 'branch' and define area 0.

```
Sw2(config)# vrf branch
Sw2(config-vrf)# router ospf 1 vrf branch
Sw2(config-ospf-1)# enable
Sw2(config-ospf-1)# area 0
Sw2(config-ospf-1)# exit
```

5. On Sw2, define SVI 46 and 47, attach them to VRF 'branch'.

```
Sw2(config)# int vlan 46
Sw2(config-if-vlan)# vrf attach branch
Sw2(config-if-vlan)# int vlan 47
```

```
Sw2(config-if-vlan)# vrf attach branch
Sw2(config-if-vlan)# exit
```

6. 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

```
Sw2(config)# int vlan 46
Sw2(config-if-vlan)# ip address 10.x.46.5/24
Sw2(config-if-vlan)# l3-counters
Sw2(config-if-vlan)# ip ospf 1 area 0
Sw2(config-if-vlan)# ip ospf cost 10
Sw2(config-if-vlan)# ip ospf network point-to-point
```

```
Sw2(config-if-vlan)# int vlan 47
Sw2(config-if-vlan)# ip address 10.x.47.5/24
Sw2(config-if-vlan)# l3-counters
Sw2(config-if-vlan)# ip ospf 1 area 0
Sw2(config-if-vlan)# ip ospf cost 100
Sw2(config-if-vlan)# ip ospf network point-to-point
Sw2(config-if-vlan)# exit
```

7. On Sw2, verify that both interfaces SVI 46 and SVI 47 have active OSPF adjacencies.

```
Sw2# show ip ospf neighbors vrf branch
OSPF Process ID 1 VRF branch
=====
Total Number of Neighbors: 2

Neighbor ID      Priority   State             Nbr Address      Interface
-----
10.x.0.4         n/a      FULL              10.x.46.4        vlan46
10.x.0.4         n/a      FULL              10.x.47.4        vlan47
```

8. 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

```
Sw2(config)# vlan 15
Sw2(config-vlan-15)# interface 1/1/4
Sw2(config-if)# vlan access 15
Sw2(config-if)# exit
```

```
Sw2(config)# interface vlan 15
Sw2(config-if-vlan)# description branch-pc
Sw2(config-if-vlan)# vrf attach branch
```

```
Sw2(config-if-vlan)# ip address 10.x.15.1/24
Sw2(config-if-vlan)# ip ospf 1 area 0
Sw2(config-if-vlan)# exit
```

```
Sw2(config)# dhcp-server vrf branch
Sw2(config-dhcp-server)# pool vlan15
Sw2(config-dhcp-server-pool)# range 10.x.15.200 10.x.15.210 prefix-len 24
Sw2(config-dhcp-server-pool)# default-router 10.x.15.1
Sw2(config-dhcp-server-pool)# exit
Sw2(config-dhcp-server)# authoritative
Sw2(config-dhcp-server)# enable
Sw2(config-dhcp-server)# exit
```

9. 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

10. 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.

11. 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

12. 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.

```
Sw2# show ip ospf routes vrf branch
Codes: i - Intra-area route, I - Inter-area route
       E1 - External type-1, E2 - External type-2

OSPF Process ID 1 VRF branch, Routing Table
-----

Total Number of Routes : 12

----- Output Omitted -----
10.x.12.0/24      (i) area: 0.0.0.0
    via 10.x.46.4 interface vlan46, cost 210 distance 110
----- Output Omitted -----
```


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.

```
Sw1# clear interface statistics
```

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.

```
Sw1# show interface vlan 46
Sw1# show interface vlan 47
Sw1# repeat id 1,2 delay 10
```

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

```
*****
Iteration : 6 Command : show interface vlan 47
*****

Interface vlan47 is up
Admin state is up
Description:
Hardware: Ethernet, MAC Address: 88:3a:30:a5:52:80
IPv4 address 10.x.47.4/24
Rx
  L3:
    0 packets, 0 bytes
Tx
  L3:
    0 packets, 0 bytes

*****
Iteration : 6 Command : show interface vlan 46
*****

Interface vlan46 is up
Admin state is up
Description:
Hardware: Ethernet, MAC Address: 88:3a:30:a5:52:80
IPv4 address 10.x.46.4/24
Rx
  L3:
    4 packets, 312 bytes
```

Tx

L3:

4 packets, 312 bytes

Test link failover

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

```
Sw2(config)# interface vlan 46
Sw2(config-if-vlan)# shutdown
```

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.

```
C:\Users\student> ping 10.x.12.1

Pinging 10.x.12.1 with 32 bytes of data:
Request timed out.
Request timed out.
Request timed out.
Request timed out.

Ping statistics for 10.x.12.1:
    Packets: Sent = 4, Received = 0, Lost = 4 (100% loss)
```

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?

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.

```
Sw2(config-if-vlan)# no shutdown
```

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

```
(config)# bfd
(config)# interface vlan 46
```

```
(config-if-vlan)# ip ospf bfd
(config-if-vlan)# interface vlan 47
(config-if-vlan)# ip ospf bfd
(config-if-vlan)# exit
```

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

```
Sw1# show bfd

Admin status: enabled
Echo source IP: <none>
Statistics:
Total number of control packets transmitted: 218
Total number of control packets received: 82
Total number of control packets dropped: 0
Session Interface VRF          Source IP          Destination IP
Echo      State      Protocol
-----
1         vlan47     default  10.x.47.4          10.x.47.5
enabled  up              ospf
2         vlan46     default  10.x.46.4          10.x.46.5
enabled  up              ospf
```

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

```
Sw1# show bfd session 1

BFD session information - Session 1
VRF: default
Min Tx interval (msec): 3000
Min Rx interval (msec): 3000
Min echo Rx interval (msec): 500
Detect multiplier: 5
Protocol(s): ospf
Local discriminator: 13211
Remote discriminator: 8519
Echo: enabled
Local diagnostic: no_diagnostic
Remote diagnostic: no_diagnostic
State flaps: 0
Interface Source IP          Destination IP
State      Pkt Rx   Pkt Tx   Pkt drop
-----
vlan47     10.x.47.4          10.x.47.5          up
49         118      0
```

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.

```
(config)# bfd detect-multiplier 3
(config)# bfd min-transmit-interval 1000
(config)# bfd min-receive-interval 1000
```

25. On Sw1, start the statistics repeat again

```
Sw1# clear interface statistics
Sw1# show int vlan 46
Sw1# show int vlan 47
Sw1# repeat id 1,2 delay 10
```

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.

```
ping 10.x.12.1 -t
```

27. On Sw2, disable SVI 46.

```
Sw2(config)# interface vlan 46
Sw2(config-if)# shutdown
```

28. On PC4, review the ping output.

```
C:\Users\student> ping 10.x.12.1 -t

Pinging 10.x.12.1 with 32 bytes of data:
Reply from 10.x.12.1: bytes=32 time<1ms TTL=62
Reply from 10.x.12.1: bytes=32 time<1ms TTL=62
Reply from 10.x.12.1: bytes=32 time<1ms TTL=62
Reply from 10.x.12.1: bytes=32 time<1ms TTL=62
Reply from 10.x.12.1: bytes=32 time<1ms TTL=62
Reply from 10.x.12.1: bytes=32 time<1ms TTL=62
Request timed out.
Reply from 10.x.12.1: bytes=32 time<1ms TTL=62
Reply from 10.x.12.1: bytes=32 time<1ms TTL=62

Ping statistics for 10.x.12.1:
    Packets: Sent = 9, Received = 8, Lost = 1 (11% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 0ms, Maximum = 0ms, Average = 0ms
```

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.

```
*****
Iteration : 11 Command : show interface vlan 47
*****

Interface vlan47 is up
Admin state is up
Description:
Hardware: Ethernet, MAC Address: 88:3a:30:a5:52:80
IPv4 address 10.x.47.4/24
Rx
  L3:
    0 packets, 0 bytes
Tx
  L3:
    0 packets, 0 bytes

*****
Iteration : 11 Command : show interface vlan 46
*****

Interface vlan46 is up
Admin state is up
Description:
Hardware: Ethernet, MAC Address: 88:3a:30:a5:52:80
IPv4 address 10.x.46.4/24
Rx
  L3:
    4 packets, 312 bytes
Tx
  L3:
    4 packets, 312 bytes
```

30. On Sw2 enable SVI 46 again.

```
Sw2(config)# interface vlan 46
Sw2(config-if)# no shutdown
```

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.

```
Sw2(config)# class ip voice
Sw2(config-class-ip)# 10 match ip any 10.x.12.0/24 count
Sw2(config-class-ip)# exit
```

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

```
Sw2(config)# pbr-action-list nexthop47
Sw2(config-pbr-action-list-nexthop47)# nexthop 10.x.47.4
Sw2(config-pbr-action-list-nexthop47)# exit
```

33. Define a new policy 'branch' that applies the PBR action 'nexthop47' to the class 'voice'.

```
Sw2(config)# policy branch
Sw2(config-policy)# 10 class ip voice action pbr nexthop47
Sw2(config-policy)# exit
```

34. Apply the new policy to the VLAN 15 interface.

```
Sw2(config)# interface vlan 15
Sw2(config-if-vlan)# apply policy branch routed-in
Sw2(config-if-vlan)# exit
```

Verification

35. On Sw2, review the PBR configuration

```
Sw2# show pbr interface vlan15
```

VRP	Port	Policy	Class	PBR	Sequence	Type	Nexthop
branch	vlan15	branch	voice	nexthop47	10	nexthop	10.x.47.4 (active)

36. On Sw2, review the policy hitcounts

```
Sw2# show policy hitcounts branch
Statistics for Policy branch:
Interface vlan15* (routed-in):
Hit Count Configuration
10 class ip voice action pbr nexthop47
```

```
387 10 match any any 10.x.12.0/255.255.255.0 count
```

* policy statistics are shared among each context type (interface, VLAN).
For routed ingress, they are only shared within the same VRF.
Use 'policy NAME copy' to create a new policy for separate statistics.

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!

```
# write memory
```

```
Copying configuration: [Success]
```

```
# copy running-config checkpoint asts-lab05-<your-name>
```

```
Copying configuration: [Success]
```

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.

```
Agg1(config)# interface 1/1/7
Agg1(config-if)# shutdown
Agg1(config-if)# exit
```

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.

```
Agg1(config)# router bgp 64500
Agg1(config-bgp)# neighbor 10.x.0.3 remote-as 64500
Agg1(config-bgp)# address-family ipv4 unicast
Agg1(config-bgp-ipv4-uc)# neighbor 10.x.0.3 activate
Agg1(config-bgp-ipv4-uc)# end
```

3. On Agg2, configure BGP AS 64500. Add Agg1 Loopback IP as iBGP peer, enable the IPv4 address family.

```
Agg2(config)# router bgp 64500
Agg2(config-bgp)# neighbor 10.x.0.2 remote-as 64500
Agg2(config-bgp)# address-family ipv4 unicast
Agg2(config-bgp-ipv4-uc)# neighbor 10.x.0.2 activate
Agg2(config-bgp-ipv4-uc)# end
```

4. On Agg1, review the current BGP peers.

```
Agg1# show bgp all summary
VRF : default
BGP Summary
-----
Local AS           : 64500           BGP Router Identifier : 10.x.0.2
```



```

Peers                : 1                Log Neighbor Changes : No
Cfg. Hold Time       : 180              Cfg. Keep Alive      : 60

Address-family : IPv4 Unicast
-----
Neighbor      Remote-AS MsgRcvd MsgSent  Up/Down Time State      AdminStatus
10.x.0.3      64500      0      1      00h:00m:00s Active      Up
----- Output Omitted -----

```

Q: What do you observe?

A: There are currently no peers active.

5. On Agg1, debug bgp to see what is happening.

```

Agg1# clear debug buffer
Agg1# debug bgp all

```

6. On Agg1, review the debug buffer. This may take a few moments based on the default BGP timers.

```

Agg1# show debug buffer

-----
show debug buffer
-----
----- Output Omitted -----
2020-10-07:12:47:15.098068|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|An OPEN message is being sent
to a neighbor.
2020-10-07:12:47:15.098075|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|NM entity index      =
268763136
2020-10-07:12:47:15.098081|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|Configured local address =
(none)
2020-10-07:12:47:15.098088|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|Configured local port   = 0
2020-10-07:12:47:15.098094|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|Selected local address  =
10.x.40.2
2020-10-07:12:47:15.098104|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|Selected local port   =
41632
2020-10-07:12:47:15.098113|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|Remote address       =
10.x.0.3
----- Output Omitted -----
2020-10-07:12:47:15.098371|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|A NOTIFICATION message has been
received from a neighbor.
2020-10-07:12:47:15.098378|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|NM entity index      =
268763136
2020-10-07:12:47:15.098427|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|Error code           = Cease
(6)
2020-10-07:12:47:15.098436|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|Error subcode        =
Connection Rejected (5)VRF Name = default.
2020-10-07:12:47:15.098557|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|A connection's FSM state has
deteriorated.
----- Output Omitted -----

```

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.

```
Agg1(config)# router bgp 64500
Agg1(config-bgp)# neighbor 10.x.0.3 update-source loopback 0
Agg1(config-bgp)# exit
```

8. On Agg2, configure the correct source interface.

```
Agg2(config)# router bgp 64500
Agg2(config-bgp)# neighbor 10.x.0.2 update-source loopback 0
Agg2(config-bgp)# exit
```

9. On Agg1, review the BGP IPv4 unicast summary. The peer should now be established.

```
Agg1# show bgp ipv4 unicast summary
VRF : default
BGP Summary
-----
Local AS           : 64500           BGP Router Identifier : 10.x.0.2
Peers              : 1               Log Neighbor Changes  : No
Cfg. Hold Time     : 180            Cfg. Keep Alive       : 60

Neighbor          Remote-AS MsgRcvd MsgSent  Up/Down Time State      AdminStatus
10.x.0.3          64500     9          12      00h:01m:03s Established Up
```

10. On Agg1, review the debug log for the 'bgp' module.

```
Agg1# show debug buffer module bgp
----- Output Omitted -----
2020-10-07:13:02:54.869693|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|BGP 268763136 established a
session with peer 10.x.0.3.
2020-10-07:13:02:54.869700|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|Configured local address and
port: 10.x.0.2:0
2020-10-07:13:02:54.869706|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|Connection local address and
port: 10.x.0.2:179
2020-10-07:13:02:54.869712|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|Configured remote address and
port: 10.x.0.3:0
----- Output Omitted -----
```

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.

```
Agg1(config)# router bgp 64500
Agg1(config-bgp)# neighbor 10.x.0.4 remote-as 64500
Agg1(config-bgp)# neighbor 10.x.0.4 update-source loopback 0
Agg1(config-bgp)# address-family ipv4 unicast
Agg1(config-bgp-ipv4-uc)# neighbor 10.x.0.4 activate
Agg1(config-bgp-ipv4-uc)# end
```

12. On Agg2, configure iBGP peering to Sw1 using the loopback interface.

```
Agg2(config)# router bgp 64500
Agg2(config-bgp)# neighbor 10.x.0.4 remote-as 64500
Agg2(config-bgp)# neighbor 10.x.0.4 update-source loopback 0
Agg2(config-bgp)# address-family ipv4 unicast
Agg2(config-bgp-ipv4-uc)# neighbor 10.x.0.4 activate
```

13. On Sw1, configure iBGP peering to Agg1 and Agg2 using the loopback interface.

```
Sw1# conf t
Sw1(config)# router bgp 64500
Sw1(config-bgp)# neighbor 10.x.0.2 remote-as 64500
Sw1(config-bgp)# neighbor 10.x.0.2 update-source loopback 0
Sw1(config-bgp)# neighbor 10.x.0.3 remote-as 64500
Sw1(config-bgp)# neighbor 10.x.0.3 update-source loopback 0
Sw1(config-bgp)# address-family ipv4 unicast
Sw1(config-bgp-ipv4-uc)# neighbor 10.x.0.2 activate
Sw1(config-bgp-ipv4-uc)# neighbor 10.x.0.3 activate
Sw1(config-bgp-ipv4-uc)# end
```

14. On Agg1, Agg2 and Sw1: verify the BGP peering is established between with the other 2 switches.

```
# show bgp ipv4 unicast summary
VRF : default
BGP Summary
-----
Local AS           : 64500           BGP Router Identifier : 10.x.0.2
Peers              : 2               Log Neighbor Changes  : No
Cfg. Hold Time     : 180            Cfg. Keep Alive       : 60

Neighbor           Remote-AS MsgRcvd MsgSent  Up/Down Time State      AdminStatus
10.x.0.3           64500    20      22     00h:10m:19s Established Up
10.x.0.4           64500    4       5     00h:01m:43s Established Up
```

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).

```
Agg1# conf t
Agg1(config)# router bgp 64500
Agg1(config-bgp)# neighbor 10.255.101.11 remote-as 64520
Agg1(config-bgp)# address-family ipv4 unicast
Agg1(config-bgp-ipv4-uc)# neighbor 10.255.101.11 activate
Agg1(config-bgp-ipv4-uc)# end
```

16. On Agg1, review the BGP peer status

```
Agg1# show bgp ipv4 unicast summary
VRF : default
BGP Summary
-----
Local AS           : 64500           BGP Router Identifier : 10.x.0.2
Peers              : 3               Log Neighbor Changes  : No
Cfg. Hold Time     : 180            Cfg. Keep Alive       : 60

Neighbor      Remote-AS  MsgRcvd  MsgSent  Up/Down Time  State      AdminStatus
10.x.0.3      64500      24       26       00h:14m:01s   Established Up
10.x.0.4      64500      8        9       00h:05m:26s   Established Up
10.255.101.11 64520      0        0       00h:00m:00s   Idle       Up
```

17. Clear the debug buffer and enable port 1/1/7.

```
Agg1# clear debug buffer
Agg1# conf t
Agg1(config)# interface 1/1/7
Agg1(config-if)# no shutdown
Agg1(config-if)# end
```

18. Review the debug buffer.

```
Agg1# show debug buffer module bgp
-----
show debug buffer
-----
----- Output Omitted -----
2020-10-07:13:18:06.180524|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|A valid OPEN message has been
received from a neighbor.
----- Output Omitted -----
2020-10-07:13:18:06.180584|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|Remote address      =
10.255.101.11
2020-10-07:13:18:06.180594|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|Remote port        = 0
2020-10-07:13:18:06.180604|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|Scope ID          = 0
```

```

2020-10-07:13:18:06.180614|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|Incoming connection?      =
False
2020-10-07:13:18:06.180624|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|Remote AS number        =
64511
----- Output Omitted -----
2020-10-07:13:18:06.180926|hpe-routing|LOG_ERR|AMM|-|BGP|BGP|A NOTIFICATION message is being
sent to a neighbor due to an unexpected
----- Output Omitted -----
2020-10-07:13:18:06.181026|hpe-routing|LOG_ERR|AMM|-|BGP|BGP|Error code                = OPEN
Message Error (2)
2020-10-07:13:18:06.181036|hpe-routing|LOG_ERR|AMM|-|BGP|BGP|Error subcode            = Bad Peer
AS (2)VRF Name = default.

```

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.

```

Agg1# conf t
Agg1(config)# router bgp 64500
Agg1(config-bgp)# no neighbor 10.255.101.11
Agg1(config-bgp)# neighbor 10.255.101.11 remote-as 64511
Agg1(config-bgp)# address-family ipv4 unicast
Agg1(config-bgp)# neighbor 10.255.101.11 activate
Agg1(config-bgp)# end

```

20. Review the peering status, it should now be established

```

Agg1# show bgp ipv4 unicast summary
VRF : default
BGP Summary
-----
Local AS           : 64500          BGP Router Identifier : 10.x.0.2
Peers              : 3              Log Neighbor Changes  : No
Cfg. Hold Time     : 180           Cfg. Keep Alive       : 60

Neighbor    Remote-AS  MsgRcvd  MsgSent  Up/Down Time  State      AdminStatus
10.x.0.3    64500        51       56       00h:36m:40s  Established Up
10.x.0.4    64500        35       39       00h:28m:04s  Established Up
10.255.101.11 64511        7        4       00h:00m:36s  Established Up

```

21. On Agg2, configure eBGP peering with routerB using the physical interface address.

```

Agg2(config)# router bgp 64500

```

```
Agg2(config-bgp)# neighbor 10.255.102.12 remote-as 64512
Agg2(config-bgp)# address-family ipv4 unicast
Agg2(config-bgp-ipv4-uc)# neighbor 10.255.102.12 activate
Agg2(config-bgp-ipv4-uc)# end
```

22. On Agg2, verify the BGP peering is established with the eBGP peer.

```
Agg2# show bgp ipv4 unicast summary
```

VRF : default
BGP Summary

Local AS	: 64500	BGP Router Identifier	: 10.x.0.3
Peers	: 3	Log Neighbor Changes	: No
Cfg. Hold Time	: 180	Cfg. Keep Alive	: 60

Neighbor	Remote-AS	MsgRcvd	MsgSent	Up/Down Time	State	AdminStatus
10.x.0.2	64500	52	54	00h:39m:55s	Established	Up
10.x.0.4	64500	37	42	00h:30m:02s	Established	Up
10.255.102.12	64512	8	4	00h:01m:09s	Established	Up

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.

```
Agg1# show bgp ipv4 unicast
```

Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default
Local Router-ID 10.x.0.2

Network	Nexthop	Metric	LocPrf	Weight	Path
* i 192.0.2.0/24	10.255.102.12	0	100	0	64512 64511 i
*>e 192.0.2.0/24	10.255.101.11	0	100	0	64511 i
* i 198.19.100.0/24	10.255.102.12	0	100	0	64512 64513 i
*>e 198.19.100.0/24	10.255.101.11	0	100	0	64511 64513 i
* i 198.19.101.0/24	10.255.102.12	0	100	0	64512 64511 64513 i
*>e 198.19.101.0/24	10.255.101.11	0	100	0	64511 64513 i
* i 198.19.102.0/24	10.255.102.12	0	100	0	64512 64513 i
*>e 198.19.102.0/24	10.255.101.11	0	100	0	64511 64512 64513 i
* i 198.51.100.0/24	10.255.102.12	0	100	0	64512 i
*>e 198.51.100.0/24	10.255.101.11	0	100	0	64511 64512 i

Total number of entries 10

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.

```
Agg1# show bgp ipv4 unicast 192.0.2.0/24
VRF : default
BGP Local AS 64500      BGP Router-id 10.x.0.2

Network      : 192.0.2.0/24      Nexthop      : 10.255.102.12
Peer         : 10.x.0.3          Origin       : IGP
Metric       : 0                Local Pref   : 100
Weight       : 0                Calc. Local Pref : 100
Best         : No               Valid        : Yes
Type         : internal         Stale        : No
Originator ID : 0.0.0.0
Aggregator ID :
Aggregator AS :
Atomic Aggregate :
RFD Flaps     : 0              RFD Penalty  : 0

AS-Path       : 64512 64511

Cluster List  :
Communities   :
Ext-Communities :
Network      : 192.0.2.0/24      Nexthop      : 10.255.101.11
Peer         : 10.255.101.11    Origin       : IGP
Metric       : 0                Local Pref   : 100
Weight       : 0                Calc. Local Pref : 100
Best         : Yes              Valid        : Yes
Type         : external         Stale        : No
Originator ID : 0.0.0.0
Aggregator ID :
Aggregator AS :
Atomic Aggregate :
RFD Flaps     : 0              RFD Penalty  : 0

AS-Path       : 64511

Cluster List  :
Communities   :
Ext-Communities :
```

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

```
Agg1# show ip route 10.255.102.12
No ipv4 routes configured
```

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

```
Agg1# show bgp ipv4 unicast 192.0.2.0/24
```

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.

```
Agg2# conf t
Agg2(config)# router bgp 64500
Agg2(config-bgp)# address-family ipv4 unicast
Agg2(config-bgp-ipv4-uc)# neighbor 10.x.0.2 next-hop-self
Agg2(config-bgp-ipv4-uc)# end
```

28. On Agg1, review the details of the 192.0.2.0/24 route.

```
Agg1# show bgp ipv4 unicast 192.0.2.0/24

VRF : default
BGP Local AS 64500      BGP Router-id 10.x.0.2

  Network          : 192.0.2.0/24      Nexthop          : 10.255.102.12
  Peer             : 10.x.0.3          Origin           : IGP
----- Output Omitted -----
```

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.

```
Agg2# debug bgp all
```

30. BGP Hard Reset: On Agg1, you can now perform a hard reset of the BGP session with Agg2 (10.x.0.3)

```
Agg1# clear bgp 10.x.0.3
```

31. On Agg2, review the debug buffer.

```
2020-10-07:14:10:57.075873|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|A connection has left
Established state.
2020-10-07:14:10:57.075884|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|Local address: 10.x.0.3
2020-10-07:14:10:57.075894|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|Local port: 0
2020-10-07:14:10:57.075904|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|Remote address: 10.x.0.2
2020-10-07:14:10:57.075914|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|Remote port: 0
2020-10-07:14:10:57.075924|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|Scope ID: 0
2020-10-07:14:10:57.075934|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|Neg hold time: 180
2020-10-07:14:10:57.075944|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|Passive:? False
2020-10-07:14:10:57.075954|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|Entity index: 268763136VRF
Name = default.
2020-10-07:14:10:57.076330|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|A BGP peer session has failed
and may restart.
```

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.

```
Agg1# show bgp ipv4 unicast

Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
               i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default
Local Router-ID 10.x.0.2

   Network          Nexthop          Metric    LocPrf    Weight Path
* i 192.0.2.0/24    10.x.0.3          0         100       0      64512 64511 i
*>e 192.0.2.0/24    10.255.101.11     0         100       0      64511 i
----- Output Omitted -----
```

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

```
Agg1# conf t
Agg1(config)# router bgp 64500
Agg1(config-bgp)# address-family ipv4 unicast
Agg1(config-bgp-ipv4-uc)# neighbor 10.x.0.3 next-hop-self
Agg1(config-bgp-ipv4-uc)# neighbor 10.x.0.4 next-hop-self
Agg1(config-bgp-ipv4-uc)# end
```

34. Complete the next-hop self configuration on Agg2

```
Agg2# conf t
Agg2(config)# router bgp 64500
Agg2(config-bgp)# address-family ipv4 unicast
Agg2(config-bgp-ipv4-uc)# neighbor 10.x.0.4 next-hop-self
Agg2(config-bgp-ipv4-uc)# end
```

35. Complete the next-hop self configuration on Sw1

```
Sw1# conf t
Sw1(config)# router bgp 64500
Sw1(config-bgp)# address-family ipv4 unicast
Sw1(config-bgp-ipv4-uc)# neighbor 10.x.0.2 next-hop-self
Sw1(config-bgp-ipv4-uc)# neighbor 10.x.0.3 next-hop-self
Sw1(config-bgp-ipv4-uc)# end
```

36. On Agg1 and Agg2, reset all sessions to make the next-hop self effective.

```
# clear bgp *
```

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.

```
Sw1# conf t
Sw1(config)# vlan 48
Sw1(config-vlan-48)# int 1/1/27
Sw1(config-if)# vlan trunk allowed 48
Sw1(config-if)# exit
```

38. On Sw1, add SVI 48 with IP address 10.x.48.4/24.

```
Sw1(config-if)# interface vlan48
Sw1(config-if-vlan)# ip address 10.x.48.4/24
Sw1(config-if-vlan)# exit
```

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.

```
Sw1# conf t
Sw1(config)# router bgp 64500
Sw1(config-bgp)# neighbor 10.x.48.5 remote-as 64600
Sw1(config-bgp)# address-family ipv4 unicast
Sw1(config-bgp-ipv4-uc)# neighbor 10.x.48.5 activate
Sw1(config-bgp-ipv4-uc)# end
```

41. On Sw1, verify the peer is in 'established' state.

```
Sw1# show bgp ipv4 unicast summary
VRF : default
BGP Summary
-----
Local AS           : 64500          BGP Router Identifier : 10.x.0.4
Peers              : 3              Log Neighbor Changes  : No
Cfg. Hold Time     : 180           Cfg. Keep Alive       : 60

Neighbor      Remote-AS  MsgRcvd  MsgSent  Up/Down Time State      AdminStatus
10.x.0.2      64500      70       67      00h:56m:09s Established Up
10.x.0.3      64500      70       70      00h:54m:11s Established Up
10.x.48.5     64600      5        8      00h:00m:46s Established Up
```

42. On Sw1, verify the BGP IPv4 routing table. You should see entries from the AS 64600.

```
Sw1# show bgp ipv4 unicast
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
               i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default
Local Router-ID 10.x.0.4

   Network          Nexthop        Metric    LocPrf    Weight Path
*i>i 192.0.2.0/24    10.x.0.2       0         100       0      64511 i
*>e 198.18.14.0/24  10.x.48.5      0         100       0      64600 i
*>e 198.18.15.0/24  10.x.48.5      0         100       0      64600 ?
*i>i 198.19.100.0/24 10.x.0.2       0         100       0      64511 64513 i
* i 198.19.100.0/24  10.x.0.3       0         100       0      64512 64513 i
*i>i 198.19.101.0/24 10.x.0.2       0         100       0      64511 64513 i
*i>i 198.19.102.0/24 10.x.0.3       0         100       0      64512 64513 i
*i>i 198.51.100.0/24 10.x.0.3       0         100       0      64512 i
Total number of entries 8
```

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.

```
Sw1# show ip route 0.0.0.0
No ipv4 routes configured
```

44. On both Agg1 and Agg2, ensure that the default route is always advertised into OSPF.

```
# conf t
(config)# router ospf 1
(config-ospf-1)# default-information originate always
(config-ospf-1)# end
```

45. On Sw1, verify in the OSPF LSDB that the 2 default routes are in the database and in the IP routing table.

```
Sw1# show ip ospf lsdb
----- Output Omitted -----
AS External Link State Advertisements
-----
LSID          ADV Router    Age           Seq#           Checksum
-----
0.0.0.0       10.x.0.2      219           0x80000001    0x00002693
0.0.0.0       10.x.0.3      144           0x80000001    0x00002098
```

```
Sw1# show ip route 0.0.0.0
Displaying ipv4 routes selected for forwarding

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

0.0.0.0/0, vrf default
    via 10.x.41.2, [110/1], ospf
    via 10.x.42.3, [110/1], ospf
```

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

```
Agg1# conf t
Agg1(config)# interface vlan 12
Agg1(config-if-vlan)# ip address 100.100.1.2/26 secondary
Agg1(config-if-vlan)# exit
```

47. On Agg2

```
Agg2# conf t
Agg2(config)# interface vlan12
Agg2(config-if-vlan)# ip address 100.100.1.3/26 secondary
Agg2(config-if-vlan)# exit
```

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.

```
Agg1(config)# router bgp 64500
Agg1(config-bgp)# address-family ipv4 unicast
Agg1(config-bgp-ipv4-uc)# network 100.100.1.0/24
Agg1(config-bgp-ipv4-uc)# end
```

49. On Agg1, review the BGP IPv4 routing table

```
Agg1# show bgp ipv4 unicast 100.100.1.0/24
```

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.

```
Agg1# conf t
Agg1(config)# ip route 100.100.1.0/24 blackhole
```

51. Verify the advertisement: On Agg2, review the BGP routing table

```
Agg2# show bgp ipv4 unicast

Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
               i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete
```

```
VRF : default
Local Router-ID 10.x.0.3
```

Network	Nexthop	Metric	LocPrf	Weight	Path
*>i 100.100.1.0/24	10.x.0.2	0	100	0	i

----- Output Omitted -----

52. Open a session to routerA. Verify the BGP routing table.

```
routerA# show bgp ipv4 unicast
```

Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

```
VRF : default
Local Router-ID 192.0.2.50
```

Network	Nexthop	Metric	LocPrf	Weight	Path
*>e 100.100.1.0/24	10.255.101.2	0	100	0	64500 i
* e 100.100.1.0/24	10.255.103.12	0	100	0	64512 64500 i

----- Output Omitted -----

53. Open a session to routerB. Verify the BGP routing table.

```
routerB# show bgp ipv4 unicast
```

Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

```
VRF : default
Local Router-ID 198.51.100.50
```

Network	Nexthop	Metric	LocPrf	Weight	Path
*>e 100.100.1.0/24	10.255.102.3	0	100	0	64500 i
* e 100.100.1.0/24	10.255.103.11	0	100	0	64511 64500 i

----- Output Omitted -----

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.

```
Agg2(config)# ip prefix-list connected_to_bgp seq 10 permit 100.100.1.0/26
```

55. On Agg2, define a route-map 'connected_to_bgp' that permits and matches the IP prefix-list 'connected_to_bgp'.

```
Agg2(config)# route-map connected_to_bgp permit seq 10
Agg2(config-route-map-connected_to_bgp-10)# match ip address prefix-list connected_to_bgp
Agg2(config-route-map-connected_to_bgp-10)# exit
```

56. On Agg2, redistribute the connected routes using the route-map 'connected_to_bgp' under the BGP IPv4 address-family.

```
Agg2(config)# router bgp 64500
Agg2(config-bgp)# address-family ipv4 unicast
Agg2(config-bgp-ipv4-uc)# redistribute connected route-map connected_to_bgp
Agg2(config-bgp-ipv4-uc)# end
```

57. On Agg2, review the BGP IPv4 routing table and verify that the /26 prefix has been added to the local BGP table.

```
Agg2# show bgp ipv4 unicast

Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
               i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default
Local Router-ID 10.x.0.3

  Network                Nexthop          Metric      LocPrf      Weight Path
```



```
*>i 100.100.1.0/24    10.x.0.2    0    100    0    i
*> 100.100.1.0/26    0.0.0.0    0    100    0    ?
----- Output Omitted -----
```

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'.

```
Agg2(config)# route-map connected_to_bgp permit seq 10
Agg2(config-route-map-connected_to_bgp-10)# set origin igp
Agg2(config-route-map-connected_to_bgp-10)# exit
```

59. On Agg2, review the route-map.

```
Agg2# show route-map
Route-map: connected_to_bgp
Seq 10, permit,
Match :
  ip prefix-list      : connected_to_bgp
Set :
  origin              : igp
```

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

```
Agg2# show bgp ipv4 unicast neighbors 10.255.102.12 advertised-routes
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
               i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default
Local Router-ID 10.x.0.3

   Network          Nexthop        Metric    LocPrf    Weight Path
*>e 100.100.1.0/24    10.255.102.3    0         0         0      64500 i
*>e 100.100.1.0/26    10.255.102.3    0         0         0      64500 i
----- Output Omitted -----
```

61. On Agg2, now configure the aggregate route under the IPv4 address-family.

```
Agg2(config)# router bgp 64500
Agg2(config-bgp)# address-family ipv4 unicast
Agg2(config-bgp-ipv4-uc)# aggregate-address 100.100.1.0/24
Agg2(config-bgp-ipv4-uc)# exit
```

62. On Agg2, review the local BGP table and the advertised routes to routerB.

```
Agg2# show bgp ipv4 unicast

Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
               i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default
Local Router-ID 10.x.0.3

   Network          Nexthop        Metric    LocPrf    Weight Path
*>i 100.100.1.0/24   10.x.0.2         0         100       0      i
*> 100.100.1.0/26    0.0.0.0          0         100       0      i
----- Output Omitted -----
```

```
Agg2# show bgp ipv4 unicast neighbors 10.255.102.12 advertised-routes

Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
               i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default
Local Router-ID 10.x.0.3

   Network          Nexthop        Metric    LocPrf    Weight Path
*>e 100.100.1.0/24   10.255.102.3     0         0         0      64500 i
*>e 100.100.1.0/26   10.255.102.3     0         0         0      64500 i
----- Output Omitted -----
```

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.

```
Agg2(config)# router bgp 64500
Agg2(config-bgp)# address-family ipv4 unicast
Agg2(config-bgp-ipv4-uc)# aggregate-address 100.100.1.0/24 summary-only
```

64. On Agg2, review the local BGP routing table.

```
Agg2# show bgp ipv4 unicast

Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
               i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default
```

Local Router-ID 10.x.0.3

Network	Nexthop	Metric	LocPrf	Weight	Path
*>i 100.100.1.0/24	10.x.0.2	0	100	0	i
s> 100.100.1.0/26	0.0.0.0	0	100	0	i

----- Output Omitted -----

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

Agg1# show bgp ipv4 unicast neighbors 10.255.101.11 advertised-routes

Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default

Local Router-ID 10.x.0.2

Network	Nexthop	Metric	LocPrf	Weight	Path
*>e 100.100.1.0/24	10.255.101.2	0	0	0	64500 i
*>e 198.18.14.0/24	10.255.101.2	0	0	0	64500 64600 i

----- Output Omitted -----

66. On Agg2, repeat this check

Agg2# show bgp ipv4 unicast neighbors 10.255.102.12 advertised-routes

Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default

Local Router-ID 10.x.0.3

Network	Nexthop	Metric	LocPrf	Weight	Path
*>e 100.100.1.0/24	10.255.102.3	0	0	0	64500 i
*>e 192.0.2.0/24	10.255.102.3	0	0	0	64500 64511 i

----- Output Omitted -----

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'.

```
Agg1(config)# router bgp 64500
Agg1(config-bgp)# neighbor internal peer-group
```

2. Configure the peer-group with remote-AS 64500.

```
Agg1(config-bgp)# neighbor internal remote-as 64500
```

3. Configure the peer-group with source interface Loopback0.

```
Agg1(config-bgp)# neighbor internal update-source loopback 0
```

4. Configure the peer-group with next-hop-self under the IPv4 unicast family.

```
Agg1(config-bgp)# address-family ipv4 unicast
Agg1(config-bgp-ipv4-uc)# neighbor internal next-hop-self
Agg1(config-bgp-ipv4-uc)# exit
```

Migrate the existing peer to the peer-group

5. Remove the existing Agg2 peer (10.x.0.3).

```
Agg1(config-bgp)# no neighbor 10.x.0.3
```

6. Add Agg2 again with the peer-group command.

```
Agg1(config-bgp)# neighbor 10.x.0.3 peer-group internal
```

7. Make sure to activate the neighbor under the IPv4 address family.

```
Agg1(config-bgp)# address-family ipv4 unicast
Agg1(config-bgp-ipv4-uc)# neighbor 10.x.0.3 activate
Agg1(config-bgp-ipv4-uc)# exit
```

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).

```
Agg1(config-bgp)# no neighbor 10.x.0.4
Agg1(config-bgp)# neighbor 10.x.0.4 peer-group internal
Agg1(config-bgp)# address-family ipv4 unicast
Agg1(config-bgp-ipv4-uc)# neighbor 10.x.0.4 activate
Agg1(config-bgp-ipv4-uc)# exit
```

9. Verify that the BGP neighbors on Agg1 are back in the established state.

```
Agg1# show bgp ipv4 unicast summary
VRF : default
BGP Summary
-----
Local AS           : 64500      BGP Router Identifier : 10.x.0.2
Peers              : 3          Log Neighbor Changes  : No
Cfg. Hold Time     : 180       Cfg. Keep Alive       : 60

Neighbor      Remote-AS  MsgRcvd  MsgSent  Up/Down Time State      AdminStatus
10.x.0.3      64500      11       15       00h:02m:33s Established Up
10.x.0.4      64500      6        7       00h:01m:14s Established Up
10.255.101.11 64511      179      186      02h:26m:59s Established Up
```

10. Review the current BGP configuration using

```
Agg1# show running bgp
router bgp 64500
  neighbor internal peer-group
  neighbor internal remote-as 64500
  neighbor internal update-source loopback 0
  neighbor 10.x.0.3 peer-group internal
  neighbor 10.x.0.4 peer-group internal
  neighbor 10.255.101.11 remote-as 64511
  address-family ipv4 unicast
    neighbor internal next-hop-self
    neighbor 10.x.0.3 activate
    neighbor 10.x.0.4 activate
    neighbor 10.255.101.11 activate
    network 100.100.1.0/24
  exit-address-family
```

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 is 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.

```
Agg1# show bgp ipv4 unicast neighbors 10.255.101.11 advertised-routes
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
               i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default
Local Router-ID 10.x.0.2
```

Network	Nexthop	Metric	LocPrf	Weight	Path
*>e 100.100.1.0/24	10.255.101.2	0	0	0	64500 i
*>e 198.18.14.0/24	10.255.101.2	0	0	0	64500 64600 i
*>e 198.18.15.0/24	10.255.101.2	0	0	0	64500 64600 ?
*>e 198.19.102.0/24	10.255.101.2	0	0	0	64500 64512 64513 i
*>e 198.51.100.0/24	10.255.101.2	0	0	0	64500 64512 i

Total number of entries 5

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.

```
Agg2# show bgp ipv4 unicast neighbors 10.255.102.12 advertised-routes
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
               i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default
Local Router-ID 10.x.0.3
```

Network	Nexthop	Metric	LocPrf	Weight	Path
*>e 100.100.1.0/24	10.255.102.3	0	0	0	64500 i
*>e 192.0.2.0/24	10.255.102.3	0	0	0	64500 64511 i

```
*>e 198.18.14.0/24      10.255.102.3      0          0          0          64500 64600 i
*>e 198.18.15.0/24      10.255.102.3      0          0          0          64500 64600 ?
*>e 198.19.101.0/24     10.255.102.3      0          0          0          64500 64511 64513 i
Total number of entries 5
```

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.

- On Agg1, define an AS-path list name 'bgp-local' that only contains an empty AS path regular expression.

```
Agg1(config)# ip aspath-list bgp-local permit ^$
```

- On Agg1, define a route-map named 'bgp-ebgp-out' that permit routes matching this AS-path list.

```
Agg1(config)# route-map bgp-ebgp-out permit
Agg1(config-route-map-bgp-ebgp-out-10)# match aspath bgp-local
Agg1(config-route-map-bgp-ebgp-out-10)# exit
```

- On Agg1, apply the route-map for outbound traffic on the routerA eBGP peer for IPv4 unicast routes.

```
Agg1(config)# router bgp 64500
Agg1(config-bgp)# address-family ipv4 unicast
Agg1(config-bgp-ipv4-uc)# neighbor 10.255.101.11 route-map bgp-ebgp-out out
Agg1(config-bgp-ipv4-uc)# exit
Agg1(config-bgp)# exit
```

- On Agg1, now review the IPv4 routes that are advertised to routerA.

```
Agg1# show bgp ipv4 unicast neighbors 10.255.101.11 advertised-routes
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
               i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default
Local Router-ID 10.x.0.2

   Network          Nexthop        Metric      LocPrf      Weight Path
*>e 100.100.1.0/24    10.255.101.2    0           0           0        64500 i
*>e 198.18.14.0/24    10.255.101.2    0           0           0        64500 64600 i
*>e 198.18.15.0/24    10.255.101.2    0           0           0        64500 64600 ?
*>e 198.19.102.0/24   10.255.101.2    0           0           0        64500 64512 64513 i
*>e 198.51.100.0/24   10.255.101.2    0           0           0        64500 64512 i
Total number of entries 5
```

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.

```
Agg1(config)# do clear bgp 10.255.101.11
```

8. Review the advertised routes to routerA.

```
Agg1# show bgp ipv4 unicast neighbors 10.255.101.11 advertised-routes
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
               i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default
Local Router-ID 10.x.0.2
```

Network	Nexthop	Metric	LocPrf	Weight	Path
*>e 100.100.1.0/24	10.255.101.2	0	0	0	64500 i

Total number of entries 1

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.

```
Agg2# conf t
Agg2(config)# ip aspath-list bgp-local permit ^$
Agg2(config)# route-map bgp-ebgp-out permit
Agg2(config-route-map-bgp-ebgp-out-10)# match aspath bgp-local
Agg2(config-route-map-bgp-ebgp-out-10)# exit
Agg2(config)# router bgp 64500
Agg2(config-bgp)# address-family ipv4 unicast
Agg2(config-bgp-ipv4-uc)# neighbor 10.255.102.12 route-map bgp-ebgp-out out
Agg2(config-bgp-ipv4-uc)# end
Agg2# clear bgp 10.255.102.12
```

10. On Agg2, verify that only the 100.100.1.0/24 prefix is advertised to routerB.

```
Agg2# show bgp ipv4 unicast neighbors 10.255.102.12 advertised-routes
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
               i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default
Local Router-ID 10.x.0.3
```

Network	Nexthop	Metric	LocPrf	Weight	Path
*>e 100.100.1.0/24	10.255.102.3	0	0	0	64500 i

Total number of entries 1

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.

```
Agg1(config)# ip aspath-list partner seq 10 permit ^64600$
```

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.

```
Agg1(config)# route-map bgp-ebgp-out permit seq 20
Agg1(config-route-map-bgp-ebgp-out-20)# match aspath partner
Agg1(config-route-map-bgp-ebgp-out-20)# end
```

13. On Agg1, verify the advertised routes to routerA now include the partner routes.

```
Agg1# show route-map bgp-ebgp-out
Route-map: bgp-ebgp-out
  Seq 10, permit,
    Match :
      aspath-list           : bgp-local
    Set :
  Seq 20, permit,
    Match :
      aspath-list           : partner
    Set :
```

```
Agg1# clear bgp 10.255.101.11
```

```
Agg1# show bgp ipv4 unicast neighbors 10.255.101.11 advertised-routes
```

Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
i internal, e external S Stale, R Removed

Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default

Local Router-ID 10.x.0.2

Network	Nexthop	Metric	LocPrf	Weight	Path
*>e 100.100.1.0/24	10.255.101.2	0	0	0	64500 i
*>e 198.18.14.0/24	10.255.101.2	0	0	0	64500 64600 i
*>e 198.18.15.0/24	10.255.101.2	0	0	0	64500 64600 ?

14. Repeat this configuration on Agg2.

```
Agg2# conf t
Agg2(config)# ip aspath-list partner seq 10 permit ^64600$
Agg2(config)# route-map bgp-ebgp-out permit seq 20
Agg2(config-route-map-bgp-ebgp-out-20)# match aspath partner
```

```
Agg2(config-route-map-bgp-ebgp-out-20)# exit
Agg2(config)# show route-map bgp-ebgp-out
```

```
Route-map: bgp-ebgp-out
  Seq 10, permit,
    Match :
      aspath-list          : bgp-local
    Set :
  Seq 20, permit,
    Match :
      aspath-list          : partner
    Set :
```

```
Agg2(config)# do clear bgp 10.255.102.12
```

```
Agg2(config)# show bgp ipv4 unicast neighbors 10.255.102.12 advertised-routes
```

```
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
               i internal, e external S Stale, R Removed
```

```
Origin codes: i - IGP, e - EGP, ? - incomplete
```

```
VRF : default
```

```
Local Router-ID 10.x.0.3
```

Network	Nexthop	Metric	LocPrf	Weight	Path
*>e 100.100.1.0/24	10.255.102.3	0	0	0	64500 i
*>e 198.18.14.0/24	10.255.102.3	0	0	0	64500 64600 i
*>e 198.18.15.0/24	10.255.102.3	0	0	0	64500 64600 ?

Total number of entries 3

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

```
Agg1# show bgp ipv4 unicast
```

Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
i internal, e external S Stale, R Removed

Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default

Local Router-ID 10.x.0.2

Network	Nextthop	Metric	LocPrf	Weight	Path
*> 100.100.1.0/24	0.0.0.0	0	100	0	i
* i 100.100.1.0/24	10.x.0.3	0	100	0	i
*>e 192.0.2.0/24	10.255.101.11	0	100	0	64511 i
*>i 198.18.14.0/24	10.x.0.4	0	100	0	64600 i
*>i 198.18.15.0/24	10.x.0.4	0	100	0	64600 ?
* i 198.19.100.0/24	10.x.0.3	0	100	0	64512 64513 i
*>e 198.19.100.0/24	10.255.101.11	0	100	0	64511 64513 i
*>e 198.19.101.0/24	10.255.101.11	0	100	0	64511 64513 i
*>i 198.19.102.0/24	10.x.0.3	0	100	0	64512 64513 i

```
* e 198.19.102.0/24 10.255.101.11 0 100 0 64511 64512 64513 i
*>i 198.51.100.0/24 10.x.0.3 0 100 0 64512 i
* e 198.51.100.0/24 10.255.101.11 0 100 0 64511 64512 i
Total number of entries 12
```

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

```
Agg2# show bgp ipv4 unicast
```

Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
i internal, e external S Stale, R Removed

Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default

Local Router-ID 10.x.0.3

Network	Nexthop	Metric	LocPrf	Weight	Path
*>i 100.100.1.0/24	10.x.0.2	0	100	0	i
s> 100.100.1.0/26	0.0.0.0	0	100	0	i
*>i 192.0.2.0/24	10.x.0.2	0	100	0	64511 i
* e 192.0.2.0/24	10.255.102.12	0	100	0	64512 64511 i
*>i 198.18.14.0/24	10.x.0.4	0	100	0	64600 i
*>i 198.18.15.0/24	10.x.0.4	0	100	0	64600 ?
* i 198.19.100.0/24	10.x.0.2	0	100	0	64511 64513 i
*>e 198.19.100.0/24	10.255.102.12	0	100	0	64512 64513 i
*>i 198.19.101.0/24	10.x.0.2	0	100	0	64511 64513 i
* e 198.19.101.0/24	10.255.102.12	0	100	0	64512 64511 64513 i
*>e 198.19.102.0/24	10.255.102.12	0	100	0	64512 64513 i
*>e 198.51.100.0/24	10.255.102.12	0	100	0	64512 i

Total number of entries 12

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

```
Sw1# show ip route bgp
```

Displaying ipv4 routes selected for forwarding

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

```
100.100.1.0/24, vrf default
    via 10.x.41.2, [200/0], bgp
    via 10.x.42.3, [200/0], bgp
192.0.2.0/24, vrf default
    via 10.x.41.2, [200/0], bgp
198.18.14.0/24, vrf default
    via 10.x.48.5, [20/0], bgp
198.18.15.0/24, vrf default
    via 10.x.48.5, [20/0], bgp
198.19.100.0/24, vrf default
    via 10.x.41.2, [200/0], bgp
198.19.101.0/24, vrf default
    via 10.x.41.2, [200/0], bgp
198.19.102.0/24, vrf default
    via 10.x.42.3, [200/0], bgp
198.51.100.0/24, vrf default
    via 10.x.42.3, [200/0], bgp
```

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

```
Agg2# show ip route bgp
```

Displaying ipv4 routes selected for forwarding

```
'[x/y]' denotes [distance/metric]
100.100.1.0/24, vrf default
    via 10.x.40.2, [200/0], bgp
    via 10.x.1.2, [200/0], bgp
192.0.2.0/24, vrf default
    via 10.x.40.2, [200/0], bgp
    via 10.x.1.2, [200/0], bgp
198.18.14.0/24, vrf default
    via 10.x.42.4, [200/0], bgp
198.18.15.0/24, vrf default
    via 10.x.42.4, [200/0], bgp
198.19.100.0/24, vrf default
    via 10.255.102.12, [20/0], bgp
198.19.101.0/24, vrf default
    via 10.x.40.2, [200/0], bgp
    via 10.x.1.2, [200/0], bgp
198.19.102.0/24, vrf default
    via 10.255.102.12, [20/0], bgp
198.51.100.0/24, vrf default
    via 10.255.102.12, [20/0], bgp
```

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'.

```
Agg1# conf t
Agg1(config)# router bgp 64500
Agg1(config-bgp)# bgp default local-preference 200

All current BGP sessions in VRF default will be restarted
Do you want to continue (y/n)? y
```

6. On Sw1, review the BGP IPv4 routing table for entries 198.19.100.0/24 and 198.19.102.0/24.

```
Sw1# show bgp ipv4 unicast

Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
               i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default
Local Router-ID 10.x.0.4
```

Network	Nexthop	Metric	LocPrf	Weight	Path
*>i 100.100.1.0/24	10.x.0.2	0	200	0	i
* i 100.100.1.0/24	10.x.0.3	0	100	0	i
*>i 192.0.2.0/24	10.x.0.2	0	200	0	64511 i

```
*>e 198.18.14.0/24 10.x.48.5 0 100 0 64600 i
*>e 198.18.15.0/24 10.x.48.5 0 100 0 64600 ?
*>i 198.19.100.0/24 10.x.0.2 0 200 0 64511 64513 i
*>i 198.19.101.0/24 10.x.0.2 0 200 0 64511 64513 i
*>i 198.19.102.0/24 10.x.0.2 0 200 0 64511 64512 64513 i
*>i 198.51.100.0/24 10.x.0.2 0 200 0 64511 64512 i
Total number of entries 9
```

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.

```
Agg2# show bgp ipv4 unicast
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default
Local Router-ID 10.x.0.3

   Network                Nexthop      Metric    LocPrf    Weight Path
*>i 100.100.1.0/24         10.x.0.2      0         200       0      i
s> 100.100.1.0/26         0.0.0.0       0         100       0      i
*>i 192.0.2.0/24          10.x.0.2      0         200       0     64511 i
* e 192.0.2.0/24         10.255.102.12 0         100       0     64512 64511 i
*>i 198.18.14.0/24        10.x.0.4      0         100       0     64600 i
*>i 198.18.15.0/24        10.x.0.4      0         100       0     64600 ?
*>i 198.19.100.0/24       10.x.0.2      0         200       0     64511 64513 i
* e 198.19.100.0/24       10.255.102.12 0         100       0     64512 64513 i
*>i 198.19.101.0/24       10.x.0.2      0         200       0     64511 64513 i
* e 198.19.101.0/24       10.255.102.12 0         100       0     64512 64511 64513 i
*>i 198.19.102.0/24       10.x.0.2      0         200       0     64511 64512 64513 i
* e 198.19.102.0/24       10.255.102.12 0         100       0     64512 64513 i
*>i 198.51.100.0/24       10.x.0.2      0         200       0     64511 64512 i
* e 198.51.100.0/24       10.255.102.12 0         100       0     64512 i
Total number of entries 14
```

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.

```
Agg1(config)# ip prefix-list net100 seq 10 permit 198.19.100.0/24 ge 24 le 24
Agg1(config)# ip prefix-list net101 seq 10 permit 198.19.101.0/24 ge 24 le 24
Agg1(config)# ip prefix-list net102 seq 10 permit 198.19.102.0/24 ge 24 le 24
```

10. On Agg1, configure a route-map that assigns each prefix the required preference.

```
Agg1(config)# route-map bgp-ebgp-in permit seq 10
Agg1(config-route-map-bgp-ebgp-in-10)# match ip address prefix-list net100
Agg1(config-route-map-bgp-ebgp-in-10)# set local-preference 300

Agg1(config-route-map-bgp-ebgp-in-10)# route-map bgp-ebgp-in permit seq 20
Agg1(config-route-map-bgp-ebgp-in-10)# match ip address prefix-list net101
Agg1(config-route-map-bgp-ebgp-in-10)# set local-preference 300

Agg1(config-route-map-bgp-ebgp-in-10)# route-map bgp-ebgp-in permit seq 30
Agg1(config-route-map-bgp-ebgp-in-10)# match ip address prefix-list net102
Agg1(config-route-map-bgp-ebgp-in-10)# set local-preference 100
Agg1(config-route-map-bgp-ebgp-in-10)# exit
```

11. Apply this route-map on the routerA BGP peer in the inbound direction.

```
Agg1(config)# router bgp 64500
Agg1(config-bgp)# address-family ipv4 unicast
Agg1(config-bgp-ipv4-uc)# neighbor 10.255.101.11 route-map bgp-ebgp-in in
Agg1(config-bgp-ipv4-uc)# exit
```

12. On Agg2, repeat these steps with the adjusted local preference values.

```
ip prefix-list net100 seq 10 permit 198.19.100.0/24 ge 24 le 24
ip prefix-list net101 seq 10 permit 198.19.101.0/24 ge 24 le 24
ip prefix-list net102 seq 10 permit 198.19.102.0/24 ge 24 le 24

route-map bgp-ebgp-in permit seq 10
match ip address prefix-list net100
set local-preference 300

route-map bgp-ebgp-in permit seq 20
match ip address prefix-list net101
set local-preference 100
```



```
route-map bgp-ebgp-in permit seq 30
match ip address prefix-list net102
set local-preference 300

router bgp 64500
address-family ipv4 unicast
neighbor 10.255.102.12 route-map bgp-ebgp-in in
exit
```

13. On Sw1, review the BGP IPv4 routing table using 'show bgp ipv4 unicast'. Check the local preference for the 3 requested prefixes.

```
Sw1# show bgp ipv4 unicast

Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
               i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default
Local Router-ID 10.x.0.4
```

Network	Nexthop	Metric	LocPrf	Weight	Path
*>i 100.100.1.0/24	10.x.0.2	0	200	0	i
* i 100.100.1.0/24	10.x.0.3	0	100	0	i
*>i 192.0.2.0/24	10.x.0.2	0	200	0	64511 i
*>e 198.18.14.0/24	10.x.48.5	0	100	0	64600 i
*>e 198.18.15.0/24	10.x.48.5	0	100	0	64600 ?
*>i 198.19.100.0/24	10.x.0.2	0	200	0	64511 64513 i
*>i 198.19.101.0/24	10.x.0.2	0	200	0	64511 64513 i
*>i 198.19.102.0/24	10.x.0.2	0	200	0	64511 64512 64513 i
*>i 198.51.100.0/24	10.x.0.2	0	200	0	64511 64512 i

Total number of entries 9

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.

```
Agg1# clear bgp 10.255.101.11
```

```
Agg2# clear bgp 10.255.102.12
```

14. On Sw1, review the BGP IPv4 routing table. Check the local preference for the 3 requested prefixes.

```
Sw1# show bgp ipv4 unicast

Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
               i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default
Local Router-ID 10.x.0.4
```

Network	Nexthop	Metric	LocPrf	Weight	Path
*>i 100.100.1.0/24	10.x.0.2	0	200	0	i
* i 100.100.1.0/24	10.x.0.3	0	100	0	i
*>e 198.18.14.0/24	10.x.48.5	0	100	0	64600 i
*>e 198.18.15.0/24	10.x.48.5	0	100	0	64600 ?
*>i 198.19.100.0/24	10.x.0.2	0	300	0	64511 64513 i
* i 198.19.100.0/24	10.x.0.3	0	300	0	64512 64513 i
*>i 198.19.101.0/24	10.x.0.2	0	300	0	64511 64513 i
*>i 198.19.102.0/24	10.x.0.3	0	300	0	64512 64513 i
Total number of entries 8					

15. On Sw1, review the IP routing table and verify that the correct next hop is selected for each of these routes.

```
Sw1# show ip route bgp
Displaying ipv4 routes selected for forwarding

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

100.100.1.0/24, vrf default
    via 10.x.41.2, [200/0], bgp
198.18.14.0/24, vrf default
    via 10.x.48.5, [20/0], bgp
198.18.15.0/24, vrf default
    via 10.x.48.5, [20/0], bgp
198.19.100.0/24, vrf default
    via 10.x.41.2, [200/0], bgp
198.19.101.0/24, vrf default
    via 10.x.41.2, [200/0], bgp
198.19.102.0/24, vrf default
    via 10.x.42.3, [200/0], bgp
```

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.

```
Sw1# show bgp ipv4 unicast
```

```
Sw1# show ip route 198.19.100.0
Displaying ipv4 routes selected for forwarding

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

198.19.100.0/24, vrf default
    via 10.x.41.2, [200/0], bgp
```

17. On Sw1, configure the AS multipath relax feature, confirm the reset of the sessions.

```
Sw1# conf t
Sw1(config)# router bgp 64500
Sw1(config-bgp)# bgp bestpath as-path multipath-relax

All current BGP sessions in VRF default will be restarted
Do you want to continue (y/n)? y

Sw1(config-bgp)# end
```

18. On Sw1, review the current BGP IPv4 routing table.

```
Sw1# show bgp ipv4 unicast

Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
               i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default
Local Router-ID 10.x.0.4
```

Network	Nexthop	Metric	LocPrf	Weight	Path
*>i 100.100.1.0/24	10.x.0.2	0	200	0	i
* i 100.100.1.0/24	10.x.0.3	0	100	0	i
*>e 198.18.14.0/24	10.x.48.5	0	100	0	64600 i
*>e 198.18.15.0/24	10.x.48.5	0	100	0	64600 ?
*>i 198.19.100.0/24	10.x.0.2	0	300	0	64511 64513 i
*=i 198.19.100.0/24	10.x.0.3	0	300	0	64512 64513 i
*>i 198.19.101.0/24	10.x.0.2	0	300	0	64511 64513 i
*>i 198.19.102.0/24	10.x.0.3	0	300	0	64512 64513 i

Total number of entries 8

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.

```
Sw1# show ip route 198.19.100.0
Displaying ipv4 routes selected for forwarding

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

```
198.19.100.0/24, vrf default
  via 10.x.41.2, [200/0], bgp
  via 10.x.42.3, [200/0], bgp
```

Optional: Test link failure and route failover

20. On Agg1, disable the uplink 1/1/7.

```
Agg1# conf t
Agg1(config)# interface 1/1/7
Agg1(config-if)# shutdown
```

21. On Sw1, verify the next-hop for the 198.19.101.0/24 network is Agg2.(this was Agg1 previously)

```
Sw1# show bgp ipv4 unicast

Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
               i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default
Local Router-ID 10.x.0.4

   Network                Nexthop          Metric    LocPrf    Weight Path
* > i 100.100.1.0/24       10.x.0.2         0         200      0      i
* i 100.100.1.0/24        10.x.0.3         0         100      0      i
* > e 198.18.14.0/24       10.x.48.5        0         100      0      64600 i
* > e 198.18.15.0/24       10.x.48.5        0         100      0      64600 ?
* > i 198.19.100.0/24      10.x.0.3         0         300      0      64512 64513 i
* > i 198.19.101.0/24      10.x.0.3         0         100      0      64512 64511 64513 i
* > i 198.19.102.0/24      10.x.0.3         0         300      0      64512 64513 i
Total number of entries 7
```

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.

```
Agg1(config-if)# no shutdown
Agg1(config-if)# end
```

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.

```
Agg2# show bgp ipv4 unicast
```

```
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
               i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete
```

```
VRF : default
```

```
Local Router-ID 10.x.0.3
```

Network	Nexthop	Metric	LocPrf	Weight	Path
*>i 100.100.1.0/24	10.x.0.2	0	200	0	i
s> 100.100.1.0/26	0.0.0.0	0	100	0	i
*>i 198.18.14.0/24	10.x.0.4	0	100	0	64600 i
*>i 198.18.15.0/24	10.x.0.4	0	100	0	64600 ?
* i 198.19.100.0/24	10.x.0.2	0	300	0	64511 64513 i
*>e 198.19.100.0/24	10.255.102.12	0	300	0	64512 64513 i
*>i 198.19.101.0/24	10.x.0.2	0	300	0	64511 64513 i
* e 198.19.101.0/24	10.255.102.12	0	100	0	64512 64511 64513 i
*>e 198.19.102.0/24	10.255.102.12	0	300	0	64512 64513 i

Total number of entries 9

26. On Agg2, check the IP routing table next hop for the 198.19.101.0/24 route.

```
Agg2# show ip route 198.19.101.0/24
```

```
Displaying ipv4 routes selected for forwarding
```

```
'[x/y]' denotes [distance/metric]
198.19.101.0/24, vrf default
    via 10.x.40.2, [200/0], bgp
    via 10.x.1.2, [200/0], bgp
```

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.

```
Agg2# show route-map
Route-map: bgp-ebgp-in
Seq 10, permit,
  Match :
    ip prefix-list          : net100
  Set :
    local-preference        : 300
Seq 20, permit,
  Match :
    ip prefix-list          : net101
  Set :
    local-preference        : 100
Seq 30, permit,
  Match :
    ip prefix-list          : net102
  Set :
    local-preference        : 300
----- Output Omitted -----
```

```
Agg2(config)# route-map bgp-ebgp-in permit seq 20
Agg2(config-route-map-bgp-ebgp-in-20)# set weight 400
Agg2(config-route-map-bgp-ebgp-in-20)# end
```

28. On Agg2, review the BGP IPv4 routing table

```
Agg2# show bgp ipv4 unicast
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default
Local Router-ID 10.x.0.3
```

Network	Nexthop	Metric	LocPrf	Weight	Path
*>i 100.100.1.0/24	10.x.0.2	0	200	0	i
s> 100.100.1.0/26	0.0.0.0	0	100	0	i
*>i 198.18.14.0/24	10.x.0.4	0	100	0	64600 i

```
*>i 198.18.15.0/24      10.x.0.4      0      100      0      64600 ?
* i 198.19.100.0/24     10.x.0.2      0      300      0      64511 64513 i
*>e 198.19.100.0/24     10.255.102.12 0      300      0      64512 64513 i
*>i 198.19.101.0/24     10.x.0.2      0      300      0      64511 64513 i
* e 198.19.101.0/24     10.255.102.12 0      100      0      64512 64511 64513 i
*>e 198.19.102.0/24     10.255.102.12 0      300      0      64512 64513 i
```

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.

```
Agg2# clear debug buffer
```

30. On Agg2, use a route refresh to refresh the routes from the routerB. This should apply the changed weight value.

```
Agg2# clear bgp ipv4 unicast 10.255.102.12 soft in
```

31. On Agg2, review the debug buffer. The BGP session remained established and the routers simply refresh the routes.

```
Agg2# show debug buffer
----- Output Omitted -----
2020-12-07:11:56:41.176526|hpe-routing|LOG_INFO|AMM|-|BGP|BGP_CONFIG|BGP Aggregate Address
added
2020-12-07:11:56:41.176544|hpe-routing|LOG_INFO|AMM|-|BGP|BGP_CONFIG|Aggregate Address:
100.100.1.0/24 options: summary-only configured
2020-12-07:11:56:41.187369|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|A ROUTE REFRESH message is
being sent to a neighbor.
2020-12-07:11:56:41.187381|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|RIB Manager entity index:
268763136
2020-12-07:11:56:41.187389|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|Local address:
10.255.102.3
2020-12-07:11:56:41.187398|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|Local port:                0
2020-12-07:11:56:41.187405|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|Remote address:
10.255.102.12
2020-12-07:11:56:41.187411|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|Remote port:                0
2020-12-07:11:56:41.187418|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|Remote AS number:          64512
2020-12-07:11:56:41.187424|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|Remote BGP ID:
0XC6336432
2020-12-07:11:56:41.187431|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|AFI:                      1
2020-12-07:11:56:41.187438|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|SAFI:                    1VRF
Name = default.
```

```
2020-12-07:11:56:41.187664|hpe-routing|LOG_INFO|AMM|-|BGP|BGP_CONFIG|BGP Aggregate Address
added
2020-12-07:11:56:41.187676|hpe-routing|LOG_INFO|AMM|-|BGP|BGP_CONFIG|Aggregate Address:
100.100.1.0/24 options: summary-only configured
2020-12-07:11:56:41.190642|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|A ROUTE REFRESH message has
been received from a neighbor.
2020-12-07:11:56:41.190654|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|NM entity index      =
268763136
----- Output Omitted -----
```

32. On Agg2, review the BGP IPv4 routing table and check the new weight value for the 192.19.101.0/24 route.

```
Agg2# show bgp ipv4 unicast
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
               i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default
Local Router-ID 10.x.0.3
```

Network	Nexthop	Metric	LocPrf	Weight	Path
*>i 100.100.1.0/24	10.x.0.2	0	200	0	i
s> 100.100.1.0/26	0.0.0.0	0	100	0	i
*>i 198.18.14.0/24	10.x.0.4	0	100	0	64600 i
*>i 198.18.15.0/24	10.x.0.4	0	100	0	64600 ?
* i 198.19.100.0/24	10.x.0.2	0	300	0	64511 64513 i
*>e 198.19.100.0/24	10.255.102.12	0	300	0	64512 64513 i
* i 198.19.101.0/24	10.x.0.2	0	300	0	64511 64513 i
*>e 198.19.101.0/24	10.255.102.12	0	100	400	64512 64511 64513 i
*>e 198.19.102.0/24	10.255.102.12	0	300	0	64512 64513 i

Total number of entries 9

33. On Agg2, check the IP routing table next hop for the 198.19.101.0/24 route.

```
Agg2# show ip route 198.19.101.0/24
Displaying ipv4 routes selected for forwarding

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

198.19.101.0/24, vrf default
    via 10.255.102.12, [20/0], bgp
```

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.

```
Sw1# show ip route 198.19.101.0/24
Displaying ipv4 routes selected for forwarding

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

198.19.101.0/24, vrf default
    via 10.x.41.2, [200/0], bgp
```

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

```
Agg1# show bgp ipv4 unicast neighbors 10.255.101.11 received-routes
```

```
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
               i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default
Local Router-ID 10.x.0.2
```

Network	Nexthop	Metric	LocPrf	Weight	Path
Total number of entries 0					

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 ?

```
Agg1# show bgp ipv4 unicast neighbors 10.255.101.11 ?

advertised-routes  Displays all routes that have been advertised to neighbors
paths              Path information of routes received by this neighbor
received-routes    Displays all received-routes that are received and
                   rejected
routes             Displays all routes that are accepted
vsx-peer           Displays VSX peer switch information
```

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.

```
Agg1# show bgp ipv4 unicast neighbors 10.255.101.11 routes

Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
               i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default
Local Router-ID 10.x.0.2
```

Network	Nexthop	Metric	LocPrf	Weight	Path
*>e 198.19.100.0/24	10.255.101.11	0	300	0	64511 64513 i
*>e 198.19.101.0/24	10.255.101.11	0	300	0	64511 64513 i
* e 198.19.102.0/24	10.255.101.11	0	100	0	64511 64512 64513 i

```
Total number of entries 3
```

38. On Agg1, clear the debug buffer and configure the peer 10.255.101.11 with the soft reconfiguration option.

```
Agg1# conf t
Agg1(config)# router bgp 64500
Agg1(config-bgp)# address-family ipv4 unicast
Agg1(config-bgp-ipv4-uc)# do clear debug buffer
Agg1(config-bgp-ipv4-uc)# neighbor 10.255.101.11 soft-reconfiguration inbound
Agg1(config-bgp-ipv4-uc)# end
```

39. On Agg1, review the debug buffer. received routes.

```
Agg1# show debug buffer
----- Output Omitted -----
2020-12-07:12:16:23.356514|hpe-routing|LOG_INFO|AMM|-|BGP|BGP_CONFIG|Peer 10.255.101.11:
Inbound Soft Reconfig Configured
2020-12-07:12:16:23.367352|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|A NOTIFICATION message is being
sent to a neighbor.
2020-12-07:12:16:23.367365|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|NM entity index      =
268763136
2020-12-07:12:16:23.367372|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|Local address      = (none)
2020-12-07:12:16:23.367379|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|Local port        = 0
2020-12-07:12:16:23.367385|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|Remote address     =
10.255.101.11
2020-12-07:12:16:23.367391|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|Remote port        = 0
2020-12-07:12:16:23.367398|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|Scope ID          = 0
2020-12-07:12:16:23.367404|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|Remote AS number   = 64511
2020-12-07:12:16:23.367410|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|Remote BGP ID     =
192.0.2.50
2020-12-07:12:16:23.367417|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|Error code        = Cease
(6)
2020-12-07:12:16:23.367424|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|Error subcode     =
Administrative Shutdown (2)VRF Name = default.
2020-12-07:12:16:23.367536|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|A connection's FSM state has
deteriorated.
----- Output Omitted -----
```

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.

```
Agg1# show bgp ipv4 unicast neighbors 10.255.101.11 received-routes
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
               i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default
Local Router-ID 10.x.0.2

   Network          Nexthop          Metric      LocPrf      Weight Path
e 192.0.2.0/24      10.255.101.11    0           0           0        64511 i
*>e 198.19.100.0/24 10.255.101.11    0           0           0        64511 64513 i
*>e 198.19.101.0/24 10.255.101.11    0           0           0        64511 64513 i
*>e 198.19.102.0/24 10.255.101.11    0           0           0        64511 64512 64513 i
e 198.51.100.0/24   10.255.101.11    0           0           0        64511 64512 i
Total number of entries 5
```

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.

```
Agg2# conf t
Agg2(config)# router bgp 64500
Agg2(config-bgp)# address-family ipv4 unicast
Agg2(config-bgp-ipv4-uc)# neighbor 10.255.102.12 soft-reconfiguration inbound
Agg2(config-bgp-ipv4-uc)# end
```

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).

```
Agg2# show route-map bgp-ebgp-in
Route-map: bgp-ebgp-in
  Seq 10, permit,
    Match :
      ip prefix-list      : net100
    Set :
      local-preference    : 300
  Seq 20, permit,
    Match :
      ip prefix-list      : net101
    Set :
      local-preference    : 100
      weight              : 400
  Seq 30, permit,
    Match :
      ip prefix-list      : net102
    Set :
      local-preference    : 300
```

```
Agg2# conf t
Agg2(config)# route-map bgp-ebgp-in permit seq 20
Agg2(config-route-map-bgp-ebgp-in-20)# set weight 500
Agg2(config-route-map-bgp-ebgp-in-20)# end
```

43. On Agg2, review the BGP IPv4 routing table.

```
Agg2# show bgp ipv4 unicast
```

Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
i internal, e external S Stale, R Removed

Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default

Local Router-ID 10.x.0.3

Network	NextHop	Metric	LocPrf	Weight	Path
*>i 100.100.1.0/24	10.x.0.2	0	200	0	i
s> 100.100.1.0/26	0.0.0.0	0	100	0	i
*>i 198.18.14.0/24	10.x.0.4	0	100	0	64600 i
*>i 198.18.15.0/24	10.x.0.4	0	100	0	64600 ?
* i 198.19.100.0/24	10.x.0.2	0	300	0	64511 64513 i
*>e 198.19.100.0/24	10.255.102.12	0	300	0	64512 64513 i
* i 198.19.101.0/24	10.x.0.2	0	300	0	64511 64513 i
*>e 198.19.101.0/24	10.255.102.12	0	100	400	64512 64511 64513 i
*>e 198.19.102.0/24	10.255.102.12	0	300	0	64512 64513 i

Total number of entries 9

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.

```
Agg2# clear debug buffer
```

```
Agg2# clear bgp ipv4 unicast 10.255.102.12 soft in
```

45. On Agg2, review the BGP IPv4 routing table.

```
Agg2# show bgp ipv4 unicast
```

Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
i internal, e external S Stale, R Removed

Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default

Local Router-ID 10.x.0.3

Network	NextHop	Metric	LocPrf	Weight	Path
*>i 100.100.1.0/24	10.x.0.2	0	200	0	i
s> 100.100.1.0/26	0.0.0.0	0	100	0	i
*>i 198.18.14.0/24	10.x.0.4	0	100	0	64600 i
*>i 198.18.15.0/24	10.x.0.4	0	100	0	64600 ?
* i 198.19.100.0/24	10.x.0.2	0	300	0	64511 64513 i
*>e 198.19.100.0/24	10.255.102.12	0	300	0	64512 64513 i
* i 198.19.101.0/24	10.x.0.2	0	300	0	64511 64513 i
*>e 198.19.101.0/24	10.255.102.12	0	100	500	64512 64511 64513 i
*>e 198.19.102.0/24	10.255.102.12	0	300	0	64512 64513 i

Total number of entries 9

Q: Was the new weight of 500 applied?

A: Yes, the weight was applied.

46. On Agg2, review the debug buffer.

```
Agg2# show debug buffer module bgp
```

```
2020-11-03:04:45:41.092261|hpe-routing|LOG_INFO|AMM|-|BGP|BGP_CONFIG|BGP Aggregate Address added
```

```
2020-11-03:04:45:41.092281|hpe-routing|LOG_INFO|AMM|-|BGP|BGP_CONFIG|Aggregate Address: 100.100.1.0/24 options: summary-only configured
```

```
2020-11-03:04:45:41.102815|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|An Adj-RIB-In is being recalculated as a result of changed import
```

```
2020-11-03:04:45:41.102825|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|policy.
```

```
2020-11-03:04:45:41.102832|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|RIB Manager entity index: 268763136
```

```
2020-11-03:04:45:41.102838|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|Neighbor Index: 0X00000004
```

```
2020-11-03:04:45:41.102845|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|AFI: 1
```

```
2020-11-03:04:45:41.102852|hpe-routing|LOG_INFO|AMM|-|BGP|BGP|SAFI: 1VRF
```

```
Name = default.
```

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.

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.

```
routerB# show bgp ipv4 unicast
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
               i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default
Local Router-ID 198.51.100.50

   Network          Nexthop        Metric    LocPrf    Weight Path
*>e 100.100.1.0/24  10.255.102.3    0         100       0      64500 i
* e 100.100.1.0/24  10.255.103.11   0         100       0      64511 64500 i
----- Output Omitted -----
```

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.

```
Agg2# show route-map bgp-ebgp-out
Route-map: bgp-ebgp-out
Seq 10, permit,
Match :
  aspath-list          : bgp-local
Set :
Seq 20, permit,
Match :
  aspath-list          : partner
Set :
```

3. On Agg2, review the advertised routes to routerB.

```
Agg2# show bgp ipv4 unicast neighbors 10.255.102.12 advertised-routes
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default
Local Router-ID 10.x.0.3
```

Network	Nexthop	Metric	LocPrf	Weight	Path
*>e 100.100.1.0/24	10.255.102.3	0	0	0	64500 i
*>e 198.18.14.0/24	10.255.102.3	0	0	0	64500 64600 i
*>e 198.18.15.0/24	10.255.102.3	0	0	0	64500 64600 ?

Total number of entries 3

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.

```
Agg2# conf t
Agg2(config)# route-map bgp-ebgp-out permit seq 10
Agg2(config-route-map-bgp-ebgp-out-10)# set as-path prepend 64500 64500 64500
Agg2(config-route-map-bgp-ebgp-out-10)# end
```

5. On Agg2, review the advertised routes to routerB, they should now have a longer AS Path list.

```
Agg2# show bgp ipv4 unicast neighbors 10.255.102.12 advertised-routes
```

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.

```
Agg2# clear bgp 10.255.102.12
```

7. On Agg2, review the advertised routes to routerB again, they should now have a longer AS Path list.

```
Agg2# show bgp ipv4 unicast neighbors 10.255.102.12 advertised-routes
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
               i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default
Local Router-ID 10.x.0.3
```

Network	Nexthop	Metric	LocPrf	Weight	Path
*>e 100.100.1.0/24	10.255.102.3	0	0	0	64500 64500 64500 64500
i					
*>e 198.18.14.0/24	10.255.102.3	0	0	0	64500 64600 i
*>e 198.18.15.0/24	10.255.102.3	0	0	0	64500 64600 ?

Total number of entries 3

8. On routerB, review the BGP routing table entry for 100.100.1.0/24 again.

```
routerB# show bgp ipv4 unicast
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
               i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default
Local Router-ID 198.51.100.50
```

Network	Nexthop	Metric	LocPrf	Weight	Path
* e 100.100.1.0/24	10.255.102.3	0	100	0	64500 64500 64500 64500
i					
*>e 100.100.1.0/24	10.255.103.11	0	100	0	64511 64500 i

----- Output Omitted -----

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.

```
Agg2# ping 10.x.12.5 repetitions 1
PING 10.x.12.5 (10.x.12.5) 100(128) bytes of data.
108 bytes from 10.x.12.5: icmp_seq=1 ttl=64 time=0.160 ms

--- 10.x.12.5 ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 0.160/0.160/0.160/0.000 ms
```

11. On Agg2, configure the eBGP peering to Sw2 (10.x.12.5)

```
Agg2# conf t
```

```
Agg2(config)# router bgp 64500
Agg2(config-bgp)# neighbor 10.x.12.5 remote-as 64600
Agg2(config-bgp)# address-family ipv4 unicast
Agg2(config-bgp-ipv4-uc)# neighbor 10.x.12.5 activate
Agg2(config-bgp-ipv4-uc)# end
```

12. On Agg2, verify the eBGP peering is established with Sw2.

```
Agg2# show bgp ipv4 unicast summary
VRF : default
BGP Summary
-----
Local AS           : 64500          BGP Router Identifier : 10.x.0.3
Peers              : 4              Log Neighbor Changes  : No
Cfg. Hold Time     : 180           Cfg. Keep Alive       : 60

Neighbor    Remote-AS  MsgRcvd  MsgSent  Up/Down Time State      AdminStatus
10.x.0.2    64500      353      364      01h:04m:05s Established Up
10.x.0.4    64500      314      361      00h:44m:50s Established Up
10.x.12.5   64600      9         10      00h:00m:59s Established Up
10.255.102.12 64512      15       14      00h:06m:40s Established Up
```

13. On Sw2, review the current BGP IPv4 routing table, check the metric values for the 100.100.1.0/24 route.

```
Sw2# show bgp vrf bgpext ipv4 unicast
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
               i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : bgpext
Local Router-ID 10.x.0.5

Network          Nexthop        Metric      LocPrf      Weight Path
*=e 100.100.1.0/24 10.x.12.3      0           100         0        64500 i
*>e 100.100.1.0/24 10.x.48.4      0           100         0        64500 i
----- Output Omitted -----
```

Apply the MED values

14. On Sw1, define a prefix list for the 100.100.1.0/24 prefix.

```
Sw1# conf t
Sw1(config)# ip prefix-list int100 seq 10 permit 100.100.1.0/24
```

15. Configure a route-map on Sw1 that applies the requested metric prefix list, make sure all other routes are still allowed.

```
Sw1(config)# route-map bgp-ebgp-partner-out permit seq 10
Sw1(config-route-map-bgp-ebgp-partner-out-10)# match ip address prefix-list int100
Sw1(config-route-map-bgp-ebgp-partner-out-10)# set metric 100
Sw1(config-route-map-bgp-ebgp-partner-out-10)# route-map bgp-ebgp-partner-out permit seq 20
```

```
Sw1(config-route-map-bgp-ebgp-partner-out-20)# exit
```

16. Apply the route-map to the Sw2 eBGP peer.

```
Sw1(config)# router bgp 64500
Sw1(config-bgp)# address-family ipv4 unicast
Sw1(config-bgp-ipv4-uc)# neighbor 10.x.48.5 route-map bgp-ebgp-partner-out out
Sw1(config-bgp-ipv4-uc)# end
```

17. Review the advertised routes to Sw2

```
Sw1# show bgp ipv4 unicast neighbors 10.x.48.5 advertised-routes

Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
               i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default
Local Router-ID 10.x.0.4
```

Network	Nexthop	Metric	LocPrf	Weight	Path
*>e 100.100.1.0/24	10.x.48.4	0	0	0	64500 i
*>e 198.19.100.0/24	10.x.48.4	0	0	0	64500 64511 64513 i
*>e 198.19.101.0/24	10.x.48.4	0	0	0	64500 64511 64513 i
*>e 198.19.102.0/24	10.x.48.4	0	0	0	64500 64512 64513 i

Total number of entries 4

18. Make sure the new policy is applied to the Sw2 peer.

```
Sw1# clear bgp 10.x.48.5
```

19. Review the advertised routes again. You should see the adjusted metric values.

```
Sw1# show bgp ipv4 unicast neighbors 10.x.48.5 advertised-routes

Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
               i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default
Local Router-ID 10.x.0.4
```

Network	Nexthop	Metric	LocPrf	Weight	Path
*>e 100.100.1.0/24	10.x.48.4	100	0	0	64500 i

----- Output Omitted -----

20. On Sw2 in VRF bgpext, verify if the MED values are shown in the BGP IPv4 routing table.

```
Sw2# show bgp vrf bgpext ipv4 unicast

Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
               i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : bgpext
Local Router-ID 10.x.0.5
```

```

      Network      Nexthop      Metric      LocPrf      Weight Path
* > e 100.100.1.0/24 10.x.12.3      0          100         0      64500 i
* e 100.100.1.0/24 10.x.48.4      100        100         0      64500 i
----- Output Omitted -----

```

21. On Agg2, define the required prefix list.

```

Agg2# conf t
Agg2(config)# ip prefix-list int100 seq 10 permit 100.100.1.0/24

```

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.

```

Agg2(config)# route-map bgp-ebgp-partner-out permit seq 10
Agg2(config-route-map-bgp-ebgp-partner-out-10)# match ip address prefix-list int100
Agg2(config-route-map-bgp-ebgp-partner-out-10)# set metric 10
Agg2(config-route-map-bgp-ebgp-partner-out-10)# route-map bgp-ebgp-partner-out permit seq 20
Agg2(config-route-map-bgp-ebgp-partner-out-20)# exit

```

23. On Agg2, apply the route-map to the Sw2 eBGP peer and make it effective.

```

Agg2(config)# router bgp 64500
Agg2(config-bgp)# address-family ipv4 unicast
Agg2(config-bgp-ipv4-uc)# neighbor 10.x.12.5 route-map bgp-ebgp-partner-out out
Agg2(config-bgp-ipv4-uc)# do clear bgp 10.x.12.5
Agg2(config-bgp-ipv4-uc)# end

```

24. On Sw2 in vrf bgpext, verify if the MED values are shown in the BGP IPv4 routing table.

```

Sw2# show bgp vrf bgpext ipv4 unicast
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : bgpext
Local Router-ID 10.x.0.5

      Network      Nexthop      Metric      LocPrf      Weight Path
* > e 100.100.1.0/24 10.x.12.3      10          100         0      64500 i
* e 100.100.1.0/24 10.x.48.4      100        100         0      64500 i
* > 198.18.14.0/24 0.0.0.0        0          100         0      i
* > 198.18.15.0/24 0.0.0.0        0          100         0      ?
* e 198.19.100.0/24 10.x.12.3      0          100         0      64500 64512 64513 i
* > e 198.19.100.0/24 10.x.48.4      0          100         0      64500 64511 64513 i
* e 198.19.101.0/24 10.x.12.3      0          100         0      64500 64512 64511 64513 i
* > e 198.19.101.0/24 10.x.48.4      0          100         0      64500 64511 64513 i
* = e 198.19.102.0/24 10.x.12.3      0          100         0      64500 64512 64513 i
* > e 198.19.102.0/24 10.x.48.4      0          100         0      64500 64512 64513 i
Total number of entries 10

```

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.

```
Agg1# show bgp ipv4 unicast neighbors 10.255.101.11 advertised-routes
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
               i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default
Local Router-ID 10.x.0.2

   Network          Nexthop        Metric      LocPrf      Weight Path
* > e 100.100.1.0/24 10.255.101.2    0           0           0       64500 i
* > e 198.18.14.0/24 10.255.101.2    0           0           0       64500 64600 i
* > e 198.18.15.0/24 10.255.101.2    0           0           0       64500 64600 ?
```


Total number of entries 3

28. On Sw2, define a prefix list 'net15' that controls the prefix 198.18.15.0/24.

```
Sw2# conf t
Sw2(config)# ip prefix-list net15 permit 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.

```
Sw2(config)# route-map bgp-ebgp-out permit seq 10
Sw2(config-route-map-bgp-ebgp-out-10)# match ip address prefix-list net15
Sw2(config-route-map-bgp-ebgp-out-10)# set community ?

AA:NN      Community number in AA:NN format
WORD       Set multiple communities for the route. Specify within double
           quotes
internet   Set INTERNET to advertise the prefix to all BGP neighbors.
local-as   Set LOCAL_AS community to not advertise outside of the AS
no-advertise Set NO_ADVERTISE community to not advertise the route
no-export   Set NO_EXPORT community to not advertise the route to eBGP
           neighbors
none       Remove BGP community attribute for the route

Sw2(config-route-map-bgp-ebgp-out-10)# set community no-export ?

additive   Append the community numbers to the route's community attribute
delete     Delete the community numbers from the route's community attribute
<cr>

Sw2(config-route-map-bgp-ebgp-out-10)# set community no-export
Sw2(config-route-map-bgp-ebgp-out-10)# exit
```

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.

```
Sw2(config)# route-map bgp-ebgp-out permit seq 20
```

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.

```
Sw2(config)# router bgp 64600
Sw2(config-bgp)# vrf bgpext
Sw2(config-bgp-vrf)# address-family ipv4 unicast
Sw2(config-bgp-vrf-ipv4-uc)# neighbor 10.x.12.3 route-map bgp-ebgp-out out
Sw2(config-bgp-vrf-ipv4-uc)# neighbor 10.x.48.4 route-map bgp-ebgp-out out
Sw2(config-bgp-vrf-ipv4-uc)# end
```

32. On Sw2, review the current bgp configuration

```
Sw2# show run bgp
```

```
router bgp 64600
!
vrf bgpext
  bgp router-id 10.x.0.5
  neighbor 10.x.12.3 remote-as 64500
  neighbor 10.x.48.4 remote-as 64500
  address-family ipv4 unicast
    neighbor 10.x.12.3 activate
    neighbor 10.x.12.3 route-map bgp-ebgp-out out
    neighbor 10.x.48.4 activate
    neighbor 10.x.48.4 route-map bgp-ebgp-out out
  redistribute connected route-map bgp-import-connected
  network 198.18.14.0/24
exit-address-family
```

33. On Sw2, clear all BGP sessions.

```
Sw2# clear bgp vrf bgpext *
```

Attempt to verify the community attribute on Sw1

34. On Sw1, review the BGP IPv4 details for the 198.18.15.0/24 route.

```
Sw1# show bgp ipv4 unicast 198.18.15.0/24
VRF : default
BGP Local AS 64500      BGP Router-id 10.x.0.4

Network      : 198.18.15.0/24      Nexthop      : 10.x.0.3
Peer         : 10.x.0.3            Origin       : incomplete
Metric       : 0                  Local Pref   : 100
Weight       : 0                  Calc. Local Pref : 100
Best         : No                 Valid        : Yes
Type         : internal           Stale        : No
Originator ID : 0.0.0.0
Aggregator ID :
Aggregator AS :
Atomic Aggregate :
RFD Flaps     : 0                  RFD Penalty  : 0

AS-Path      : 64600

Cluster List :
Communities  :
Ext-Communities :
Network      : 198.18.15.0/24      Nexthop      : 10.x.48.5
Peer         : 10.x.48.5          Origin       : incomplete
Metric       : 0                  Local Pref   : 100
Weight       : 0                  Calc. Local Pref : 100
Best         : Yes                Valid        : Yes
Type         : external           Stale        : No
Originator ID : 0.0.0.0
```

```

Aggregator ID      :
Aggregator AS      :
Atomic Aggregate   :
RFD Flaps          : 0                      RFD Penalty      : 0

AS-Path            : 64600

Cluster List       :
Communities        :
Ext-Communities    :

```

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.

```

Sw2# conf t
Sw2(config)# router bgp 64600
Sw2(config-bgp)# vrf bgpext
Sw2(config-bgp-vrf)# address-family ipv4 unicast
Sw2(config-bgp-vrf-ipv4-uc)# neighbor 10.x.12.3 send-community
Sw2(config-bgp-vrf-ipv4-uc)# neighbor 10.x.48.4 send-community
Sw2(config-bgp-vrf-ipv4-uc)# end

```

NOTE: The command 'send-community' without any options will enable both standard and extended attributes.

36. On Sw2, verify the BGP configuration.

```

Sw2# show run bgp
router bgp 64600
!
vrf bgpext
  bgp router-id 10.x.0.5
  neighbor 10.x.12.3 remote-as 64500
  neighbor 10.x.48.4 remote-as 64500
  address-family ipv4 unicast
    neighbor 10.x.12.3 activate
    neighbor 10.x.12.3 route-map bgp-ebgp-out out
    neighbor 10.x.12.3 send-community both
    neighbor 10.x.48.4 activate
    neighbor 10.x.48.4 route-map bgp-ebgp-out out
    neighbor 10.x.48.4 send-community both

```

```
redistribute connected route-map bgp-import-connected
network 198.18.14.0/24
exit-address-family
```

37. On Sw1, review the BGP IPv4 details for the 198.18.15.0/24 route.

```
Sw1# show bgp ipv4 unicast 198.18.15.0/24
```

```
VRF : default
BGP Local AS 64500      BGP Router-id 10.x.0.4

  Network      : 198.18.15.0/24      Nexthop      : 10.x.0.3
  Peer         : 10.x.0.3            Origin       : incomplete
  Metric       : 0                  Local Pref   : 100
  Weight       : 0                  Calc. Local Pref : 100
  Best         : No                 Valid        : Yes
  Type         : internal           Stale        : No
  Originator ID : 0.0.0.0
  Aggregator ID :
  Aggregator AS :
  Atomic Aggregate :
  RFD Flaps     : 0                  RFD Penalty  : 0

  AS-Path      : 64600

  Cluster List :
  Communities   :
  Ext-Communities :
  Network      : 198.18.15.0/24      Nexthop      : 10.x.48.5
  Peer         : 10.x.48.5            Origin       : incomplete
  Metric       : 0                  Local Pref   : 100
  Weight       : 0                  Calc. Local Pref : 100
  Best         : Yes                 Valid        : Yes
  Type         : external           Stale        : No
  Originator ID : 0.0.0.0
  Aggregator ID :
  Aggregator AS :
  Atomic Aggregate :
  RFD Flaps     : 0                  RFD Penalty  : 0

  AS-Path      : 64600

  Cluster List :
  Communities   : no-export
  Ext-Communities :
```

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.

```
Agg2# show bgp ipv4 unicast 198.18.15.0/24
```

VRF : default

BGP Local AS 64500

BGP Router-id 10.x.0.3

Network	: 198.18.15.0/24	Nexthop	: 10.x.0.4
Peer	: 10.x.0.4	Origin	: incomplete
Metric	: 0	Local Pref	: 100
Weight	: 0	Calc. Local Pref	: 100
Best	: No	Valid	: Yes
Type	: internal	Stale	: No
Originator ID	: 0.0.0.0		
Aggregator ID	:		
Aggregator AS	:		
Atomic Aggregate	:		
RFD Flaps	: 0	RFD Penalty	: 0
AS-Path	: 64600		
Cluster List	:		
Communities	:		
Ext-Communities	:		
Network	: 198.18.15.0/24	Nexthop	: 10.x.12.5
Peer	: 10.x.12.5	Origin	: incomplete
Metric	: 0	Local Pref	: 100
Weight	: 0	Calc. Local Pref	: 100
Best	: Yes	Valid	: Yes
Type	: external	Stale	: No
Originator ID	: 0.0.0.0		
Aggregator ID	:		
Aggregator AS	:		
Atomic Aggregate	:		
RFD Flaps	: 0	RFD Penalty	: 0
AS-Path	: 64600		
Cluster List	:		
Communities	: no-export		
Ext-Communities	:		

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.

```
Agg2# conf t
Agg2(config)# router bgp 64500
Agg2(config-bgp)# address-family ipv4 unicast
Agg2(config-bgp-ipv4-uc)# neighbor 10.x.0.2 send-community
Agg2(config-bgp-ipv4-uc)# neighbor 10.x.0.4 send-community
Agg2(config-bgp-ipv4-uc)# end
```

40. On Sw1, enable Agg1 and Agg2 peers with the 'send-community' attribute.

```
Sw1# conf t
Sw1(config)# router bgp 64500
Sw1(config-bgp)# address-family ipv4 unicast
Sw1(config-bgp-ipv4-uc)# neighbor 10.x.0.2 send-community
Sw1(config-bgp-ipv4-uc)# neighbor 10.x.0.3 send-community
Sw1(config-bgp-ipv4-uc)# end
```

41. On Agg1, apply the configuration to the peer group 'internal'

```
Agg1# conf t
Agg1(config)# router bgp 64500
Agg1(config-bgp)# address-family ipv4 unicast
Agg1(config-bgp-ipv4-uc)# neighbor internal send-community
Agg1(config-bgp-ipv4-uc)# end
```

42. On Agg1, verify the BGP neighbor details for peer Agg2 (10.x.0.3)

```
Agg1# show bgp ipv4 unicast neighbors 10.x.0.3
----- Output Omitted -----

Address Family : IPv4 Unicast
-----

Rt. Reflect. Client : No
Allow-AS in       : 0
Max. Prefix       : 15000
NextHop-Self      : Yes^
Send Community    : both^
Advt. Interval    : 30
Soft Reconfig In  :
Default-Originate :
```

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

```
Agg1# show bgp ipv4 unicast 198.18.15.0/24
```

VRF : default

BGP Local AS 64500 BGP Router-id 10.x.0.2

Network	: 198.18.15.0/24	Nexthop	: 10.x.0.3
Peer	: 10.x.0.3	Origin	: incomplete
Metric	: 0	Local Pref	: 100
Weight	: 0	Calc. Local Pref	: 100
Best	: Yes	Valid	: Yes
Type	: internal	Stale	: No
Originator ID	: 0.0.0.0		
Aggregator ID	:		
Aggregator AS	:		
Atomic Aggregate	:		
RFD Flaps	: 0	RFD Penalty	: 0

AS-Path : 64600

Cluster List :

Communities : no-export

Ext-Communities :

Network	: 198.18.15.0/24	Nexthop	: 10.x.0.4
Peer	: 10.x.0.4	Origin	: incomplete
Metric	: 0	Local Pref	: 100
Weight	: 0	Calc. Local Pref	: 100
Best	: No	Valid	: Yes
Type	: internal	Stale	: No
Originator ID	: 0.0.0.0		
Aggregator ID	:		
Aggregator AS	:		
Atomic Aggregate	:		
RFD Flaps	: 0	RFD Penalty	: 0

AS-Path : 64600

Cluster List :

Communities : no-export

Ext-Communities :

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.

```
Agg1# show bgp ipv4 unicast neighbors 10.255.101.11 advertised-routes
```

Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
i internal, e external S Stale, R Removed

Origin codes: i - IGP, e - EGP, ? - incomplete

```
VRF : default  
Local Router-ID 10.x.0.2
```

Network	Nexthop	Metric	LocPrf	Weight	Path
*>e 100.100.1.0/24	10.255.101.2	0	0	0	64500 i
*>e 198.18.14.0/24	10.255.101.2	0	0	0	64500 64600 i
Total number of entries 2					

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

```
Agg1# show bgp ipv4 unicast neighbors 10.255.101.11 advertised-routes
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
               i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default
Local Router-ID 10.x.0.2
```

Network	Nexthop	Metric	LocPrf	Weight	Path
*>e 100.100.1.0/24	10.255.101.2	0	0	0	64500 i
*>e 198.18.14.0/24	10.255.101.2	0	0	0	64500 64600 i

Total number of entries 2

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

```
Agg1# show run bgp
Agg1# show bgp ipv4 unicast neighbor 10.255.101.11
```

2. On Agg1, review the contents of the route-map to routerA

```
Agg1# show route-map bgp-ebgp-out
Route-map: bgp-ebgp-out
  Seq 10, permit,
    Match :
      aspath-list          : bgp-local
    Set :
  Seq 20, permit,
```

```
Match :
  aspath-list      : partner
Set :
```

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.

```
Agg1# conf t
Agg1(config)# route-map bgp-ebgp-out permit seq 20
Agg1(config-route-map-bgp-ebgp-out-20)# set as-path exclude 64600
```

4. On Agg1, while in the route-map context, use 'show run current' to review the route-map configuration.

```
Agg1(config-route-map-bgp-ebgp-out-20)# show run current
route-map bgp-ebgp-out permit seq 10
  match aspath-list bgp-local
route-map bgp-ebgp-out permit seq 20
  match aspath-list partner
  set as-path exclude 64600
Agg1(config-route-map-bgp-ebgp-out-20)# end
```

5. On Agg1, now ensure that the route-map changes are applied to the routerA BGP session.

```
Agg1# clear bgp 10.255.101.11
```

6. Review the advertised routes to routerA.

```
Agg1# show bgp ipv4 unicast neighbors 10.255.101.11 advertised-routes
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
               i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default
Local Router-ID 10.x.0.2
```

Network	Nexthop	Metric	LocPrf	Weight	Path
*>e 100.100.1.0/24	10.255.101.2	0	0	0	64500 i
*>e 198.18.14.0/24	10.255.101.2	0	0	0	64500 i

```
Total number of entries 2
```

7. On Agg2, repeat this procedure.

```
show bgp ipv4 unicast neighbors 10.255.102.12 advertised-routes
route-map bgp-ebgp-out permit seq 20
  set as-path exclude 64600
exit
do clear bgp 10.255.102.12
show bgp ipv4 unicast neighbors 10.255.102.12 advertised-routes
```

8. On RouterC, verify that the AS 64600 no longer appears in the AS Path for 198.18.14.0/24 prefixes.

```
routerC# show bgp ipv4 unicast
```

```
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,  
i internal, e external S Stale, R Removed
```

```
Origin codes: i - IGP, e - EGP, ? - incomplete
```

```
VRF : default
```

```
Local Router-ID 203.0.113.50
```

Network	Nexthop	Metric	LocPrf	Weight	Path
*>e 100.100.1.0/24	10.255.104.11	0	100	0	64511 64500 i
* e 100.100.1.0/24	10.255.105.12	0	100	0	64512 64511 64500 i
*>e 192.0.2.0/24	10.255.104.11	0	100	0	64511 i
* e 192.0.2.0/24	10.255.105.12	0	100	0	64512 64511 i
*>e 198.18.14.0/24	10.255.104.11	0	100	0	64511 64500 i
* e 198.18.14.0/24	10.255.105.12	0	100	0	64512 64500 i
*> 198.19.100.0/24	0.0.0.0	0	100	0	i
*> 198.19.101.0/24	0.0.0.0	0	100	0	i
*> 198.19.102.0/24	0.0.0.0	0	100	0	i
* e 198.51.100.0/24	10.255.104.11	0	100	0	64511 64512 i
*>e 198.51.100.0/24	10.255.105.12	0	100	0	64512 i
*> 203.0.113.0/24	0.0.0.0	0	100	0	i

Total number of entries 12

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)

```
Sw1# conf t
Sw1(config)# router bgp 64500
Sw1(config-bgp)# no neighbor 10.x.0.3
```

2. On Sw1, verify the peering has been removed.

```
Sw1# show bgp ipv4 unicast summary

VRF : default
BGP Summary
-----
Local AS           : 64500           BGP Router Identifier : 10.x.0.4
Peers              : 2               Log Neighbor Changes  : No
Cfg. Hold Time     : 180            Cfg. Keep Alive       : 60

Neighbor    Remote-AS  MsgRcvd  MsgSent  Up/Down Time State  AdminStatus
10.x.0.2    64500      137     132     01h:47m:41s Established Up
10.x.48.5   64600      76      77      00h:23m:20s Established Up
```

3. On Agg2, remove the neighbor for Sw1 (10.x.0.4)

```
Agg2(config)# router bgp 64500
Agg2(config-bgp)# no neighbor 10.x.0.4
```

Review the result of the broken topology

- 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.

```
Sw1# show bgp ipv4 unicast
```

Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default

Local Router-ID 10.x.0.4

Network	Nexthop	Metric	LocPrf	Weight	Path
*>i 100.100.1.0/24	10.x.0.2	0	200	0	i
*>e 198.18.14.0/24	10.x.48.5	0	100	0	64600 i
*>e 198.18.15.0/24	10.x.48.5	0	100	0	64600 ?
*>i 198.19.100.0/24	10.x.0.2	0	300	0	64511 64513 i
*>i 198.19.101.0/24	10.x.0.2	0	300	0	64511 64513 i

Total number of entries 5

- On Agg1, check the BGP IPv4 routing table for entry 198.19.102.0/24

```
Agg1# show bgp ipv4 unicast
```

Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default

Local Router-ID 10.x.0.2

Network	Nexthop	Metric	LocPrf	Weight	Path
*> 100.100.1.0/24	0.0.0.0	0	200	0	i
* i 100.100.1.0/24	10.x.0.3	0	100	0	i
*>i 198.18.14.0/24	10.x.0.3	0	100	0	64600 i
*=i 198.18.14.0/24	10.x.0.4	0	100	0	64600 i
*>i 198.18.15.0/24	10.x.0.3	0	100	0	64600 ?
*=i 198.18.15.0/24	10.x.0.4	0	100	0	64600 ?
* i 198.19.100.0/24	10.x.0.3	0	300	0	64512 64513 i
*>e 198.19.100.0/24	10.255.101.11	0	300	0	64511 64513 i
* i 198.19.101.0/24	10.x.0.3	0	100	0	64512 64511 64513 i
*>e 198.19.101.0/24	10.255.101.11	0	300	0	64511 64513 i
*>i 198.19.102.0/24	10.x.0.3	0	300	0	64512 64513 i
* e 198.19.102.0/24	10.255.101.11	0	100	0	64511 64512 64513 i

Total number of entries 12

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.

- On Agg1, under the IPv4 address-family, configure the Sw1 (10.x.0.4) as a route-reflector client.

```
Agg1# conf t
Agg1(config)# router bgp 64500
Agg1(config-bgp)# address-family ipv4 unicast
Agg1(config-bgp-ipv4-uc)# neighbor 10.x.0.4 route-reflector-client

BGP Session with this peer will be restarted

Agg1(config-bgp-ipv4-uc)# end
```

- 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

```
Sw1# show bgp ipv4 unicast 198.19.102.0/24

VRF : default
BGP Local AS 64500          BGP Router-id 10.x.0.4

Network      : 198.19.102.0/24      Nexthop      : 10.x.0.3
Peer         : 10.x.0.2              Origin       : IGP
Metric       : 0                    Local Pref   : 300
Weight       : 0                    Calc. Local Pref : 300
Best         : Yes                  Valid        : Yes
Type         : internal             Stale        : No
Originator ID : 10.x.0.3
Aggregator ID :
Aggregator AS :
Atomic Aggregate :
RFD Flaps    : 0                    RFD Penalty  : 0
```

```
AS-Path          : 64512 64513
Cluster List     : 10.x.0.2
Communities      :
Ext-Communities  :
```

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.

```
Agg1# conf t
Agg1(config)# ip route 10.255.21.0/24 10.255.101.11
Agg1(config)# ip route 10.255.22.0/24 10.255.101.11
```

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.

```
Agg1# ping 10.255.21.1 source loopback0 repetitions 2
PING 10.255.21.1 (10.255.21.1) from 10.x.0.2 : 100(128) bytes of data.
108 bytes from 10.255.21.1: icmp_seq=1 ttl=63 time=3.56 ms
108 bytes from 10.255.21.1: icmp_seq=2 ttl=63 time=3.36 ms

--- 10.255.21.1 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1001ms
rtt min/avg/max/mdev = 3.363/3.461/3.560/0.114 ms
```

3. On Agg1, redistribute all static routes into OSPF.

```
Agg1(config)# router ospf 1
Agg1(config-ospf-1)# redistribute static
Agg1(config-ospf-1)# exit
```

4. On Agg2, define the same static routes, but use routerB next hop (10.255.102.12)

```
Agg2# conf t
Agg2(config)# ip route 10.255.21.0/24 10.255.102.12
Agg2(config)# ip route 10.255.22.0/24 10.255.102.12
```

5. Repeat the redistribution on Agg2.

```
Agg2(config)# router ospf 1
Agg2(config-ospf-1)# redistribute static
Agg2(config-ospf-1)# exit
```

6. On Sw1, check the IP routing table and metric for the 10.255.21.0 route.

```
Sw1# show ip route 10.255.21.0/24
```

Displaying ipv4 routes selected for forwarding

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

```
10.255.21.0/24, vrf default
    via 10.x.41.2, [110/25], ospf
    via 10.x.42.3, [110/25], ospf
```

```
Sw1# show ip ospf routes 10.255.21.0/24
```

Codes: i - Intra-area route, I - Inter-area route
E1 - External type-1, E2 - External type-2

OSPF Process ID 1 VRF default, Routing Table for prefixes 10.255.21.0/24

Total Number of Routes : 2

```
10.255.21.0/24      (E2)
    via 10.x.41.2 interface vlan41, cost 25 distance 110
10.255.21.0/24      (E2)
    via 10.x.42.3 interface vlan42, cost 25 distance 110
```

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.

```
Agg1(config)# router ospf 1
Agg1(config-ospf-1)# default-metric 6000
Agg1(config-ospf-1)# exit
```

8. On Sw1, verify the result in the IP routing table and the OSPF routing table.

```
Sw1# show ip route 10.255.21.0/24
```

Displaying ipv4 routes selected for forwarding

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

```
10.255.21.0/24, vrf default
    via 10.x.42.3, [110/25], ospf
```

Sw1# show ip ospf route 10.255.21.0/24

Codes: i - Intra-area route, I - Inter-area route
E1 - External type-1, E2 - External type-2

OSPF Process ID 1 VRF default, Routing Table for prefixes 10.255.21.0/24

Total Number of Routes : 1

```
10.255.21.0/24      (E2)
    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.

Sw1# show ip ospf lsdb external

OSPF Router with ID (10.x.0.4) (Process ID 1 VRF default)

=====

AS External Link State Advertisements

LSID	ADV Router	Age	Seq#	Checksum
10.252.0.0	10.x.0.2	127	0x80000002	0x0000ef3b
10.252.0.0	10.x.0.3	330	0x80000001	0x0000b2e6
10.255.21.0	10.x.0.2	127	0x80000002	0x0000e32f
10.255.21.0	10.x.0.3	330	0x80000001	0x0000a6da
10.255.22.0	10.x.0.2	127	0x80000002	0x0000d839
10.255.22.0	10.x.0.3	330	0x80000001	0x00009be4

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.

```
Agg2(config)# router ospf 1
Agg2(config-ospf-1)# default-metric 7000
Agg2(config-ospf-1)# exit
```

11. On Sw1, verify the result of the Agg2 configuration in the IP routing table and the OSPF routing table.

```
Sw1# show ip route 10.255.21.0/24

Displaying ipv4 routes selected for forwarding

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

10.255.21.0/24, vrf default, tag 0
    via 10.x.41.2, [110/6000], ospf
```

```
Sw1# show ip ospf routes 10.255.21.0/24

Codes: i - Intra-area route, I - Inter-area route
       E1 - External type-1, E2 - External type-2

OSPF Process ID 1 VRF default, Routing Table for prefixes 10.255.21.0/24
-----

Total Number of Routes : 1

10.255.21.0/24      (E2)
    via 10.x.41.2 interface vlan41, cost 6000 distance 110
```

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.

```
(config)# ip prefix-list net21 seq 10 permit 10.255.21.0/24
(config)# ip prefix-list net22 seq 10 permit 10.255.22.0/24
```

14. On Agg1, define a route-map named 'static_to_ospf1' and apply the correct initial metrics.

```
Agg1(config)# route-map static_to_ospf1 permit seq 10
Agg1(config-route-map-static_to_ospf1-10)# match ip address prefix-list net21
Agg1(config-route-map-static_to_ospf1-10)# set metric 1000
Agg1(config-route-map-static_to_ospf1-10)# route-map static_to_ospf1 permit seq 20
Agg1(config-route-map-static_to_ospf1-10)# match ip address prefix-list net22
Agg1(config-route-map-static_to_ospf1-10)# set metric 2000
Agg1(config-route-map-static_to_ospf1-10)# exit
```

15. On Agg1, apply this new route-map to the OSPF static route redistribution

```
Agg1(config)# router ospf 1
Agg1(config-ospf-1)# redistribute static route-map static_to_ospf1
```

16. On Sw1, verify the cost on the 2 destination subnets

```
Sw1# show ip ospf route 10.255.21.0/24
Codes: i - Intra-area route, I - Inter-area route
       E1 - External type-1, E2 - External type-2

OSPF Process ID 1 VRF default, Routing Table for prefixes 10.255.21.0/24
-----

Total Number of Routes : 1

10.255.21.0/24      (E2)
  via 10.x.41.2 interface vlan41, cost 1000 distance 110
```

```
Sw1# show ip ospf route 10.255.22.0/24
Codes: i - Intra-area route, I - Inter-area route
       E1 - External type-1, E2 - External type-2

OSPF Process ID 1 VRF default, Routing Table for prefixes 10.255.22.0/24
-----

Total Number of Routes : 1

10.255.22.0/24      (E2)
  via 10.x.41.2 interface vlan41, cost 2000 distance 110
```

```
Sw1# show ip route
```

```
Displaying ipv4 routes selected for forwarding
----- Output Omitted -----
10.255.21.0/24, vrf default
    via 10.x.41.2, [110/1000], ospf
10.255.22.0/24, vrf default
    via 10.x.41.2, [110/2000], ospf
```

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.

```
Agg2(config)# route-map static_to_ospf1 permit seq 10
Agg2(config-route-map-static_to_ospf1-10)# match ip address prefix-list net21
Agg2(config-route-map-static_to_ospf1-10)# set metric 2000
Agg2(config-route-map-static_to_ospf1-10)# route-map static_to_ospf1 permit seq 20
Agg2(config-route-map-static_to_ospf1-10)# match ip address prefix-list net22
Agg2(config-route-map-static_to_ospf1-10)# set metric 1000
Agg2(config-route-map-static_to_ospf1-10)# exit
```

18. On Agg2, apply the new route-map to the OSPF static route redistribution.

```
Agg2(config-ospf-1)# router ospf 1
Agg2(config-ospf-1)# redistribute static route-map static_to_ospf1
Agg2(config-ospf-1)# exit
```

19. On Sw1, review the result. Both external routes should now be listed with metric 1000, 1 to each Aggregation switch.

```
Sw1# show ip ospf route | begin 10.255
10.255.21.0/24      (E2)
    via 10.x.41.2 interface vlan41, cost 1000 distance 110
10.255.22.0/24      (E2)
    via 10.x.42.3 interface vlan42, cost 1000 distance 110
```

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.

```
Agg1(config)# show route-map
Route-map: static_to_ospf1
  Seq 10, permit,
    Match :
      ip prefix-list          : net21
    Set :
      metric                  : 1000
  Seq 20, permit,
    Match :
      ip prefix-list          : net22
    Set :
      metric                  : 2000
```

```
Agg1(config)# route-map static_to_ospf1 permit seq 20
Agg1(config-route-map-static_to_ospf1-20)# set metric-type external type-1
Agg1(config-route-map-static_to_ospf1-20)# exit
```

21. On Sw1, review the impact on the routes. Compare this with the previous output.

```
Sw1# show ip ospf route | begin 10.255
10.255.21.0/24      (E2)
   via 10.x.41.2 interface vlan41, cost 1000 distance 110
10.255.22.0/24      (E1)
   via 10.x.41.2 interface vlan41, cost 2100 distance 110
```

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.

```
Sw1# show ip ospf interface brief
OSPF Process ID 1 VRF default
=====
Total Number of Interfaces: 5
```

Interface	Area	IP Address/Mask	Cost	State	Status
loopback0	0.0.0.0	10.x.0.4/32	0	Loopback	up
vlan41	0.0.0.0	10.x.41.4/24	100	Point-to-point	up
vlan42	0.0.0.0	10.x.42.4/24	100	Point-to-point	up
vlan46	0.0.0.0	10.x.46.4/24	10	Point-to-point	up
vlan47	0.0.0.0	10.x.47.4/24	100	Point-to-point	up

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.

```
Agg2(config)# route-map static_to_ospf1 permit seq 20
Agg2(config-route-map-static_to_ospf1-20)# set metric-type external type-1
Agg2(config-route-map-static_to_ospf1-20)# exit
```

24. On Sw1, verify that the path to Agg2 is now the best path with a metric of 1100.

```
Sw1# show ip ospf route | begin 10.255
10.255.21.0/24      (E2)
    via 10.x.41.2 interface vlan41, cost 1000 distance 110
10.255.22.0/24      (E1)
    via 10.x.42.3 interface vlan42, cost 1100 distance 110
```

Test failover impact on the route metrics

25. On Sw1, verify the cost on the 2 destination subnets and repeat the output.

```
Sw1# show ip ospf route | begin 10.255
Sw1# repeat
```

26. On Agg1, disable the uplink port 1/1/7.

```
Agg1(config)# interface 1/1/7
Agg1(config-if)# shutdown
```

27. On Sw1, the 10.255.21.0/24 should now have cost 2000.

```
*****
Iteration : 11 Command : show ip ospf route | begin 10.255
*****
10.255.21.0/24      (E2)
    via 10.x.42.3 interface vlan42, cost 2000 distance 110
10.255.22.0/24      (E1)
    via 10.x.42.3 interface vlan42, cost 1100 distance 110
```

28. On Agg1, enable the uplink 1/1/7 again.

```
Agg1(config-if)# no shutdown
```

29. On Agg2, disable the uplink port 1/1/7.

```
Agg2(config)# interface 1/1/7
Agg2(config-if)# shutdown
```

30. On Sw1, the 10.255.22.0/24 route should now have cost 2100.

```
*****
Iteration : 33 Command : show ip ospf route | begin 10.255
*****
10.255.21.0/24      (E2)
   via 10.x.41.2 interface vlan41, cost 1000 distance 110
10.255.22.0/24      (E1)
   via 10.x.41.2 interface vlan41, cost 2100 distance 110
```

31. On Agg2, enable the uplink 1/1/7 again.

```
Agg2(config-if)# no shutdown
```

32. On Sw1, verify that the original costs are shown again, stop the repeat.

```
*****
Iteration : 56 Command : show ip ospf route | begin 10.255
*****
10.255.21.0/24      (E2)
   via 10.x.41.2 interface vlan41, cost 1000 distance 110
10.255.22.0/24      (E1)
   via 10.x.42.3 interface vlan42, cost 1100 distance 110
```

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.

```
Agg1(config)# ip prefix-list 102556 permit 10.255.64.0/23
Agg1(config)# route-map static_to_ospf1 permit seq 30
Agg1(config-route-map-static_to_ospf1-30)# match ip address prefix-list 102556
```

```
Sw1# show ip route | in 10.255.6
10.255.64.0/23, vrf default
```

36. Scenario2: All the routes under the 10.255.64.0/23 prefix should appear in the Sw1 routing table.

```
Agg1(config)# ip prefix-list 102556 seq 10 permit 10.255.64.0/23 ge 23
```

```
Sw1# show ip route | in 10.255.6
10.255.64.0/23, vrf default
10.255.64.0/26, vrf default
10.255.64.1/32, vrf default
10.255.64.2/32, vrf default
10.255.64.64/26, vrf default
10.255.64.128/26, vrf default
10.255.64.192/26, vrf default
10.255.65.0/26, vrf default
10.255.65.1/32, vrf default
10.255.65.2/32, vrf default
10.255.65.64/26, vrf default
10.255.65.128/26, vrf default
10.255.65.192/26, vrf default
```

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.

```
Agg1(config)# ip prefix-list 102556 seq 10 permit 10.255.64.0/23 ge 24
```

```
Sw1# show ip route | in 10.255.6
10.255.64.0/26, vrf default
10.255.64.1/32, vrf default
10.255.64.2/32, vrf default
10.255.64.64/26, vrf default
10.255.64.128/26, vrf default
10.255.64.192/26, vrf default
10.255.65.0/26, vrf default
10.255.65.1/32, vrf default
10.255.65.2/32, vrf default
10.255.65.64/26, vrf default
10.255.65.128/26, vrf default
```

```
10.255.65.192/26, vrf default
```

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.

```
Agg1(config)# ip prefix-list 102556 seq 10 permit 10.255.64.0/23 ge 24 le 31
```

```
Sw1# show ip route | in 10.255.6
```

```
10.255.64.0/26, vrf default  
10.255.64.64/26, vrf default  
10.255.64.128/26, vrf default  
10.255.64.192/26, vrf default  
10.255.65.0/26, vrf default  
10.255.65.64/26, vrf default  
10.255.65.128/26, vrf default  
10.255.65.192/26, vrf default
```

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.

```
Agg1(config)# no ip route 10.255.21.0/24 10.255.101.11
```

2. On Agg2, remove the static route 10.255.21.0/24 to the partner network.

```
Agg2(config)# no ip route 10.255.21.0/24 10.255.102.12
```

Verify current routes and OSPF neighbors

3. On Agg1, verify that there is no 10.255.21.0/24 route in the routing table.

```
Agg1# show ip route 10.255.21.0
No ipv4 routes configured
```

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.

```
Agg1(config)# router bgp 64500
Agg1(config-bgp)# neighbor 10.255.101.11 remote-as 65007
Agg1(config-bgp)# address-family ipv4 unicast
Agg1(config-bgp-ipv4-uc)# neighbor 10.255.101.11 activate
Agg1(config-bgp-ipv4-uc)# exit
```

5. On Agg1, verify the eBGP peering state is established.

```
Agg1(config-bgp)# show bgp ipv4 unicast summary
```

VRF : default
BGP Summary

Local AS	: 64500	BGP Router Identifier	: 10.x.0.2
Peers	: 1	Log Neighbor Changes	: No
Cfg. Hold Time	: 180	Cfg. Keep Alive	: 60

Neighbor	Remote-AS	MsgRcvd	MsgSent	Up/Down	Time	State	AdminStatus
10.255.101.11	65007	3	3	00h:00m:05s	Established	Up	

6. On Agg1, review the BGP IPv4 routing table.

```
Agg1(config-bgp)# show bgp ipv4 unicast
```

Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default
Local Router-ID 10.x.0.2

Network	Nexthop	Metric	LocPrf	Weight	Path
*>e 10.255.0.12/32	10.255.101.11	0	100	0	65007 ?
*>e 10.255.0.13/32	10.255.101.11	0	100	0	65007 ?
*>e 10.255.21.0/24	10.255.101.11	0	100	0	65007 ?
*>e 10.255.22.0/24	10.255.101.11	0	100	0	65007 ?
*>e 10.255.23.0/24	10.255.101.11	0	100	0	65007 ?
*>e 10.255.103.0/24	10.255.101.11	0	100	0	65007 ?
*>e 10.255.104.0/24	10.255.101.11	0	100	0	65007 ?
*>e 10.255.105.0/24	10.255.101.11	0	100	0	65007 ?

Total number of entries 8

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

```
Agg1(config-bgp)# show bgp ipv4 unicast neighbors 10.255.101.11 advertised-routes
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
               i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default
Local Router-ID 10.x.0.2

   Network                Nexthop          Metric      LocPrf      Weight Path
Total number of entries 0
```

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.

```
Agg1(config)# router bgp 64500
Agg1(config-bgp)# address-family ipv4 unicast
Agg1(config-bgp-ipv4-uc)# redistribute ospf
Agg1(config-bgp-ipv4-uc)# exit
```

9. On Agg1, review the BGP IPv4 routing table. You should see all the local OSPF routes in the BGP routing table.

```
Agg1# show bgp ipv4 unicast
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
               i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete
```

VRF : default
Local Router-ID 10.x.0.2

Network	Nexthop	Metric	LocPrf	Weight	Path
*> 10.x.0.3/32	10.x.1.3	0	100	0	?
*> 10.x.0.4/32	10.x.41.4	0	100	0	?
*> 10.x.0.6/32	10.x.40.6	0	100	0	?
*> 10.x.1.0/24	0.0.0.0	0	100	0	?
*> 10.x.11.0/24	0.0.0.0	0	100	0	?
*> 10.x.12.0/24	0.0.0.0	0	100	0	?
*> 10.x.40.0/24	0.0.0.0	0	100	0	?
*> 10.x.41.0/24	0.0.0.0	0	100	0	?
*> 10.x.42.0/24	10.x.41.4	0	100	0	?
*> 10.x.46.0/24	10.x.41.4	0	100	0	?
*> 10.x.47.0/24	10.x.41.4	0	100	0	?
*>e 10.255.0.12/32	10.255.101.11	0	100	0	65007 ?
*>e 10.255.0.13/32	10.255.101.11	0	100	0	65007 ?
*>e 10.255.21.0/24	10.255.101.11	0	100	0	65007 ?
*> 10.255.22.0/24	10.x.40.3	0	100	0	?
* e 10.255.22.0/24	10.255.101.11	0	100	0	65007 ?
*>e 10.255.23.0/24	10.255.101.11	0	100	0	65007 ?
*>e 10.255.103.0/24	10.255.101.11	0	100	0	65007 ?
*>e 10.255.104.0/24	10.255.101.11	0	100	0	65007 ?
*>e 10.255.105.0/24	10.255.101.11	0	100	0	65007 ?
Total number of entries 20					

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.

Agg1(config-bgp)# show bgp ipv4 unicast neighbors 10.255.101.11 advertised-routes

Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default
Local Router-ID 10.x.0.2

Network	Nexthop	Metric	LocPrf	Weight	Path
*>e 10.x.0.3/32	10.255.101.2	0	0	0	64500 ?
*>e 10.x.0.4/32	10.255.101.2	0	0	0	64500 ?
*>e 10.x.0.6/32	10.255.101.2	0	0	0	64500 ?
*>e 10.x.1.0/24	10.255.101.2	0	0	0	64500 ?
*>e 10.x.11.0/24	10.255.101.2	0	0	0	64500 ?
*>e 10.x.12.0/24	10.255.101.2	0	0	0	64500 ?
*>e 10.x.40.0/24	10.255.101.2	0	0	0	64500 ?
*>e 10.x.41.0/24	10.255.101.2	0	0	0	64500 ?
*>e 10.x.42.0/24	10.255.101.2	0	0	0	64500 ?

```
*>e 10.x.46.0/24      10.255.101.2    0          0          0          64500 ?
*>e 10.x.47.0/24      10.255.101.2    0          0          0          64500 ?
*>e 10.255.22.0/24    10.255.101.2    0          0          0          64500 ?
Total number of entries 12
```

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.

```
Sw1(config)# show ip route 10.255.21.0
```

```
No ipv4 routes configured
```

```
Sw1(config)# show ip ospf lsdb external
```

```
OSPF Router with ID (10.x.0.4) (Process ID 1 VRF default)
```

```
=====
```

```
AS External Link State Advertisements
```

```
-----
```

LSID	ADV Router	Age	Seq#	Checksum
10.255.22.0	10.x.0.2	1061	0x80000001	0x0000e26b
10.255.22.0	10.x.0.3	1746	0x80000009	0x00009997
10.255.64.0	10.x.0.2	1061	0x80000001	0x0000e74b
10.255.64.64	10.x.0.2	1061	0x80000001	0x0000658d
10.255.64.128	10.x.0.2	1061	0x80000001	0x0000e2cf
10.255.64.192	10.x.0.2	1061	0x80000001	0x00006012
10.255.65.0	10.x.0.2	1061	0x80000001	0x0000dc55
10.255.65.64	10.x.0.2	1061	0x80000001	0x00005a97
10.255.65.128	10.x.0.2	1061	0x80000001	0x0000d7d9
10.255.65.192	10.x.0.2	1061	0x80000001	0x0000551c

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.

```
Agg1(config)# router ospf 1
```

```
Agg1(config-ospf-1)# redistribute bgp
```

```
Agg1(config-ospf-1)# exit
```

13. On Sw1, review the routing table and the OSPF LSDB for the 10.255.21.0/24 route.

```
Sw1(config)# show ip route 10.255.21.0
```

Displaying ipv4 routes selected for forwarding

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

```
10.255.21.0/24, vrf default, tag 0
      via 10.x.41.2, [110/6000], ospf
```

```
Sw1# show ip ospf route 10.255.21.0/24
```

Codes: i - Intra-area route, I - Inter-area route
E1 - External type-1, E2 - External type-2

OSPF Process ID 1 VRF default, Routing Table for prefixes 10.255.21.0/24

Total Number of Routes : 1

```
10.255.21.0/24      (E2)
      via 10.x.41.2 interface vlan41, cost 6000 distance 110
```

14. On Sw1, attempt to ping and traceroute to the partner network 10.255.21.1

```
Sw1# ping 10.255.21.1 repetitions 2
```

```
PING 10.255.21.1 (10.255.21.1) 100(128) bytes of data.
108 bytes from 10.255.21.1: icmp_seq=1 ttl=62 time=3.13 ms
108 bytes from 10.255.21.1: icmp_seq=2 ttl=62 time=2.72 ms
```

```
--- 10.255.21.1 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 999ms
rtt min/avg/max/mdev = 2.725/2.929/3.133/0.204 ms
```

```
Sw1# traceroute 10.255.21.1
```

```
traceroute to 10.255.21.1 (10.255.21.1), 1 hops min, 30 hops max, 3 sec. timeout, 3 probes
 1  10.x.41.2  0.142ms  0.071ms  0.069ms
 2  10.255.101.11  2.026ms  1.585ms  1.305ms
 3  10.255.21.1  2.331ms  2.078ms  1.980ms
```

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.

```
Agg2(config)# router bgp 64500
Agg2(config-bgp)# neighbor 10.255.102.12 remote-as 65007
Agg2(config-bgp)# address-family ipv4 unicast
Agg2(config-bgp-ipv4-uc)# neighbor 10.255.102.12 activate
Agg2(config-bgp-ipv4-uc)# exit
Agg2(config-bgp)# exit
```

2. Verify the peering is established and that Agg2 has received BGP routes from routerB.

```
Agg2(config)# show bgp ipv4 unicast summary
```

VRF : default
BGP Summary

```
-----
Local AS           : 64500           BGP Router Identifier : 10.x.0.3
Peers              : 1               Log Neighbor Changes  : No
Cfg. Hold Time     : 180            Cfg. Keep Alive       : 60
```

Neighbor	Remote-AS	MsgRcvd	MsgSent	Up/Down	Time	State	AdminStatus
10.255.102.12	65007	5	5	00h:01m:05s	Established	Up	

```
Agg2(config)# show bgp ipv4 unicast
```

Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default
Local Router-ID 10.x.0.3

Network	Nexthop	Metric	LocPrf	Weight	Path
*>e 10.251.0.0/16	10.255.102.12	0	100	0	65007 ?
*>e 10.252.1.0/24	10.255.102.12	0	100	0	65007 ?
*>e 10.255.0.11/32	10.255.102.12	0	100	0	65007 ?
*>e 10.255.0.13/32	10.255.102.12	0	100	0	65007 ?
*>e 10.255.21.0/24	10.255.102.12	0	100	0	65007 ?
*>e 10.255.22.0/24	10.255.102.12	0	100	0	65007 ?
*>e 10.255.23.0/24	10.255.102.12	0	100	0	65007 ?
*>e 10.255.103.0/24	10.255.102.12	0	100	0	65007 ?
*>e 10.255.104.0/24	10.255.102.12	0	100	0	65007 ?
*>e 10.255.105.0/24	10.255.102.12	0	100	0	65007 ?

Total number of entries 10

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.

```
routerA# show ip route 10.255.21.0
```

Displaying ipv4 routes selected for forwarding

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

10.255.21.0/24, vrf default

via 10.255.104.13, [110/125], ospf

4. On Agg1, list the routes that are advertised to routerA. You will review this list again in a few moments.

Agg1(config)# show bgp ipv4 unicast neighbors 10.255.101.11 advertised-routes

Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
i internal, e external S Stale, R Removed

Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default

Local Router-ID 10.x.0.2

Network	NextHop	Metric	LocPrf	Weight	Path
*>e 10.x.0.3/32	10.255.101.2	0	0	0	64500 ?
*>e 10.x.0.4/32	10.255.101.2	0	0	0	64500 ?
*>e 10.x.0.6/32	10.255.101.2	0	0	0	64500 ?
*>e 10.x.1.0/24	10.255.101.2	0	0	0	64500 ?
*>e 10.x.11.0/24	10.255.101.2	0	0	0	64500 ?
*>e 10.x.12.0/24	10.255.101.2	0	0	0	64500 ?
*>e 10.x.40.0/24	10.255.101.2	0	0	0	64500 ?
*>e 10.x.41.0/24	10.255.101.2	0	0	0	64500 ?
*>e 10.x.42.0/24	10.255.101.2	0	0	0	64500 ?
*>e 10.x.46.0/24	10.255.101.2	0	0	0	64500 ?
*>e 10.x.47.0/24	10.255.101.2	0	0	0	64500 ?
*>e 10.255.22.0/24	10.255.101.2	0	0	0	64500 ?

Total number of entries 12

5. On Agg2, review the OSPF LSDB for external routes.

Agg2(config)# show ip ospf lsdb external

OSPF Router with ID (10.x.0.3) (Process ID 1 VRF default)

=====

AS External Link State Advertisements

LSID	ADV Router	Age	Seq#	Checksum
...				
10.255.21.0	10.x.0.2	589	0x80000001	0x00003ee0
...				

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

```
Agg2(config)# router ospf 1
Agg2(config-ospf-1)# redistribute bgp
Agg2(config-ospf-1)# exit
```

Review routing after the redistribution

7. On Agg2, review the OSPF LSDB.

```
Agg2(config)# show ip ospf lsdb external
OSPF Router with ID (10.x.0.3) (Process ID 1 VRF default)
=====
AS External Link State Advertisements
-----
LSID                ADV Router      Age      Seq#           Checksum
-----
...
10.255.21.0         10.x.0.2        973      0x80000001    0x00003ee0
10.255.21.0         10.x.0.3        266      0x80000001    0x00006bc6
...
```

8. On Agg1, review the eBGP advertised routes to routerA

```
Agg1(config)# show bgp ipv4 unicast neighbors 10.255.101.11 advertised-routes
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
               i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default
Local Router-ID 10.x.0.2
```

Network	Nexthop	Metric	LocPrf	Weight	Path
*>e 10.x.0.3/32	10.255.101.2	0	0	0	64500 ?
*>e 10.x.0.4/32	10.255.101.2	0	0	0	64500 ?
*>e 10.x.0.6/32	10.255.101.2	0	0	0	64500 ?
*>e 10.x.1.0/24	10.255.101.2	0	0	0	64500 ?
*>e 10.x.11.0/24	10.255.101.2	0	0	0	64500 ?
*>e 10.x.12.0/24	10.255.101.2	0	0	0	64500 ?
*>e 10.x.40.0/24	10.255.101.2	0	0	0	64500 ?
*>e 10.x.41.0/24	10.255.101.2	0	0	0	64500 ?
*>e 10.x.42.0/24	10.255.101.2	0	0	0	64500 ?

```
*>e 10.x.46.0/24      10.255.101.2    0      0      0      64500 ?
*>e 10.x.47.0/24      10.255.101.2    0      0      0      64500 ?
*>e 10.251.0.0/16     10.255.101.2    0      0      0      64500 ?
*>e 10.252.1.0/24     10.255.101.2    0      0      0      64500 ?
*>e 10.255.0.11/32    10.255.101.2    0      0      0      64500 ?
*>e 10.255.0.13/32    10.255.101.2    0      0      0      64500 ?
*>e 10.255.21.0/24    10.255.101.2    0      0      0      64500 ?
*>e 10.255.22.0/24    10.255.101.2    0      0      0      64500 ?
*>e 10.255.23.0/24    10.255.101.2    0      0      0      64500 ?
*>e 10.255.103.0/24   10.255.101.2    0      0      0      64500 ?
*>e 10.255.104.0/24   10.255.101.2    0      0      0      64500 ?
*>e 10.255.105.0/24   10.255.101.2    0      0      0      64500 ?
Total number of entries 21
```

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

```
routerA# show ip route 10.255.21.0

Displaying ipv4 routes selected for forwarding

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

10.255.21.0/24, vrf default
    via 10.255.101.2, [20/0], bgp
```

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

```
Sw1# ping 10.255.21.1

PING 10.255.21.1 (10.255.21.1) 100(128) bytes of data.
From 10.255.101.11 icmp_seq=1 Time to live exceeded
From 10.255.101.11 icmp_seq=1 Time to live exceeded
From 10.255.101.11 icmp_seq=1 Time to live exceeded
From 10.255.101.11 icmp_seq=1 Time to live exceeded
From 10.255.101.11 icmp_seq=5 Time to live exceeded

--- 10.255.21.1 ping statistics ---
5 packets transmitted, 0 received, +5 errors, 100% packet loss, time 4061ms
```

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.

```
Agg2(config)# route-map bgp_to_ospf1 permit seq 10
Agg2(config-route-map-bgp_to_ospf1-10)# set tag 65007
Agg2(config-route-map-bgp_to_ospf1-10)# exit

Agg2(config)# router ospf 1
Agg2(config-ospf-1)# redistribute bgp route-map bgp_to_ospf1
Agg2(config-ospf-1)# exit
```

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.

```
Agg1(config)# route-map ospf1_to_bgp deny seq 10
Agg1(config-route-map-ospf1_to_bgp-10)# match tag 65007
Agg1(config-route-map-ospf1_to_bgp-10)# exit
Agg1(config)# route-map ospf1_to_bgp permit seq 20
Agg1(config-route-map-ospf1_to_bgp-20)# exit

Agg1(config)# router bgp 64500
Agg1(config-bgp)# address-family ipv4 unicast
Agg1(config-bgp-ipv4-uc)# redistribute ospf route-map ospf1_to_bgp
Agg1(config-bgp-ipv4-uc)# exit
Agg1(config-bgp)# exit
```

13. On Agg1, clear the BGP sessions to apply the changes

```
Agg1(config)# clear bgp *
```

14. On Agg1, review the advertised routes to routerA. Compare this with the previous output of the same command you ran earlier.

```
Agg1(config)# show bgp ipv4 unicast neighbors 10.255.101.11 advertised-routes
```

Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
i internal, e external S Stale, R Removed

Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default

Local Router-ID 10.x.0.2

Network	Nexthop	Metric	LocPrf	Weight	Path
*>e 10.x.0.3/32	10.255.101.2	0	0	0	64500 ?
*>e 10.x.0.4/32	10.255.101.2	0	0	0	64500 ?
*>e 10.x.0.6/32	10.255.101.2	0	0	0	64500 ?
*>e 10.x.1.0/24	10.255.101.2	0	0	0	64500 ?
*>e 10.x.11.0/24	10.255.101.2	0	0	0	64500 ?
*>e 10.x.12.0/24	10.255.101.2	0	0	0	64500 ?
*>e 10.x.40.0/24	10.255.101.2	0	0	0	64500 ?
*>e 10.x.41.0/24	10.255.101.2	0	0	0	64500 ?
*>e 10.x.42.0/24	10.255.101.2	0	0	0	64500 ?
*>e 10.x.46.0/24	10.255.101.2	0	0	0	64500 ?
*>e 10.x.47.0/24	10.255.101.2	0	0	0	64500 ?
*>e 10.255.22.0/24	10.255.101.2	0	0	0	64500 ?

Total number of entries 12

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.

```
Agg2(config)# show bgp ipv4 unicast neighbors 10.255.102.12 advertised-routes
```

Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
i internal, e external S Stale, R Removed

Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default

Local Router-ID 10.x.0.3

Network	Nexthop	Metric	LocPrf	Weight	Path
Total number of entries 0					

16. On Agg2, verify that a simple redistribution would cause the route feedback

```
Agg2(config)# router bgp 64500
Agg2(config-bgp)# address-family ipv4 unicast
Agg2(config-bgp-ipv4-uc)# redistribute ospf
Agg2(config-bgp-ipv4-uc)# exit
Agg2(config-bgp)# exit
```

17. On Agg2, review the advertised routes, this currently includes 10.255.21.0/24. That will be prevented in the next steps.

```
Agg2(config)# show bgp ipv4 unicast neighbors 10.255.102.12 advertised-routes
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
               i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default
Local Router-ID 10.x.0.3

   Network                Nexthop          Metric    LocPrf    Weight Path
...
*>e 10.255.21.0/24        10.255.102.3      0          0         0      64500 ?
...
Total number of entries 27
```

Now prepare the filtering using route tags.

18. On Agg1, mark the routes when they are redistributed into OSPF

```
Agg1(config)# route-map bgp_to_ospf1 permit seq 10
Agg1(config-route-map-bgp_to_ospf1-10)# set tag 65007
Agg1(config-route-map-bgp_to_ospf1-10)# exit
Agg1(config)# router ospf 1
Agg1(config-ospf-1)# redistribute bgp route-map bgp_to_ospf1
Agg1(config-ospf-1)# exit
```

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.

```
Agg2(config)# route-map ospf1_to_bgp deny seq 10
Agg2(config-route-map-ospf1_to_bgp-10)# match tag 65007
Agg2(config-route-map-ospf1_to_bgp-10)# exit
Agg2(config)# route-map ospf1_to_bgp permit seq 20
Agg2(config-route-map-ospf1_to_bgp-20)# exit
Agg2(config)#
Agg2(config)# router bgp 64500
Agg2(config-bgp)# address-family ipv4 unicast
```



```
Agg2(config-bgp-ipv4-uc)# redistribute ospf route-map ospf1_to_bgp
Agg2(config-bgp-ipv4-uc)# exit
Agg2(config-bgp)# exit
```

20. On Agg2, clear the BGP session to apply the route-map changes.

```
Agg2(config)# clear bgp *
```

21. On Agg2, review the advertised routes to routerB.

```
Agg2(config)# show bgp ipv4 unicast neighbors 10.255.102.12 advertised-routes

Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
               i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default
Local Router-ID 10.x.0.3
```

Network	Nexthop	Metric	LocPrf	Weight	Path
*>e 10.x.0.2/32	10.255.102.3	0	0	0	64500 ?
*>e 10.x.0.4/32	10.255.102.3	0	0	0	64500 ?
*>e 10.x.0.6/32	10.255.102.3	0	0	0	64500 ?
*>e 10.x.1.0/24	10.255.102.3	0	0	0	64500 ?
*>e 10.x.11.0/24	10.255.102.3	0	0	0	64500 ?
*>e 10.x.12.0/24	10.255.102.3	0	0	0	64500 ?
*>e 10.x.40.0/24	10.255.102.3	0	0	0	64500 ?
*>e 10.x.41.0/24	10.255.102.3	0	0	0	64500 ?
*>e 10.x.42.0/24	10.255.102.3	0	0	0	64500 ?
*>e 10.x.46.0/24	10.255.102.3	0	0	0	64500 ?
*>e 10.x.47.0/24	10.255.102.3	0	0	0	64500 ?
*>e 10.255.22.0/24	10.255.102.3	0	0	0	64500 ?
*>e 10.255.64.0/26	10.255.102.3	0	0	0	64500 ?
*>e 10.255.64.64/26	10.255.102.3	0	0	0	64500 ?
*>e 10.255.64.128/26	10.255.102.3	0	0	0	64500 ?
*>e 10.255.64.192/26	10.255.102.3	0	0	0	64500 ?
*>e 10.255.65.0/26	10.255.102.3	0	0	0	64500 ?
*>e 10.255.65.64/26	10.255.102.3	0	0	0	64500 ?
*>e 10.255.65.128/26	10.255.102.3	0	0	0	64500 ?
*>e 10.255.65.192/26	10.255.102.3	0	0	0	64500 ?

Total number of entries 20

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.

```
Sw1(config)# show ip route 10.255.21.0
```

Displaying ipv4 routes selected for forwarding

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

```
10.255.21.0/24, vrf default, tag 65007
    via 10.x.41.2, [110/6000], ospf
```

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

```
routerA# show running-config
route-map bgp_to_ospf1 permit seq 10
    set tag 64500
!
router ospf 1
    router-id 10.255.0.11
    redistribute static
    redistribute bgp route-map bgp_to_ospf1
    area 0.0.0.0

route-map ospf1_to_bgp deny seq 10
    match tag 64500
route-map ospf1_to_bgp permit seq 20

router bgp 65007
    bgp fast-external-fallover
    neighbor 10.255.101.2 remote-as 64500
    address-family ipv4 unicast
        neighbor 10.255.101.2 activate
    redistribute ospf route-map ospf1_to_bgp
    exit-address-family
```

24. routerB

```
routerB# show running-config
```

```
route-map bgp_to_ospf1 permit seq 10
    set tag 64500
!
router ospf 1
    router-id 10.255.0.12
    redistribute bgp route-map bgp_to_ospf1
    area 0.0.0.0

route-map ospf1_to_bgp deny seq 10
    match tag 64500
route-map ospf1_to_bgp permit seq 20

router bgp 65007
    bgp fast-external-fallover
    neighbor 10.255.102.3 remote-as 64500
    address-family ipv4 unicast
        neighbor 10.255.102.3 activate
    redistribute ospf route-map ospf1_to_bgp
    exit-address-family
```

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.

Aruba Training-Confidential

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'

```
Sw2# conf t
Sw2(config)# 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

```
interface 1/1/27
vlan trunk allowed 15
```

Aruba Training-Confidential

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.

```
Sw2# show ip ospf neighbors vrf branch
```

```
OSPF Process ID 1 VRF branch
```

```
=====
```

```
Total Number of Neighbors: 2
```

Neighbor ID	Priority	State	Nbr Address	Interface
10.x.0.4	n/a	FULL	10.x.46.4	vlan46
10.x.0.4	n/a	FULL	10.x.47.4	vlan47

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.

```
Sw2# show ip route vrf branch
```

```
Displaying ipv4 routes selected for forwarding
```

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

```
10.x.0.2/32, vrf branch
    via 10.x.46.4, [110/110], ospf
10.x.0.3/32, vrf branch
    via 10.x.46.4, [110/110], ospf
```

```
----- Output Omitted -----
```

Verify VRF Route Isolation

3. On Sw2, review the routing table for VRF iot.

```
Sw2# show ip route vrf iot
```

```
Displaying ipv4 routes selected for forwarding
```

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

```
10.x.61.0/24, vrf iot
    via vlan61, [0/0], connected
10.x.61.1/32, vrf iot
```

```

        via vlan61, [0/0], local
10.x.62.0/24, vrf iot
        via vlan62, [0/0], connected
10.x.62.1/32, vrf iot
        via vlan62, [0/0], local

```

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.

```

Sw2# ping 10.x.62.1 source 10.x.61.1 vrf iot repetitions 2

PING 10.x.62.1 (10.x.62.1) from 10.x.61.1 : 100(128) bytes of data.
108 bytes from 10.x.62.1: icmp_seq=1 ttl=64 time=0.057 ms
108 bytes from 10.x.62.1: icmp_seq=2 ttl=64 time=0.049 ms

--- 10.x.62.1 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1030ms
rtt min/avg/max/mdev = 0.049/0.053/0.057/0.004 ms

```

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.

```

Sw2# ping 10.x.61.1 vrf branch

connect: Network is unreachable

```

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.

```
Sw2(config)# ip route 10.x.61.0/24 vlan61 vrf branch
```

NOTE: AOS-CX will automatically notice that the configured destination interface belongs to another VRF.

2. Review the IP routing table in VRF `branch`.

```
Sw2# show ip route 10.x.61.0/24 vrf branch
Displaying ipv4 routes selected for forwarding

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

10.x.61.0/24, vrf branch
    via vlan61[vrf iot], [1/0], static
```

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.

```
Sw2(config)# ip route 10.x.15.0/24 vlan15 vrf iot
```

4. On Sw2, review the IP routing table of VRF iot. Check the 10.x.15.0/24 route.

```
Sw2# show ip route 10.x.15.0/24 vrf iot
```

Displaying ipv4 routes selected for forwarding

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

```
10.x.15.0/24, vrf iot
      via vlan15[vrf branch], [1/0], static
```

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.

```
Sw2(config)# ip route 10.x.12.0/24 vlan46 vrf iot
```

6. Review the IP routing table of VRF iot. Check the 10.x.12.0/24 route and the destination.

```
Sw2# show ip route 10.x.12.0/24 vrf iot
```

Displaying ipv4 routes selected for forwarding

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

```
10.x.12.0/24, vrf iot
      via vlan46[vrf branch], [1/0], static
```

7. Local test1: On Sw2, ping 10.x.61.1 from the VRF branch using source SVI 15.

```
Sw2# ping 10.x.61.1 source 10.x.15.1 vrf branch repetitions 2
```

```
PING 10.x.61.1 (10.x.61.1) from 10.x.15.1 : 100(128) bytes of data.
108 bytes from 10.x.61.1: icmp_seq=1 ttl=64 time=0.085 ms
108 bytes from 10.x.61.1: icmp_seq=2 ttl=64 time=0.075 ms
```

```
--- 10.x.61.1 ping statistics ---
```

```
2 packets transmitted, 2 received, 0% packet loss, time 1014ms
rtt min/avg/max/mdev = 0.075/0.080/0.085/0.005 ms
```

8. Local test2: On Sw2, ping 10.x.15.1 from the VRF iot using source SVI 61.

```
Sw2# ping 10.x.15.1 source 10.x.61.1 vrf iot repetitions 2
```

```
PING 10.x.15.1 (10.x.15.1) from 10.x.61.1 : 100(128) bytes of data.
108 bytes from 10.x.15.1: icmp_seq=1 ttl=64 time=0.089 ms
108 bytes from 10.x.15.1: icmp_seq=2 ttl=64 time=0.081 ms
```

```
--- 10.x.15.1 ping statistics ---
```

```
2 packets transmitted, 2 received, 0% packet loss, time 1022ms
rtt min/avg/max/mdev = 0.081/0.085/0.089/0.004 ms
```

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.

```
Sw2# ping 10.x.15.1 source 10.x.62.1 vrf iot repetitions 2
PING 10.x.15.1 (10.x.15.1) from 10.x.62.1 : 100(128) bytes of data.

--- 10.x.15.1 ping statistics ---
2 packets transmitted, 0 received, 100% packet loss, time 1018ms
```

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.

```
Agg1# ping 10.x.61.1 source vlan12 repetitions 2
PING 10.x.61.1 (10.x.61.1) from 10.x.12.2 : 100(128) bytes of data.
ping: sendmsg: Network is unreachable
ping: sendmsg: Network is unreachable

--- 10.x.61.1 ping statistics ---
2 packets transmitted, 0 received, 100% packet loss, time 1020ms
```

Q: Did the ping succeed? If not, why?

A: No, check the routing table on Agg1 for route 10.x.61.0.

```
Agg1# show ip route 10.x.61.0
No ipv4 routes configured
```

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.

```
Sw2(config)# router ospf 2 vrf branch
Sw2(config-ospf-2)# redistribute static
```

13. On Agg1, review the routing table, the route 10.x.61.0 should now be in the routing table.

```
Agg1# show ip route 10.x.61.0/24
Displaying ipv4 routes selected for forwarding

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

10.x.61.0/24, vrf default
```

```
via 10.x.41.4, [110/25], ospf
```

14. On Agg1, Attempt to ping 10.x.61.1 using source SVI 12.

```
Agg1# ping 10.x.61.1 source vlan12 repetitions 2
PING 10.x.61.1 (10.x.61.1) from 10.x.12.2 : 100(128) bytes of data.
108 bytes from 10.x.61.1: icmp_seq=1 ttl=61 time=0.230 ms
108 bytes from 10.x.61.1: icmp_seq=2 ttl=61 time=0.193 ms
--- 10.x.61.1 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1056ms
rtt min/avg/max/mdev = 0.193/0.211/0.230/0.023 ms
```

Q: Did this work?

A: Yes, this should work.

15. Attempt to ping 10.x.61.1 using source SVI 1.

```
Agg1# ping 10.x.61.1 source vlan1 repetitions 2
PING 10.x.61.1 (10.x.61.1) from 10.x.1.2 : 100(128) bytes of data.
--- 10.x.61.1 ping statistics ---
2 packets transmitted, 0 received, 100% packet loss, time 1018ms
```

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.

```
Sw2# diagnostic
Sw2# diag utilities tcpdump proto 1 vrf iot
```

17. On Agg1, do another attempt to ping 10.x.61.1 using source SVI 1.

```
Agg1# ping 10.x.61.1 source vlan1 repetitions 2
```

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.

```
tcpdump: listening on any, link-type LINUX_SLL (Linux cooked), capture size 262144 bytes
1 14:47:17.535541 IP (tos 0x0, ttl 62, id 3082, offset 0, flags [DF], proto ICMP (1),
length 128)
10.x.1.2 > TX-Sw2: ICMP echo request, id 14815, seq 1, length 108
2 14:47:17.631171 IP (tos 0xc0, ttl 64, id 51792, offset 0, flags [none], proto ICMP
(1), length 96)
localhost > localhost: ICMP localhost udp port domain unreachable, length 76
IP (tos 0x0, ttl 64, id 52875, offset 0, flags [DF], proto UDP (17), length 68)
localhost.38574 > localhost.domain: 36892+ PTR? 2.1.12.10.in-addr.arpa. (40)
```

```
3 14:47:17.729487 IP (tos 0xc0, ttl 64, id 51803, offset 0, flags [none], proto ICMP (1), length 96)
```

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).

```
Sw2# conf t
Sw2(config)# no ip route 10.x.61.0/24 vlan61 vrf branch
Sw2(config)# no ip route 10.x.12.0/24 vlan46 vrf iot
Sw2(config)# int 1/1/4
Sw2(config-if)# vlan access 12
```

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

```
Sw2(config)# vrf branch
Sw2(config-vrf)# rd 64500:1
```

2. Configure VRF iot with route-distinguisher 64500:2

```
Sw2(config-vrf)# vrf iot
Sw2(config-vrf)# rd 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).

```
Sw2(config)# ip prefix-list branch_leak seq 10 permit 10.x.12.0/24
```

5. Configure a route-map 'bgp_branch_from_ospf' on Sw2 that permits the 'branch_leak' routes.

```
Sw2(config)# route-map bgp_branch_from_ospf permit seq 10
Sw2(config-route-map-bgp_branch_from_ospf-10)# match ip address prefix-list branch_leak
Sw2(config-route-map-bgp_branch_from_ospf-10)# exit
```

6. On Sw2, configure BGP with AS 64500 and enter the vrf branch context.

```
Sw2(config)# router bgp 64500
Sw2(config-bgp)# vrf branch
```

7. Redistribute the OSPF routes using the configured route-map

```
Sw2(config-bgp-vrf)# address-family ipv4 unicast
Sw2(config-bgp-vrf-ipv4-uc)# redistribute ospf route-map bgp_branch_from_ospf
Sw2(config-bgp-vrf-ipv4-uc)# end
```

8. Review the BGP IPv4 routing table on Sw2

```
Sw2# show bgp ipv4 unicast
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default
Local Router-ID 10.x.1.5

   Network          Nexthop          Metric      LocPrf      Weight Path
Total number of entries 0
```

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.

```
Sw2# show bgp vrf branch ipv4 unicast
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : branch
Local Router-ID 10.x.47.5

   Network          Nexthop          Metric      LocPrf      Weight Path
Route Distinguisher: 64500:1
*> 10.x.12.0/24      10.x.46.4          0           100         0           ?
Total number of entries 1
```

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

```
Sw2# show bgp all-vrf all
```

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

```
ip prefix-list iot_leak seq 10 permit 10.x.61.0/24

route-map bgp_iot_from_connected permit seq 10
match ip address prefix-list iot_leak

router bgp 64500
vrf iot
address-family ipv4 unicast
redistribute connected route-map bgp_iot_from_connected
```

12. Verify the BGP VRF iot tables.

```
Sw2# show bgp vrf iot ipv4 unicast

Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
               i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : iot
Local Router-ID 10.x.62.1

   Network          Nexthop          Metric    LocPrf    Weight Path
   Route Distinguisher: 64500:2
*> 10.x.61.0/24      0.0.0.0          0          100        0        ?
Total number of entries 1
```

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.

```
Sw2# show run bgp
router bgp 64500
!
  vrf branch
    address-family ipv4 unicast
      redistribute ospf route-map bgp_branch_from_ospf
    exit-address-family
!
  vrf iot
    address-family ipv4 unicast
      redistribute connected route-map bgp_iot_from_connected
    exit-address-family
```

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

```
Sw2# show bgp vrf branch ipv4 unicast 10.x.12.0/24

VRF : branch
BGP Local AS 64500      BGP Router-id 10.x.47.5
```

```

Network          : 10.x.12.0/24          Nexthop          : 10.x.46.4
Peer             : 0.0.0.0              Origin           : incomplete
----- Output Omitted -----
Communities      :
Ext-Communities  :

```

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.

```

Sw2# conf t
Sw2(config)# vrf branch
Sw2(config-vrf)# address-family ipv4 unicast
Sw2(config-vrf-ipv4-uc)# route-target export 64500:1
Sw2(config-vrf-ipv4-uc)# exit
Sw2(config-vrf)# exit

```

16. Review the BGP route details of the 10.x.12.0/24 route again.

```

Sw2# show bgp vrf branch ipv4 unicast 10.x.12.0/24

VRF : branch
BGP Local AS 64500          BGP Router-id 10.x.47.5

Network          : 10.x.12.0/24          Nexthop          : 10.x.46.4
Peer             : 0.0.0.0              Origin           : incomplete
----- Output Omitted -----
Communities      :
Ext-Communities  : RT: 64500:1

```

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.

```

Sw2# show ip route vrf iot

Displaying ipv4 routes selected for forwarding

```

```
'[x/y]' denotes [distance/metric]
10.x.15.0/24, vrf iot
    via vlan15[vrf branch], [1/0], static
10.x.61.0/24, vrf iot
    via vlan61, [0/0], connected
10.x.61.1/32, vrf iot
    via vlan61, [0/0], local
10.x.62.0/24, vrf iot
    via vlan62, [0/0], connected
10.x.62.1/32, vrf iot
    via vlan62, [0/0], local
```

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

```
Sw2(config)# vrf iot
Sw2(config-vrf)# address-family ipv4 unicast
Sw2(config-vrf-ipv4-uc)# route-target import 64500:1
Sw2(config-vrf-ipv4-uc)# exit
Sw2(config-vrf)# exit
```

19. Review the VRF iot routing table.

```
Sw2# show ip route vrf iot
Displaying ipv4 routes selected for forwarding

'[x/y]' denotes [distance/metric]
10.x.12.0/24, vrf iot
    via 10.x.46.4[vrf branch], [200/0], bgp
10.x.15.0/24, vrf iot
    via vlan15[vrf branch], [1/0], static
10.x.61.0/24, vrf iot
    via vlan61, [0/0], connected
10.x.61.1/32, vrf iot
    via vlan61, [0/0], local
10.x.62.0/24, vrf iot
    via vlan62, [0/0], connected
10.x.62.1/32, vrf iot
    via vlan62, [0/0], local
```

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.

```
Sw2(config)# vrf iot
Sw2(config-vrf)# address-family ipv4 unicast
Sw2(config-vrf-ipv4-uc)# route-target export 64500:2
Sw2(config-vrf-ipv4-uc)# vrf branch
Sw2(config-vrf)# address-family ipv4 unicast
Sw2(config-vrf-ipv4-uc)# route-target import 64500:2
Sw2(config-vrf-ipv4-uc)# end
```

21. Verify by checking BGP VRF iot routes and RT

```
Sw2# show bgp vrf iot ipv4 unicast
Sw2# show bgp vrf iot ipv4 unicast 10.x.61.0/24
```

22. Review the route 10.x.61.0 in the VRF branch IP routing table, check the 'via' field of the route.

```
Sw2# show ip route 10.x.61.0 vrf branch
Displaying ipv4 routes selected for forwarding

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

10.x.61.0/24, vrf branch
    via vlan61[vrf iot], [200/0], bgp
```

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.

```
Agg1# ping 10.x.61.1 source vlan12 repetitions 2
PING 10.x.61.1 (10.x.61.1) from 10.x.12.2 : 100(128) bytes of data.
ping: sendmsg: Network is unreachable
ping: sendmsg: Network is unreachable

--- 10.x.61.1 ping statistics ---
2 packets transmitted, 0 received, 100% packet loss, time 1025ms
```

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.

```
Sw2# conf t
Sw2(config)# router ospf 2 vrf branch
Sw2(config-ospf-2)# redistribute bgp
Sw2(config-ospf-2)# exit
```

25. Verify: On Agg1, review the IP routing table. Make sure the 10.x.61.0/24 route is listed.

```
Agg1# show ip route 10.x.61.0
Displaying ipv4 routes selected for forwarding

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

10.x.61.0/24, vrf default
    via 10.x.41.4, [110/25], ospf
```

26. On Agg1, attempt to ping the 10.x.61.1 IP address using source SVI 12.

```
Agg1# ping 10.x.61.1 source vlan12 repetitions 2
PING 10.x.61.1 (10.x.61.1) from 10.x.12.2 : 100(128) bytes of data.
108 bytes from 10.x.61.1: icmp_seq=1 ttl=61 time=0.235 ms
108 bytes from 10.x.61.1: icmp_seq=2 ttl=61 time=0.196 ms

--- 10.x.61.1 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1042ms
rtt min/avg/max/mdev = 0.196/0.215/0.235/0.024 ms
```

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)

```
Sw2# show run vrf
vrf branch
  rd 64500:1
  address-family ipv4 unicast
    route-target export 64500:1
    route-target import 64500:2
  exit-address-family
vrf iot
  rd 64500:2
  address-family ipv4 unicast
    route-target export 64500:2
    route-target import 64500:1
  exit-address-family
```

```
Sw2(config)# vrf iot
Sw2(config-vrf)# address-family ipv4 unicast
Sw2(config-vrf-ipv4-uc)# no route-target import 64500:1
```

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.

```
Sw2(config-vrf-ipv4-uc)# route-target import 64500:101
```

29. Review the VRF configuration on Sw2

```
Sw2# show run vrf
vrf branch
  rd 64500:1
  address-family ipv4 unicast
    route-target export 64500:1
    route-target import 64500:2
  exit-address-family
vrf iot
  rd 64500:2
  address-family ipv4 unicast
    route-target export 64500:2
    route-target import 64500:101
  exit-address-family
```

30. Review the Sw2's VRF iot routing table, verify that the route to 10.x.12.0/24 has disappeared.

```
Sw2# show ip route 10.x.12.0/24 vrf iot
No ipv4 routes configured
```

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'.

```
Sw2# show route-map
Route-map: bgp_branch_from_ospf
  Seq 10, permit,
  Match :
    ip prefix-list          : branch_leak
  Set :

Route-map: bgp_iot_from_connected
  Seq 10, permit,
  Match :
    ip prefix-list          : iot_leak
  Set :
```

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".

```
Sw2# conf t
Sw2(config)# route-map bgp_branch_from_ospf permit seq 10
Sw2(config-route-map-bgp_branch_from_ospf-10)# set extcommunity rt "64500:101"
Sw2(config-route-map-bgp_branch_from_ospf-10)# exit
```

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.

```
Sw2# show route-map
Route-map: bgp_branch_from_ospf
  Seq 10, permit,
  Match :
    ip prefix-list          : branch_leak
  Set :
    extcommunity rt         : "64500:101"

Route-map: bgp_iot_from_connected
  Seq 10, permit,
  Match :
    ip prefix-list          : iot_leak
  Set :
```

34. Review the details of the BGP VRF branch 10.x.12.0/24 entry.

```
Sw2# show bgp vrf branch ipv4 unicast 10.x.12.0/24
VRF : branch
BGP Local AS 64500          BGP Router-id 10.x.47.5

Network      : 10.x.12.0/24      Nexthop      : 10.x.46.4
Peer         : 0.0.0.0           Origin       : incomplete
Metric       : 0                 Local Pref   : 100
Weight       : 0                 Calc. Local Pref : 100
Best         : Yes               Valid        : Yes
Type         : external          Stale        : No
Originator ID : 0.0.0.0
Aggregator ID :
Aggregator AS :
Atomic Aggregate :
RFD Flaps     : 0                RFD Penalty  : 0

AS-Path       :
Cluster List  :
Communities   :
Ext-Communities : RT: 64500:1 RT: 64500:101
```

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.

```
Sw2# show ip route vrf iot
Displaying ipv4 routes selected for forwarding

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

10.x.12.0/24, vrf iot
    via 10.x.46.4[vrf branch], [200/0], bgp
10.x.15.0/24, vrf iot
    via vlan15[vrf branch], [1/0], static
10.x.61.0/24, vrf iot
    via vlan61, [0/0], connected
10.x.61.1/32, vrf iot
    via vlan61, [0/0], local
10.x.62.0/24, vrf iot
    via vlan62, [0/0], connected
10.x.62.1/32, vrf iot
    via vlan62, [0/0], local
```

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

```
Agg1# ping 10.x.61.1 source vlan12 repetitions 2
PING 10.x.61.1 (10.x.61.1) from 10.x.12.2 : 100(128) bytes of data.
108 bytes from 10.x.61.1: icmp_seq=1 ttl=61 time=0.242 ms
108 bytes from 10.x.61.1: icmp_seq=2 ttl=61 time=0.227 ms

--- 10.x.61.1 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1025ms
rtt min/avg/max/mdev = 0.227/0.234/0.242/0.017 ms
```

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

```
Sw1# show ip ospf neighbors
```

```
OSPF Process ID 1 VRF default
```

```
=====
```

```
Total Number of Neighbors: 4
```

Neighbor ID	Priority	State	Nbr Address	Interface
10.x.200.2	n/a	FULL	10.x.241.2	vlan241
10.x.200.3	n/a	FULL	10.x.242.3	vlan242

```
--- <output omitted> ---
```

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.

```
Agg1# show ip route ospf
```

```
Displaying ipv4 routes selected for forwarding
```

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

```
10.x.0.3/32, vrf default
```

```
via 10.x.1.3, [110/100], ospf
```

```
10.x.0.4/32, vrf default
```

```

        via 10.x.41.4, [110/100], ospf
10.x.42.0/24, vrf default
        via 10.x.1.3, [110/200], ospf
        via 10.x.41.4, [110/200], ospf
10.x.200.2/32, vrf default
        via 10.x.41.4, [110/200], ospf
10.x.200.3/32, vrf default
        via 10.x.41.4, [110/200], ospf
10.x.212.0/24, vrf default
        via 10.x.41.4, [110/300], ospf
10.x.241.0/24, vrf default
        via 10.x.41.4, [110/200], ospf
10.x.242.0/24, vrf default
        via 10.x.41.4, [110/200], ospf

```

4. On PC3, verify the IP address. It should have an IP address in the 10.x.11.0/24 subnet.

```

C:\Users\student> ipconfig
----- Output Omitted -----
Ethernet adapter Lab NIC - 6300-A Port 3:

    Connection-specific DNS Suffix  . : arubatraining.com
    IPv4 Address. . . . . : 10.x.11.103
    Subnet Mask . . . . . : 255.255.255.0
    Default Gateway . . . . . : 10.x.11.1

```

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)

```

C:\Users\student> ipconfig /renew
----- Output Omitted -----
Ethernet adapter Lab NIC 6300-B Port 4:

    Connection-specific DNS Suffix  . :
    IPv4 Address. . . . . : 10.x.212.z
    Subnet Mask . . . . . : 255.255.255.0
    Default Gateway . . . . . : 10.x.212.1

```

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):

```

C:\Users\student> tracert -d 10.x.212.z
Tracing route to 10.x.212.z over a maximum of 30 hops
 1  <1 ms  <1 ms  <1 ms  10.x.11.3
 2  <1 ms  <1 ms  <1 ms  10.x.41.4

```



```
3    <1 ms    <1 ms    <1 ms  10.x.241.2
4    <1 ms    <1 ms    <1 ms  10.x.212.z
Trace complete.
```

- 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.

```
Sw1# conf t
Sw1(config)# interface loopback 0
Sw1(config-loopback-if)# ip pim-sparse enable
Sw1(config-loopback-if)# router pim
Sw1(config-pim)# bsr-candidate priority 200
Sw1(config-pim)# bsr-candidate source-ip-interface loopback0
Sw1(config-pim)# enable
Sw1(config-pim)# exit
```

2. Enable all Layer3 interfaces for multicast on Sw1.

```
Sw1(config)# interface vlan41
Sw1(config-if-vlan)# ip pim-sparse enable
Sw1(config-if-vlan)# interface vlan42
Sw1(config-if-vlan)# ip pim-sparse enable
Sw1(config-if-vlan)#
Sw1(config-if-vlan)# interface vlan241
Sw1(config-if-vlan)# ip pim-sparse enable
Sw1(config-if-vlan)# interface vlan242
Sw1(config-if-vlan)# ip pim-sparse enable
Sw1(config-if-vlan)# exit
```

3. Review the enabled PIM interfaces on Sw1

```
Sw1# show ip pim interface
```

```
Sw1# show ip pim interface brief
```

```
-----
VRF : default                                Total number of interfaces : 5
-----
Interface      IP Address      DR Address      Neighbor      Mode      VSX
count                                     Role
-----
vlan241         10.x.241.4/24    10.x.241.4      0             sparse    N/A
loopback0       10.x.0.4/32      N/A             N/A           sparse    N/A
```

vlan242	10.x.242.4/24	10.x.242.4	0	sparse	N/A
vlan41	10.x.41.4/24	10.x.41.4	0	sparse	N/A
vlan42	10.x.42.4/24	10.x.42.4	0	sparse	N/A

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)

```
Agg1# conf t
Agg1(config)# interface loopback 0
Agg1(config-loopback-if)# ip pim-sparse enable
Agg1(config-loopback-if)# router pim
Agg1(config-pim)# enable
Agg1(config-pim)# interface vlan1
Agg1(config-if-vlan)# ip pim-sparse enable
Agg1(config-if-vlan)# interface vlan 11
Agg1(config-if-vlan)# ip pim-sparse enable
Agg1(config-if-vlan)# interface vlan 41
Agg1(config-if-vlan)# ip pim-sparse enable
Agg1(config-if-vlan)# exit
```

5. On Agg1, verify the enabled PIM interfaces

```
Agg1# show ip pim interface brief
```

VRF : default			Total number of interfaces : 4		
Interface	IP Address	DR Address	Neighbor count	Mode	VSX Role
vlan1	10.x.1.2/24	10.x.1.2	0	sparse	N/A
vlan41	10.x.41.2/24	10.x.41.4	1	sparse	N/A
loopback0	10.x.0.2/32	N/A	N/A	sparse	N/A
vlan11	10.x.11.2/24	10.x.11.2	0	sparse	N/A

6. On Agg2(default), enable the interfaces and verify the enabled interfaces.

```
Agg2# conf t
Agg2(config)# interface loopback 0
Agg2(config-loopback-if)# ip pim-sparse enable
Agg2(config-loopback-if)# router pim
Agg2(config-pim)# enable
```

```
Agg2(config-pim)# interface vlan1
Agg2(config-if-vlan)# ip pim-sparse enable
Agg2(config-if-vlan)# interface vlan 11
Agg2(config-if-vlan)# ip pim-sparse enable
Agg2(config-if-vlan)# interface vlan 42
Agg2(config-if-vlan)# ip pim-sparse enable
Agg2(config-if-vlan)# exit
```

```
Agg2(config)# show ip pim interface brief
```

VRF : default			Total number of interfaces : 4		
Interface	IP Address	DR Address	Neighbor count	Mode	VSX Role
vlan11	10.x.11.3/24	10.x.11.3	0	sparse	N/A
vlan42	10.x.42.3/24	10.x.42.3	0	sparse	N/A
loopback0	10.x.0.3/32	N/A	N/A	sparse	N/A
vlan1	10.x.1.3/24	10.x.1.3	0	sparse	N/A

7. On Agg1(b2) , enable the interfaces and verify the enabled interfaces.

```
Agg1(config)# interface loopback 2
Agg1(config-loopback-if)# ip pim-sparse enable
Agg1(config-loopback-if)# router pim vrf b2
Agg1(config-pim)# enable
Agg1(config-pim)# interface vlan 212
Agg1(config-if-vlan)# ip pim-sparse enable
Agg1(config-if-vlan)# interface vlan 241
Agg1(config-if-vlan)# ip pim-sparse enable
Agg1(config-if-vlan)# exit
```

```
Agg1(config)# show ip pim interface brief vrf b2
```

VRF : b2			Total number of interfaces : 3		
Interface	IP Address	DR Address	Neighbor count	Mode	VSX Role
loopback2	10.x.200.2/32	N/A	N/A	sparse	N/A
vlan212	10.x.212.2/24	10.x.212.2	0	sparse	N/A
vlan241	10.x.241.2/24	10.x.241.2	0	sparse	N/A

8. On Agg2(b2), enable the interfaces and verify the enabled interfaces.

```
Agg2(config)# interface loopback 2
Agg2(config-loopback-if)# ip pim-sparse enable
Agg2(config-loopback-if)# router pim vrf b2
```

```
Agg2(config-pim)# enable
Agg2(config-pim)# interface vlan 212
Agg2(config-if-vlan)# ip pim-sparse enable
Agg2(config-if-vlan)# interface vlan 242
Agg2(config-if-vlan)# ip pim-sparse enable
Agg2(config-if-vlan)# exit
```

```
Agg2(config)# show ip pim interface brief vrf b2
```

```
-----
VRF : b2                                     Total number of interfaces : 3
-----
Interface      IP Address      DR Address      Neighbor      Mode      VSX
count                                     Role
-----
loopback2      10.x.200.3/32    N/A             N/A           sparse    N/A
vlan212        10.x.212.3/24    10.x.212.3      0             sparse    N/A
vlan242        10.x.242.3/24    10.x.242.3      0             sparse    N/A
-----
```

9. On Sw1, verify you have 4 PIM neighbors.

```
Sw1# show ip pim neighbor brief
```

```
-----
VRF : default                               Total number of neighbors : 4
-----
Interface      Neighbor      Uptime      Expires      DR      Hold
Time           Secondary Address      (HH:MM:SS)
(HH:MM:SS)     (IPV4)        (IPV4)
-----
vlan241        10.x.241.2    00:03:45    00:01:32    1        00:01:45
Null
vlan242        10.x.242.3    00:02:32    00:01:44    1        00:01:45
Null
vlan41         10.x.41.2     00:07:43    00:01:35    1        00:01:45
Null
vlan42         10.x.42.3     00:05:51    00:01:27    1        00:01:45
Null
-----
```

10. On Agg1(default) review the elected BSR

```
Agg1# show ip pim bsr
```

```
Status and Counters- PIM-SM Bootstrap Router Information
VRF : default
E-BSR Address : 10.x.0.4
E-BSR Priority : 200
```

```
E-BSR Hash Mask Length : 30
E-BSR Up Time           : 7 mins 39 secs
Next Bootstrap Message  : 1 mins 31 secs
C-BSR Admin Status      : This system is not a Candidate-BSR
C-RP Admin Status       : This system is not a Candidate-RP
```

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

```
Agg1# show ip pim bsr vrf b2
Status and Counters- PIM-SM Bootstrap Router Information
VRF                  : b2
E-BSR Address        : 10.x.0.4
E-BSR Priority        : 200
E-BSR Hash Mask Length : 30
E-BSR Up Time        : 4 mins 53 secs
Next Bootstrap Message : 1 mins 17 secs
C-BSR Admin Status    : This system is not a Candidate-BSR
C-RP Admin Status     : This system is not a Candidate-RP
```

```
Agg1# show ip pim bsr all-vrfs
Status and Counters- PIM-SM Bootstrap Router Information
VRF                  : b2
E-BSR Address        : 10.x.0.4
E-BSR Priority        : 200
E-BSR Hash Mask Length : 30
E-BSR Up Time        : 5 mins 27 secs
Next Bootstrap Message : 1 mins 43 secs
C-BSR Admin Status    : This system is not a Candidate-BSR
C-RP Admin Status     : This system is not a Candidate-RP

VRF                  : default
E-BSR Address        : 10.x.0.4
E-BSR Priority        : 200
E-BSR Hash Mask Length : 30
E-BSR Up Time        : 9 mins 27 secs
Next Bootstrap Message : 1 mins 43 secs
C-BSR Admin Status    : This system is not a Candidate-BSR
C-RP Admin Status     : This system is not a Candidate-RP
```

Configure Sw1 as C-RP

12. On Agg1(default), review the current RP list.

```
Agg1# show ip pim rp-set
```

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.

```
Sw1(config)# router pim
Sw1(config-pim)# rp-candidate source-ip-interface loopback0
Sw1(config-pim)# rp-candidate priority 210
Sw1(config-pim)# exit
```

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.

```
Agg1# show ip pim rp-set
Agg1# repeat
```

15. Verify that Sw1 is listed as an RP.

```
*****
Iteration : 6 Command : show ip pim rp-set
*****

VRF: default

Status and Counters - PIM-SM Learned RP-Set Information
Group Address      Group Mask      RP Address      Hold Time      Expire Time
-----
224.0.0.0          240.0.0.0      10.x.0.4        150            150
```

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.

```
Sw1# show ip multicast summary pim-sm

Total number of PIM SM enabled VRFs      : 1
Total number of PIM SM interfaces         : 5
```

```
Total number of PIM SM neighbors : 4
Total number of PIM A/A enabled VRF : 0
Total number of PIM A/A enabled interfaces : 0
Total number of PIM SM Mroutes in route state : 0
Total number of PIM SM Mroutes in bridge state : 0
Total number of local multicast flows registered in this RP : 0
Total number of PIMv6 SM enabled VRFs : 0
Total number of PIMv6 SM interfaces : 0
Total number of PIMv6 SM neighbors : 0
Total number of PIMv6 A/A enabled VRF : 0
Total number of PIMv6 A/A enabled interfaces : 0
Total number of PIMv6 SM Mroutes in route state : 0
Total number of PIMv6 SM Mroutes in bridge state : 0
----- Output Omitted -----
```

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).

```
Agg1# conf t
Agg1(config)# interface vlan 11
Agg1(config-if-vlan)# ip igmp enable
Agg1(config-if-vlan)# interface vlan 212
Agg1(config-if-vlan)# ip igmp enable
Agg1(config-if-vlan)# end
```

2. Repeat this configuration on Agg2.

```
Agg2# conf t
Agg2(config)# interface vlan 11
Agg2(config-if-vlan)# ip igmp enable
Agg2(config-if-vlan)# interface vlan 212
Agg2(config-if-vlan)# ip igmp enable
Agg2(config-if-vlan)# end
```

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.

```
Agg1# show ip igmp all-vrfs
Agg1# repeat

*****
Iteration : 3 Command : show ip igmp all-vrfs
*****

VRF Name      : default
Interface     : vlan11
IGMP Configured Version : 3
IGMP Operating Version  : 3
Querier State      : Querier
Querier IP [this switch] : 10.x.11.1
Querier Uptime      : 1m 52s
Querier Expiration Time : 1m 18s
```

IGMP Snoop Enabled on VLAN : False

Active Group Address	Vers	Mode	Uptime	Expires
239.255.255.250	3	EXC	1m 46s	3m 41s

VRF Name : b2
Interface : vlan212
IGMP Configured Version : 3
IGMP Operating Version : 3
Querier State : Querier
Querier IP [this switch] : 10.x.212.2
Querier Uptime : 1m 53s
Querier Expiration Time : 1m 17s
IGMP Snoop Enabled on VLAN : False

Active Group Address	Vers	Mode	Uptime	Expires
239.255.255.250	3	EXC	1m 52s	3m 34s

4. Review the IGMP querier state on Agg2 as well.

Agg2# show ip igmp all-vrfs

Agg2# repeat

Iteration : 3 Command : show ip igmp all-vrfs

VRF Name : default
Interface : vlan11
IGMP Configured Version : 3
IGMP Operating Version : 3
Querier State : Non-Querier
Querier IP [this switch] : 10.x.11.1
Querier Uptime : 2m 2s
Querier Expiration Time : 2m 14s
IGMP Snoop Enabled on VLAN : False

Active Group Address	Vers	Mode	Uptime	Expires
239.255.255.250	3	EXC	2m 15s	2m 28s

VRF Name : b2
Interface : vlan212
IGMP Configured Version : 3
IGMP Operating Version : 3
Querier State : Non-Querier
Querier IP : 10.x.212.2
Querier Uptime : 2m 3s
Querier Expiration Time : 2m 13s
IGMP Snoop Enabled on VLAN : False

Active Group Address	Vers	Mode	Uptime	Expires

239.255.255.250	3	EXC	2m 16s	2m 19s

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.

- On Agg1, review the current configuration of SVI 11 (so you can copy/paste the active gateway commands in the next step)

```
Agg1# show run int vlan11
interface vlan11
  vsx-sync active-gateways policies
  ip address 10.x.11.1/24
  active-gateway ip mac 02:02:00:00:00:00
  active-gateway ip 10.x.11.1
  13-counters
  ip ospf 1 area 0.0.0.0
  ip ospf passive
  ip helper-address 10.252.1.21
  ip pim-sparse enable
  ip igmp enable
  exit
```

- On Agg1, disable Active Gateway IP on the SVI 11 and apply the new IP address.

```
Agg1# conf t
Agg1(config)# interface vlan 11
Agg1(config-if-vlan)# no active-gateway ip
Agg1(config-if-vlan)# ip address 10.x.11.2/24
```

```
Agg1(config-if-vlan)# exit
```

7. On Agg2, apply the new IP address on SVI 11

```
Agg2# conf t
Agg2(config)# interface vlan 11
Agg2(config-if-vlan)# ip address 10.x.11.3/24
Agg2(config-if-vlan)# exit
```

8. On Agg1, apply the active-gateway configuration

```
Agg1(config)# interface vlan 11
Agg1(config-if-vlan)# active-gateway ip mac 02:00:00:00:00:01
Agg1(config-if-vlan)# active-gateway ip 10.x.11.1
Agg1(config-if-vlan)# exit
```

9. On Agg1, review the IGMP Querier state is now 'Querier' for both SVI 11 and SVI 242.

```
Agg1# show ip igmp all-vrfs
Agg1# repeat
*****
Iteration : 3 Command : show ip igmp all-vrfs
*****
VRF Name      : default
Interface     : vlan11
IGMP Configured Version : 3
IGMP Operating Version  : 3
Querier State      : Querier
Querier IP [this switch] : 10.x.11.2
Querier Uptime      : 8m 34s
Querier Expiration Time : 0m 54s
IGMP Snoop Enabled on VLAN : False

Active Group Address  Vers Mode Uptime    Expires
-----
239.255.255.250      3    EXC  8m 27s    3m 14s

VRF Name      : b2
Interface     : vlan212
IGMP Configured Version : 3
IGMP Operating Version  : 3
Querier State      : Querier
Querier IP [this switch] : 10.x.212.2
Querier Uptime      : 8m 35s
Querier Expiration Time : 0m 53s
IGMP Snoop Enabled on VLAN : False

Active Group Address  Vers Mode Uptime    Expires
-----
239.255.255.250      3    EXC  8m 34s    3m 10s
```

10. On Agg2, review the IGMP Querier state is now 'Non-Querier' for both SVI 11 and SVI 242.

```

Agg2# show ip igmp all-vrfs
Agg2# repeat
*****
Iteration : 3 Command : show ip igmp all-vrfs
*****
VRF Name : default
Interface : vlan11
IGMP Configured Version : 3
IGMP Operating Version : 3
Querier State : Non-Querier
Querier IP : 10.x.11.2
Querier Uptime : 13m 58s
Querier Expiration Time : 2m 53s
IGMP Snoop Enabled on VLAN : False

Active Group Address  Vers Mode Uptime    Expires
-----
239.255.255.250      3    EXC  14m 11s   3m 8s

VRF Name : b2
Interface : vlan212
IGMP Configured Version : 3
IGMP Operating Version : 3
Querier State : Non-Querier
Querier IP : 10.x.212.2
Querier Uptime : 13m 59s
Querier Expiration Time : 2m 52s
IGMP Snoop Enabled on VLAN : False

Active Group Address  Vers Mode Uptime    Expires
-----
239.255.255.250      3    EXC  14m 15s   3m 1s

```

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.

```

Agg1# show ip igmp snooping
IGMP Snooping Protocol Info

Total VLANs with IGMP enabled : 0
Current count of multicast groups joined : 0

IGMP Drop Unknown Multicast : Disabled

```

12. On Agg1, enable IGMP Snooping on VLANs 11 and 212.

TIP: You can use a VLAN range command for the IGMP snooping configuration.

```
Agg1(config)# vlan 11,212
Agg1(config-vlan-<11,212>)# ip igmp snooping enable
Agg1(config-vlan-<11,212>)# exit
```

13. On Agg1, verify the IGMP snooping status again, you should see that VLAN 11 and 242 are now enabled for IGMP snooping.

```
Agg1# show ip igmp snooping
IGMP Snooping Protocol Info

Total VLANs with IGMP enabled      : 2
Current count of multicast groups joined : 2

IGMP Drop Unknown Multicast        : Global
VLAN ID : 11
VLAN Name : VLAN11
IGMP Configured Version : 3
IGMP Operating Version : 3
Querier Address [this switch] : 10.x.11.2
Querier Port :
Querier UpTime :0m 26s
Querier Expiration Time :0m 6s

Active Group Address  Tracking  Vers Mode Uptime    Expires
-----
239.255.255.250      Filter   3    EXC  0m 22s    3m 58s

VLAN ID : 212
VLAN Name : VLAN212
IGMP Configured Version : 3
IGMP Operating Version : 3
Querier Address [this switch] : 10.x.212.2
Querier Port :
Querier UpTime :0m 29s
Querier Expiration Time :0m 3s

Active Group Address  Tracking  Vers Mode Uptime    Expires
-----
239.255.255.250      Filter   3    EXC  0m 21s    4m 0s
```

14. On Agg1, review the IGMP Querier status, it should also show that IGMP Snooping is enabled for the VLAN:

```
Agg1# show ip igmp interface vlan 11
IGMP Configured Version : 3
IGMP Operating Version : 3
Querier State           : Querier
Querier IP [this switch] : 10.x.11.2
```

```
Querier Uptime          : 1m 14s
Querier Expiration Time : 1m 56s
IGMP Snoop Enabled on VLAN : True
```

Active Group Address	Vers	Mode	Uptime	Expires
239.255.255.250	3	EXC	1m 9s	4m 11s

15. On Agg2, enable IGMP Snooping on VLANs 11,212 and verify the status.

```
Agg2(config)# vlan 11,212
Agg2(config-vlan-<11,212>)# ip igmp snooping enable
```

```
Agg2(config-vlan-<11,212>)# show ip igmp snooping
```

IGMP Snooping Protocol Info

```
Total VLANs with IGMP enabled      : 2
Current count of multicast groups joined : 1
```

```
IGMP Drop Unknown Multicast      : Global
VLAN ID : 11
VLAN Name : VLAN11
IGMP Configured Version : 3
IGMP Operating Version : 3
Querier Address : 10.x.11.2
Querier Port : lag256
Querier UpTime :1m 40s
Querier Expiration Time :3m 36s
```

Active Group Address	Tracking	Vers	Mode	Uptime	Expires
239.255.255.250	Filter	3	EXC	1m 38s	3m 42s

```
VLAN ID : 212
VLAN Name : VLAN212
IGMP Configured Version : 3
IGMP Operating Version : 3
Querier Address [this switch] : 10.x.212.3
Querier Port :
Querier UpTime :0m 2s
Querier Expiration Time :0m 30s
```

16. On Sw2, enable IGMP Snooping on VLANs 11,212 and verify the status.

```
Sw2# conf t
Sw2(config)# vlan 11,212
Sw2(config-vlan-<11,212>)# ip igmp snooping enable
```

```
Sw2(config-vlan-<11,212>)# show ip igmp snooping
```

IGMP Snooping Protocol Info

```
Total VLANs with IGMP enabled      : 2
Current count of multicast groups joined : 0

IGMP Drop Unknown Multicast      : Global
VLAN ID : 11
VLAN Name : VLAN11
IGMP Configured Version : 3
IGMP Operating Version :
Querier Address :
Querier Port :
Querier UpTime :
Querier Expiration Time :
```

```
VLAN ID : 212
VLAN Name : VLAN212
IGMP Configured Version : 3
IGMP Operating Version :
Querier Address :
Querier Port :
Querier UpTime :
Querier Expiration Time :
```

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.

```
Sw2# show ip igmp snooping
Sw2# repeat

*****
Iteration : 7 Command : show ip igmp snooping
*****

IGMP Snooping Protocol Info

Total VLANs with IGMP enabled      : 2
Current count of multicast groups joined : 2

IGMP Drop Unknown Multicast      : Global
VLAN ID : 11
VLAN Name : VLAN11
IGMP Configured Version : 3
IGMP Operating Version : 3
Querier Address : 10.x.11.2
Querier Port : lag255
Querier UpTime : 0m 17s
Querier Expiration Time : 3m 58s
```

Active Group Address	Tracking	Vers	Mode	Uptime	Expires
239.255.255.250	Filter	3	EXC	2m 16s	4m 10s


```
VLAN ID : 212
VLAN Name : VLAN212
IGMP Configured Version : 3
IGMP Operating Version : 3
Querier Address : 10.x.212.2
Querier Port : lag255
Querier UpTime : 0m 18s
Querier Expiration Time : 3m 57s
```

Active Group Address	Tracking	Vers	Mode	Uptime	Expires
239.255.255.250	Filter	3	EXC	0m 13s	4m 8s

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

- 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.

```
Sw1# show ip pim interface loopback0 counters
```

```
Interface : loopback0
VRF       : default
```

```
Tx Counters :
```

```

Hello                                0
State Refresh                        0
Join/Prune                          0
RP Advertisement                     0
Graft                               0
Graft Ack                           0
Assert                              0
Bsm                                  0
Register                             0
Register Stop                        0

```

Rx Counters :

```

Hello                                0
State Refresh                        0
Join/Prune                          0
RP Advertisement                     0
Graft                               0
Graft Ack                           0
Assert                              0
Bsm                                  0
Register                             0
Register Stop                        0
Register Drops(Register ACL hitcount) 0
Join/Prune Drops(RP ACL hitcount)    0

```

Rx Drop Counters :

```

Hello                                0
State Refresh                        0
Join/Prune                          0
RP Advertisement                     0
Graft                               0
Graft Ack                           0
Assert                              0
Bsm                                  0

```

19. On Sw1, review the multicast summary, look for the number of multicast flows on this RP.

```
Sw1# show ip multicast summary pim-sm
```

```

Total number of PIM SM enabled VRFs      : 1
Total number of PIM SM interfaces         : 5
Total number of PIM SM neighbors          : 4
Total number of PIM A/A enabled VRF      : 0
Total number of PIM A/A enabled interfaces : 0
Total number of PIM SM Mroutes in route state : 0
Total number of PIM SM Mroutes in bridge state : 0
Total number of local multicast flows registered in this RP : 0
Total number of PIMv6 SM enabled VRFs    : 0
Total number of PIMv6 SM interfaces       : 0
Total number of PIMv6 SM neighbors        : 0

```

```
Total number of PIMv6 A/A enabled VRF : 0
Total number of PIMv6 A/A enabled interfaces : 0
Total number of PIMv6 SM Mroutes in route state : 0
Total number of PIMv6 SM Mroutes in bridge state : 0
----- Output Omitted -----
```

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.

```
# show ip mroute brief vrf b2
```

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.

```
# show ip mroute brief vrf b2
```

```
-----
VRF : b2                                     Total number of entries : 1
-----
Group Address      Source Address      Neighbor            Interface           Uptime
State
(HH:MM:SS)
-----
239.1.212.1        10.x.212.z          10.x.212.3          vlan212             00:00:06
bridge
-----
```

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.

```
Sw1# show ip multicast summary pim-sm
```

```
Total number of PIM SM enabled VRFs : 1
Total number of PIM SM interfaces : 5
Total number of PIM SM neighbors : 4
Total number of PIM A/A enabled VRF : 0
Total number of PIM A/A enabled interfaces : 0
Total number of PIM SM Mroutes in route state : 0
Total number of PIM SM Mroutes in bridge state : 0
Total number of local multicast flows registered in this RP : 1
Total number of PIMv6 SM enabled VRFs : 0
Total number of PIMv6 SM interfaces : 0
Total number of PIMv6 SM neighbors : 0
Total number of PIMv6 A/A enabled VRF : 0
Total number of PIMv6 A/A enabled interfaces : 0
Total number of PIMv6 SM Mroutes in route state : 0
Total number of PIMv6 SM Mroutes in bridge state : 0
```

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

```
Sw1# show ip pim interface loopback0 counters
```

```
Interface : loopback0
VRF       : default
```

Tx Counters :

```
----- Output Omitted -----
Register          0
Register Stop     4
```

Rx Counters :

```
----- Output Omitted -----
Register          5
Register Stop     0
----- Output Omitted -----
```

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.

```
Sw1# show ip mroute brief
```

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.

```
Agg1# show ip mroute brief
Agg1# show ip mroute brief vsx-peer
```

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.

```
Sw2# show ip igmp snooping vlan 11 detail
```

IGMP Snooping Protocol Info

Total VLANs with IGMP enabled : 2
Current count of multicast groups joined : 3

IGMP Drop Unknown Multicast : Global

VLAN ID : 11

VLAN Name : VLAN11

IGMP Configured Version : 3

IGMP Operating Version : 3

Querier Address : 10.x.11.2

Querier Port : lag255

Querier UpTime : 22m 26s

Querier Expiration Time : 2m 49s

Router Detected Port(s) : lag255

Active Group Address	Tracking	Vers	Mode	Uptime	Expires	Ports
239.1.212.1	Filter	3	EXC	1m 25s	2m 55s	1/1/27
239.255.255.250	Filter	3	EXC	24m 25s	2m 59s	1/1/27

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

```
Agg1# show ip igmp int vlan 11
```

IGMP Configured Version : 3

IGMP Operating Version : 3

Querier State : Querier

Querier IP [this switch] : 10.x.11.2

Querier Uptime : 29m 10s

Querier Expiration Time : 1m 18s

IGMP Snoop Enabled on VLAN : True

Active Group Address	Vers	Mode	Uptime	Expires
239.1.212.1	3	EXC	2m 52s	3m 36s
239.255.255.250	3	EXC	29m 5s	3m 42s

32. On Agg1, repeat the command for the VSX peer. Agg2 should also have detected the group.

```
Agg1# show ip igmp int vlan 11 vsx-peer
```

```

IGMP Configured Version : 3
IGMP Operating Version  : 3
Querier State           : Non-Querier
Querier IP              : 10.x.11.2
Querier Uptime          : 29m 50s
Querier Expiration Time : 2m 44s
IGMP Snoop Enabled on VLAN : True

```

Active Group Address	Vers	Mode	Uptime	Expires
239.1.212.1	3	EXC	3m 35s	2m 54s
239.255.255.250	3	EXC	29m 48s	2m 59s

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.

```

Agg1# show ip igmp snooping vlan 11 detail
IGMP Snooping Protocol Info

Total VLANs with IGMP enabled : 2
Current count of multicast groups joined : 3

IGMP Drop Unknown Multicast : Global
VLAN ID : 11
VLAN Name : VLAN11
IGMP Configured Version : 3
IGMP Operating Version : 3
Querier Address [this switch] : 10.x.11.2
Querier Port :
Querier UpTime : 30m 26s
Querier Expiration Time : 0m 2s
Router Detected Port(s) : lag256

```

Active Group Address	Tracking	Vers	Mode	Uptime	Expires	Ports
239.1.212.1	Filter	3	EXC	4m 8s	2m 20s	lag2
239.255.255.250	Filter	3	EXC	30m 20s	2m 26s	lag2

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# show ip igmp snooping vlan 11 detail vsx-peer

```

IGMP Snooping Protocol Info

```
Total VLANs with IGMP enabled      : 2
Current count of multicast groups joined : 3

IGMP Drop Unknown Multicast        : Global
VLAN ID : 11
VLAN Name : VLAN11
IGMP Configured Version : 3
IGMP Operating Version : 3
Querier Address : 10.x.11.2
Querier Port : lag256
Querier UpTime : 31m 0s
Querier Expiration Time : 3m 40s
Router Detected Port(s) : lag256
```

Active Group Address	Tracking	Vers	Mode	Uptime	Expires	Ports
239.1.212.1	Filter	3	EXC	4m 45s	3m 47s	lag2
239.255.255.250	Filter	3	EXC	30m 58s	3m 55s	lag2

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

```
Agg1# show ip pim interface vlan11
```

PIM Interfaces

VRF: default

```
Interface : vlan11
Neighbor count : 1
IP Address : 10.x.11.2/24
Mode : sparse
```

Designated Router : 10.x.11.3

```
Proxy DR : false
Hello Interval (sec) : 30
Hello Delay (sec) : 5
```

```
Override Interval (msec) : 2500      Lan Prune Delay : Yes
Propagation Delay (msec) : 500        DR Priority : 1
Neighbor Timeout : 81
```

```
Agg1# show ip pim interface vlan11 vsx-peer
```


PIM Interfaces

VRF: default

Interface : vlan11
Neighbor count : 1
IP Address : 10.x.11.3/24
Mode : sparse

Designated Router : 10.x.11.3

Proxy DR : false
Hello Interval (sec) : 30
Hello Delay (sec) : 5

Override Interval (msec) : 2500	Lan Prune Delay : Yes
Propagation Delay (msec) : 500	DR Priority : 1
Neighbor Timeout : 100	

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.

Agg1# show ip mroute brief

VRF : default			Total number of entries : 1	
Group Address	Source Address	Neighbor	Interface	Uptime
State				(HH:MM:SS)
239.1.212.1	10.x.212.z	10.x.11.3	vlan11	00:06:53
bridge				

Agg1# show ip mroute brief **vsx-peer**

VRF : default			Total number of entries : 1	
Group Address	Source Address	Neighbor	Interface	Uptime
State				(HH:MM:SS)
239.1.212.1	10.x.212.z	10.x.42.4	vlan42	00:07:31
route				

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.

```
Agg1# show ip mroute 239.1.212.1 10.x.212.z
IP Multicast Route Entries
VRF : default
Total number of entries : 1

Group Address          : 239.1.212.1
Source Address         : 10.x.212.z   # PC4 IP
Neighbor              : 10.x.11.3
Incoming interface     : vlan11
Multicast Routing Protocol : PIM-SM
Unicast Routing Protocol : OSPF
Metric                : 300
Metric Pref           : 110
Uptime (HH:MM:SS)     : 00:10:16
```

```
Agg1# show ip mroute 239.1.212.1 10.x.212.z vsx-peer
IP Multicast Route Entries
VRF : default
Total number of entries : 1

Group Address          : 239.1.212.1
Source Address         : 10.x.212.z   # PC4 IP
Neighbor              : 10.x.42.4
Incoming interface     : vlan42
Multicast Routing Protocol : PIM-SM
Unicast Routing Protocol : OSPF
Metric                : 300
Metric Pref           : 110
Uptime (HH:MM:SS)     : 00:10:58
Downstream Interface
Interface              State
-----
```

vlan11 forwarding

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.

```
Sw1# show ip mroute brief
```

VRF : default				Total number of entries : 1	
Group Address State	Source Address	Neighbor	Interface	Uptime (HH:MM:SS)	
239.1.212.1 route	10.x.212.z	10.x.241.2	vlan241	00:12:09	

Q: What is the incoming interface?

A: SVI 241 or SVI 242.

39. On Sw1, review the details of the route.

```
Sw1# show ip mroute 239.1.212.1 10.x.212.z
```

IP Multicast Route Entries

VRF : default
Total number of entries : 1

Group Address	: 239.1.212.1	
Source Address	: 10.x.212.z	# PC4 IP
Neighbor	: 10.x.241.2	
Incoming interface	: vlan241	
Multicast Routing Protocol	: PIM-SM	
Unicast Routing Protocol	: OSPF	
Metric	: 200	
Metric Pref	: 110	
Uptime (HH:MM:SS)	: 00:13:17	
Downstream Interface		
Interface	State	
-----	-----	
vlan42	forwarding	

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.

```
Agg1# show ip mroute brief vrf b2
```

VRF : b2				Total number of entries : 1
Group Address State	Source Address	Neighbor	Interface	Uptime (HH:MM:SS)
239.1.212.1 route	10.x.212.z	10.x.212.3	vlan212	00:23:01

```
Agg1# show ip mroute brief vrf b2 vsx-peer
```

VRF : b2				Total number of entries : 1
Group Address State	Source Address	Neighbor	Interface	Uptime (HH:MM:SS)
239.1.212.1 bridge	10.x.212.z	10.x.212.2	vlan212	00:23:03

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.

```
Agg1# show ip mroute 239.1.212.1 10.x.212.z vrf b2
```

IP Multicast Route Entries

VRF : b2

Total number of entries : 1

```
Group Address      : 239.1.212.1
Source Address     : 10.x.212.z   # PC4 IP
Neighbor           : 10.x.212.3
```

```
Incoming interface      : vlan212
Multicast Routing Protocol : PIM-SM
Unicast Routing Protocol : connected
Metric                  : 0
Metric Pref              : 0
Uptime (HH:MM:SS)       : 00:24:53
Downstream Interface
Interface                State
-----                -
vlan241                  forwarding
```

```
Agg1# show ip mroute 239.1.212.1 10.x.212.z vrf b2 vsx-peer
IP Multicast Route Entries

VRF : b2
Total number of entries : 1

Group Address           : 239.1.212.1
Source Address           : 10.x.212.z   # PC4 IP
Neighbor                : 10.x.212.2
Incoming interface       : vlan212
Multicast Routing Protocol : PIM-SM
Unicast Routing Protocol : connected
Metric                   : 0
Metric Pref               : 0
Uptime (HH:MM:SS)        : 00:25:16
```

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.

```
Agg2# mtrace 10.x.212.z 239.1.212.1
```

43. On Agg1, repeat this command.

```
Agg1# mtrace 10.x.212.z 239.1.212.1
Type escape sequence to abort.
Mtrace from Self for 10.x.212.z and Group 239.1.212.1
Querying full reverse path...
-0 Flow in bridged state
```

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)

```
Sw1# mtrace 10.x.212.z 239.1.212.1
Type escape sequence to abort.
Mtrace from Self for 10.x.212.z and Group 239.1.212.1
Querying full reverse path...
-0 10.x.241.4 PIM 1ms
-1 10.x.212.2 PIM 32ms
```

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.

```
Agg1# mtrace 10.x.212.z 239.1.212.1 vrf b2
Type escape sequence to abort.
Mtrace from Self for 10.x.212.z and Group 239.1.212.1
Querying full reverse path...
-0 Source directly connected
```

```
Agg2# mtrace 10.x.212.z 239.1.212.1 vrf b2
Type escape sequence to abort.
Mtrace from Self for 10.x.212.z and Group 239.1.212.1
Querying full reverse path...
-0 Flow 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

```
Agg2# mtrace 10.x.212.z 239.1.212.1 lhr 10.x.11.3 vrf b2
Type escape sequence to abort.
Mtrace from 10.x.11.3 for 10.x.212.z and Group 239.1.212.1
Querying full reverse path...
```

```
-0 10.x.42.3      PIM 1ms
-1 10.x.241.4     PIM 30ms
-2 10.x.212.2     PIM 42ms
```

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.

```
Agg2# mtrace 10.x.212.z 239.1.212.1 lhr 10.x.41.4 vrf b2
Type escape sequence to abort.
Mtrace from 10.x.41.4 for 10.x.212.z and Group 239.1.212.1
Querying full reverse path...
-0 10.x.241.4      PIM 1ms
-1 10.x.212.2      PIM 4ms
```

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.

```
Sw1# diagnostics
Sw1# diag-dump pim basic
```

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.

```
Sw1# diag-dump pim basic local-file
Diagnostic-dump for pim was captured into 'local-file'. Please use the 'copy diag-dump local-file' command to copy-out this file.
Sw1# copy diag-dump local-file tftp://10.251.1.91/pim-sw1.txt vrf mgmt
```

% Total	% Received	% Xferd	Average Speed	Time	Time	Time	Current
			Dload Upload	Total	Spent	Left	Speed
100 48858	0	0 100 48858	0 840k	--:--:--	--:--:--	--:--:--	840k
100 48858	0	0 100 48858	0 838k	--:--:--	--:--:--	--:--:--	838k
TX-Sw1#							

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.

Example output:

```
-----
piminfo vrf default rp
-----
```

```
rp::mgr
rpTable.count=2, rpEntryTable.count=2
rp ip= state={noRp}
rp ip=10.x.0.4 state={self}
rpEntry group=239.255.255.250 rp=10.x.0.4 time=168
rpEntry group=239.1.212.1 rp=10.x.0.4 time=168
```


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.

```
Agg1# conf t
Agg1(config)# router pim vrf b2
Agg1(config-pim)# rp-candidate source-ip-interface loopback2
Agg1(config-pim)# rp-candidate group-prefix 239.1.212.0/24
Agg1(config-pim)# exit
```

2. Repeat this for Agg2(b2)

```
Agg2# conf t
Agg2(config)# router pim vrf b2
Agg2(config-pim)# rp-candidate source-ip-interface loopback2
Agg2(config-pim)# rp-candidate group-prefix 239.1.212.0/24
Agg2(config-pim)# exit
```

3. Verify on Sw1 that the RP-SET was updated. Use repeat until the list is updated, this may take a few moments.

```
Sw1# show ip pim rp-set
Sw1# repeat

*****
Iteration : 21 Command : show ip pim rp-set
*****

VRF: default

Status and Counters - PIM-SM Learned RP-Set Information
Group Address      Group Mask          RP Address           Hold Time  Expire Time
-----
224.0.0.0          240.0.0.0          10.x.200.3          150        147
224.0.0.0          240.0.0.0          10.x.0.4             150        147
224.0.0.0          240.0.0.0          10.x.200.2          150        147
239.1.212.0        255.255.255.0      10.x.200.2          150        147
239.1.212.0        255.255.255.0      10.x.200.3          150        147
```

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

```
Sw1# show ip pim interface loopback0 counters
```

```
Interface : loopback0
VRF       : default
```

```
Tx Counters :
```

```
----- Output Omitted -----
```

```
Register           0
Register Stop      44
```

```
Rx Counters :
```

```
----- Output Omitted -----
```

```
Register           45
Register Stop      0
```

```
----- Output Omitted -----
```

5. On Agg1, review the current Rx RP Register counters of Loopback2, repeat this for the VSX peer

```
Agg1# show ip pim interface loopback2 counters
```

```
Interface : loopback2
VRF       : b2
```

```
Tx Counters :
```

```
----- Output Omitted -----
```

```
RP Advertisement    7
```

```
----- Output Omitted -----
```

```
Register           0
Register Stop      5
```

```
Rx Counters :
```

```
----- Output Omitted -----
```

```
Register           5
Register Stop      0
```

----- Output Omitted -----

Agg1# show ip pim interface loopback2 counters vsx-peer

Interface : loopback2
VRF : b2

Tx Counters :

----- Output Omitted -----

RP Advertisement 7

----- Output Omitted -----

Register 0

Register Stop 0

Rx Counters :

----- Output Omitted -----

Register 0

Register Stop 0

----- Output Omitted -----

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.

Sw1# show ip pim interface loopback0 counters

Interface : loopback0
VRF : default

Tx Counters :

----- Output Omitted -----

Register 0

Register Stop 44

Rx Counters :

----- Output Omitted -----

Register 45

Register Stop 0

----- Output Omitted -----

8. On Agg1, repeat the counters command and the VSX-peer counters command. The RX Register counter should have increased.

Agg1# show ip pim interface loopback2 counters | in Register

Register 0

Register Stop 12

Register 12

Register Stop 0

Register Drops(Register ACL hitcount)	0

Agg1# show ip pim interface loopback2 counters vsx | in Register

Register	0
Register Stop	0
Register	0
Register Stop	0
Register Drops(Register ACL hitcount)	0

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.

No.	Time	Source	Destination	Protocol	Length	Info
1	0.000000	10.12.212.205	239.1.1.1	UDP	1042	52388 → 5001 Len=1000
2	0.999548	10.12.212.205	239.1.1.1	UDP	1042	52388 → 5001 Len=1000
3	2.000414	10.12.212.205	239.1.1.1	UDP	1042	52388 → 5001 Len=1000
4	3.000028	10.12.212.205	239.1.1.1	UDP	1042	52388 → 5001 Len=1000

> Frame 1: 1042 bytes on wire (8336 bits), 1042 bytes captured (8336 bits) on interface 0
> Ethernet II, Src: Vmware_b1:95:30 (00:50:56:b1:95:30), Dst: IPv4mcast_01:01:01 (01:00:5e:01:01:01)
> Internet Protocol Version 4, Src: 10.12.212.205, Dst: 239.1.1.1
> User Datagram Protocol, Src Port: 52388, Dst Port: 5001
> Data (1000 bytes)

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.

No.	Time	Source	Destination	Protocol	Length	Info
2	0.998079	10.12.212.205	238.1.1.1	UDP	1042	52389 → 5001 Len=1000
3	1.998012	10.12.212.205	238.1.1.1	UDP	1042	52389 → 5001 Len=1000
4	2.042572	b8:d4:e7:41:68:00	Broadcast	ARP	60	Gratuitous ARP for 10.12.212.3 (Request)
5	2.998318	10.12.212.205	238.1.1.1	UDP	1042	52389 → 5001 Len=1000

```

> Frame 1: 1042 bytes on wire (8336 bits), 1042 bytes captured (8336 bits) on interface 0
> Ethernet II, Src: Vmware_b1:95:30 (00:50:56:b1:95:30), Dst: IPv4mcast_01:01:01 (01:00:5e:01:01:01)
> Internet Protocol Version 4, Src: 10.12.212.205, Dst: 238.1.1.1
> User Datagram Protocol, Src Port: 52389, Dst Port: 5001
> Data (1000 bytes)

```

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.

```
Agg1# conf t
Agg1(config)# router pim
Agg1(config-pim)# active-active
Agg1(config-pim)# router pim vrf b2
Agg1(config-pim)# active-active
Agg1(config-pim)# exit
```

2. On Agg2, repeat this configuration

```
Agg2# conf t
Agg2(config)# router pim
Agg2(config-pim)# active-active
Agg2(config-pim)# router pim vrf b2
Agg2(config-pim)# active-active
Agg2(config-pim)# exit
```

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.

```
Agg1# show interface vlan 11 | include MAC
Hardware: Ethernet, MAC Address: b8:d4:e7:41:38:00
```

4. On Agg1, repeat this for the VSX peer and note the MAC address of Agg2 SVI11.

```
Agg1# show interface vlan 11 vsx-peer | include MAC
Hardware: Ethernet, MAC Address: b8:d4:e7:41:68:00
```

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.

```
Agg1# show ip pim interface vlan11
PIM Interfaces
VRF: default

Interface : vlan11
Neighbor count : 1
IP Address : 10.x.11.2/24
Mode      : sparse

Designated Router : 10.x.11.2
Proxy DR          : false
Hello Interval (sec) : 30
Hello Delay (sec)   : 5

Override Interval (msec) : 2500      Lan Prune Delay      : Yes
Propagation Delay (msec) : 500       DR Priority           : 1
Neighbor Timeout      : 87
```

```
Agg1# show ip pim interface vlan11 vsx-peer
```

```
PIM Interfaces
```


VRF: default

Interface : vlan11
Neighbor count : 1
IP Address : 10.x.11.3/24
Mode : sparse

Designated Router : 10.x.11.3
Proxy DR : true
Hello Interval (sec) : 30
Hello Delay (sec) : 5

Override Interval (msec) : 2500 Lan Prune Delay : Yes
Propagation Delay (msec) : 500 DR Priority : 1
Neighbor Timeout : 88

6. On Agg1, you can also use the brief option

Agg1# show ip pim interface brief

VRF : default			Total number of interfaces : 4		
Interface	IP Address	DR Address	Neighbor count	Mode	VSX Role
vlan1	10.x.1.2/24	10.x.1.3	1	sparse	N/A
vlan41	10.x.41.2/24	10.x.41.4	1	sparse	N/A
loopback0	10.x.0.2/32	N/A	N/A	sparse	N/A
vlan11	10.x.11.2/24	10.x.11.2	1	sparse	DR

Agg1# show ip pim interface brief vsx-peer

VRF : default			Total number of interfaces : 4		
Interface	IP Address	DR Address	Neighbor count	Mode	VSX Role
vlan11	10.x.11.3/24	10.x.11.3	1	sparse	Proxy DR
vlan42	10.x.42.3/24	10.x.42.4	1	sparse	N/A
loopback0	10.x.0.3/32	N/A	N/A	sparse	N/A
vlan1	10.x.1.3/24	10.x.1.3	1	sparse	N/A

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.

No.	Time	Source	Destination	Protocol	Length	Info
1	0.000000	10.12.11.3	224.0.0.13	PIMv2	68	Hello
2	0.023910	10.12.212.205	239.1.212.1	UDP	1042	52390 → 5001 Len=1000
3	1.022867	10.12.212.205	239.1.212.1	UDP	1042	52390 → 5001 Len=1000
4	2.022849	10.12.212.205	239.1.212.1	UDP	1042	52390 → 5001 Len=1000
5	3.022534	10.12.212.205	239.1.212.1	UDP	1042	52390 → 5001 Len=1000
6	3.122112	Vmware_b1:04:25	02:00:00:00:00:01	ARP	42	Who has 10.12.11.1? Tell 10.12.11.103
7	3.123211	02:00:00:00:00:01	Vmware_b1:04:25	ARP	60	10.12.11.1 is at 02:00:00:00:00:01
8	3.123418	02:00:00:00:00:01	Vmware_b1:04:25	ARP	60	10.12.11.1 is at 02:00:00:00:00:01
9	4.022357	10.12.212.205	239.1.212.1	UDP	1042	52390 → 5001 Len=1000
10	5.024155	10.12.212.205	239.1.212.1	UDP	1042	52390 → 5001 Len=1000

> Frame 2: 1042 bytes on wire (8336 bits), 1042 bytes captured (8336 bits) on interface 0
 > Ethernet II, Src: b8:d4:e7:41:38:00 (b8:d4:e7:41:38:00), Dst: IPv4mcast_01:d4:01 (01:00:5e:01:d4:01)
 > Internet Protocol Version 4, Src: 10.12.212.205, Dst: 239.1.212.1
 > User Datagram Protocol, Src Port: 52390, Dst Port: 5001
 > Data (1000 bytes)

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.

```
Sw1# show ip mroute brief
```

VRF : default			Total number of entries : 1	
Group Address	Source Address	Neighbor	Interface	Uptime
State				(HH:MM:SS)
239.1.212.1	10.x.212.z	10.x.241.2	vlan241	00:03:59
route				

12. On Sw1, now review the details of this source group combination.

```
Sw1# show ip mroute 239.1.212.1 10.x.212.z
```

IP Multicast Route Entries

VRF : default

Total number of entries : 1

```

Group Address      : 239.1.212.1
Source Address     : 10.x.212.z
Neighbor          : 10.x.241.2
Incoming interface : vlan241
Multicast Routing Protocol : PIM-SM
Unicast Routing Protocol : OSPF
Metric            : 200
Metric Pref       : 110
Uptime (HH:MM:SS) : 00:05:12
Downstream Interface
Interface          State
-----
vlan41             forwarding
vlan42             forwarding
  
```

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.

```
Agg1# show ip mroute 239.1.212.1 10.x.212.z
```

IP Multicast Route Entries

VRF : default

Total number of entries : 1

```

Group Address      : 239.1.212.1
Source Address     : 10.x.212.z
Neighbor          : 10.x.41.4
Incoming interface : vlan41
Multicast Routing Protocol : PIM-SM
Unicast Routing Protocol : OSPF
Metric            : 300
Metric Pref       : 110
Uptime (HH:MM:SS) : 00:07:06
Downstream Interface
Interface          State          VSX Role
-----
vlan11             forwarding    DR
  
```

```
Agg1# show ip mroute 239.1.212.1 10.x.212.z vsx-peer
```

IP Multicast Route Entries

```
VRF : default
Total number of entries : 2

Group Address      : 239.1.212.1
Source Address     : 10.x.212.z
Neighbor           : 10.x.11.2
Incoming interface : vlan11
Multicast Routing Protocol : PIM-SM
Unicast Routing Protocol : OSPF
Metric             : 300
Metric Pref        : 110
Uptime (HH:MM:SS) : 00:07:22

Group Address      : 239.1.212.1
Source Address     : 10.x.212.z
Neighbor           : 10.x.42.4
Incoming interface : vlan42
Multicast Routing Protocol : PIM-SM
Unicast Routing Protocol : OSPF
Metric             : 300
Metric Pref        : 110
Uptime (HH:MM:SS) : 00:07:35

Interface      State      VSX Role
-----
vlan11         forwarding Proxy DR
```

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:

```
# show ip mroute 239.1.212.1 10.x.212.z
# repeat
```

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.

```
# write mem
```

```
# boot system
y
```

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.

No.	Time	Src	Source	Destination	Protocol	Length
173	102.999997	b8:d4:e7:41:38:00	10.12.212.205	239.1.212.1	UDP	
174	103.999681	b8:d4:e7:41:38:00	10.12.212.205	239.1.212.1	UDP	
175	104.339896	02:00:00:00:00:01	02:00:00:00:00:01	Broadcast	ARP	
176	105.015327	b8:d4:e7:41:38:00	10.12.212.205	239.1.212.1	UDP	
177	105.999722	b8:d4:e7:41:38:00	10.12.212.205	239.1.212.1	UDP	
178	106.999713	b8:d4:e7:41:38:00	10.12.212.205	239.1.212.1	UDP	
179	107.624159	Vmware_b1:04:25	Vmware_b1:04:25	02:00:00:00:00:01	ARP	
180	108.015651	b8:d4:e7:41:68:00	10.12.212.205	239.1.212.1	UDP	
181	108.624281	Vmware_b1:04:25	Vmware_b1:04:25	02:00:00:00:00:01	ARP	
182	108.627872	02:00:00:00:00:01	02:00:00:00:00:01	Vmware_b1:04:25	ARP	
183	108.999677	b8:d4:e7:41:68:00	10.12.212.205	239.1.212.1	UDP	
184	109.999895	b8:d4:e7:41:68:00	10.12.212.205	239.1.212.1	UDP	

```
> Frame 180: 1042 bytes on wire (8336 bits), 1042 bytes captured (8336 bits) on interface 0
> Ethernet II, Src: b8:d4:e7:41:68:00 (b8:d4:e7:41:68:00), Dst: IPv4mcast_01:d4:01 (01:00:5e:01:d4:01)
> Internet Protocol Version 4, Src: 10.12.212.205, Dst: 239.1.212.1
> User Datagram Protocol, Src Port: 52390, Dst Port: 5001
> Data (1000 bytes)
```

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.

```
*****
Iteration : 2 Command : show ip mroute 239.1.212.1 10.x.212.z
*****

----- Output Omitted -----

Interface      State      VSX Role
-----
vlan11         forwarding Proxy DR

*****
Iteration : 3 Command : show ip mroute 239.1.212.1 10.x.212.z
*****

----- Output Omitted -----

Interface      State      VSX Role
```

vlan11 forwarding DR

Q: Did the 'proxy DR' state change?

A: Yes, this system is now the active forwarder.

Aruba Training-Confidential

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.

```
Agg1# conf t
Agg1(config)# access-list ip igmpfilter
Agg1(config-acl-ip)# deny any any 239.1.2.0/255.255.255.0
Agg1(config-acl-ip)# permit any any any
Agg1(config-acl-ip)# vsx-sync
Agg1(config-acl-ip)# exit
```

2. Apply this ACL on the SVI 11 to IGMP.

```
Agg1(config)# interface vlan 11
Agg1(config-if-vlan)# ip igmp apply access-list igmpfilter
Agg1(config-if-vlan)# exit
```

3. Repeat this configuration on Agg2.

```
Agg2# conf t
Agg2(config)# interface vlan 11
Agg2(config-if-vlan)# ip igmp apply access-list igmpfilter
Agg2(config-if-vlan)# exit
```

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.

```
Agg1# show ip igmp interface vlan11 counters
```

IGMP Counters

```
Interface Name      : vlan11
VRF Name            : default
Membership Timeout   : 0
```

	Rx	Tx

V1 All Hosts Queries	0	0
V2 All Hosts Queries	0	0
V3 All Hosts Queries	0	12
V2 Group Specific Queries	0	0
V3 Group Specific Queries	0	0
Group And Source Specific Queries	0	0
V3 Member Reports	64	N/A
V2 Member Reports	0	N/A
V1 Member Reports	0	N/A
V2 Member Leaves	0	N/A
Packets dropped by ACL	2	N/A

```
Agg1# show ip igmp interface vlan11 counters vsx
```

IGMP Counters

```
Interface Name      : vlan11
VRF Name            : default
Membership Timeout   : 0
```

	Rx	Tx

V1 All Hosts Queries	0	0
V2 All Hosts Queries	0	0
V3 All Hosts Queries	62	2
V2 Group Specific Queries	0	0
V3 Group Specific Queries	6	0
Group And Source Specific Queries	0	0
V3 Member Reports	227	N/A
V2 Member Reports	0	N/A
V1 Member Reports	0	N/A
V2 Member Leaves	0	N/A
Packets dropped by ACL	2	N/A

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.

```
Sw2# conf t
Sw2(config)# access-list ip igmpsnoop
Sw2(config-acl-ip)# deny any any 239.1.3.0/255.255.255.0
Sw2(config-acl-ip)# permit any any any
Sw2(config-acl-ip)# exit
```

10. On Sw2, apply this 'igmpsnoop' ACL to the VLAN 11 IGMP snooping

```
Sw2(config)# vlan 11
Sw2(config-vlan-11)# ip igmp snooping apply access-list igmpsnoop
Sw2(config-vlan-11)# exit
```

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.

```
Sw2# show ip igmp snooping vlan 11 counters
IGMP Snooping VLAN Counters

VLAN ID    :    11
VLAN Name  : VLAN11

Rx Counters :

V1 All Hosts Queries           0
V2 All Hosts Queries           0
V3 All Hosts Queries          59
V2 Group Specific Queries       0
V3 Group Specific Queries       6
Group And Source Specific Queries 0
V1 Member Reports              0
V2 Member Reports              0
V3 Member Reports             291
V2 Member Leaves               0

Tx Counters :

Flood on vlan                  65
Forward to router              291
V2 Group Specific Queries       0
V3 Group Specific Queries       0

Errors:

Unknown Message Type           0
Malformed Packets              0
Bad Checksum                   0
Packet received on IGMP-disabled Interface 0
```

Interface Wrong Version Queries	0
Packets dropped by ACL	2
Port Counters:	
Membership Timeout	0

IGMP debug

14. On Sw2, enable debugging for 'igmp config'.

```
Sw2# debug igmp config
```

15. On Sw2, enable terminal-monitor for severity debug and filter on '239.1.3'

```
Sw2# terminal-monitor severity debug filter 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.

```
2020-10-09T21:05:36.233137-0400 hpe-mgmd[3297] <DEBUG>
debug|LOG_DEBUG|MSTR|1|IGMP|IGMP_CONFIG|Trying to find a match for src: '' groups:
'239.1.3.1' from ACL igmpsnoop
2020-10-09T21:05:36.233498-0400 hpe-mgmd[3297] <DEBUG>
debug|LOG_DEBUG|MSTR|1|IGMP|IGMP_CONFIG|ACL Rule match found for node with group: 239.1.3.1
and action: 0
```

This concludes the IGMP ACL task.

18. On Sw2, disable all debugging.

```
Sw2# no terminal-monitor
Sw2# no debug all
```

19. On Sw2, assign PC4 port to VLAN 12

```
Sw2# conf t
Sw2(config)# interface 1/1/4
Sw2(config-if)# vlan access 12
Sw2(config-lag-if)# exit
```

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.

Example output:

Server listening on 8000

Accepted connection from 10.x.12.104, port 2347

[5] local 10.6.11.103 port 8000 connected to 10.6.12.104 port 2348

[7] local 10.6.11.103 port 8000 connected to 10.6.12.104 port 2349

ID	Interval	Transfer	Bandwidth
[5]	0.00-5.01 sec	35.6 MBytes	59.7 Mbits/sec
[7]	0.00-5.01 sec	35.1 MBytes	58.8 Mbits/sec
[SUM]	0.00-5.01 sec	70.8 MBytes	118 Mbits/sec

...

4. **On both Agg1 and Agg2**, clear the int 1/1/2 statistics.

```
# clear interface 1/1/2 statistics
```

5. **On both Agg1 and Agg2**, show the interface 1/1/2 queue statistics.

```
# show interface 1/1/2 queues
```

Interface 1/1/2 is up

Admin state is up

	Tx Bytes	Tx Packets	Tx Drops
Q0	0	0	0
Q1	6828004673	4952923	0
Q2	0	0	0
Q3	0	0	0
Q4	0	0	0
Q5	0	0	0
Q6	21646	245	0
Q7	50753	435	0

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.

```
Agg1# show qos cos-map
```

code_point	local_priority	color	name
0	1	green	Best_Effort
1	0	green	Background
2	2	green	Excellent_Effort
3	3	green	Critical_Applications
4	4	green	Video
5	5	green	Voice
6	6	green	Internetwork_Control
7	7	green	Network_Control

Q: 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.

```
Agg1# show qos dscp-map
```

DSCP	code_point	local_priority	color	name
000000	0	1	green	CS0
000001	1	1	green	
000010	2	1	green	
000011	3	1	green	
000100	4	1	green	
000101	5	1	green	
000110	6	1	green	

```
000111 7 1 green
001000 8 0 green CS1
----- Output Omitted -----
```

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.

```
Agg1# show qos queue-profile
profile_status profile_name
-----
applied        factory-default
```

```
Agg1# show qos queue-profile factory-default
queue_num local_priorities name
-----
0          0                  Scavenger_and_backup_data
1          1
2          2
3          3
4          4
5          5
6          6
7          7
```

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.

```
Sw2# conf t
Sw2(config)# class ip bulk-data
Sw2(config-class-ip)# 10 match tcp any any eq 8000 count
Sw2(config-class-ip)# exit
```

10. On Sw2, define a policy 'asts-qos' that assigns the class 'bulk-data' the DSCP value CS1 and PCP value 2.

```
Sw2(config)# policy asts-qos
```



```
Sw2(config-policy)# 10 class ip bulk-data action dscp CS1 action pcp 2
Sw2(config-policy)# exit
```

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.

```
Sw2(config)# interface 1/1/4
Sw2(config-if)# apply policy asts-qos in
Sw2(config-if)# exit
```

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.

```
Sw2(config)# mirror session 1
Sw2(config-mirror-1)# destination cpu
Sw2(config-mirror-1)# enable
```

13. On Sw2, verify the current configuration of the mirror 1

```
Sw2(config-mirror-1)# show running current

mirror session 1
  destination cpu
  enable
```

14. On Sw2, define a policy 'qos-mirror' that matches class 'bulk-data' and applies action 'mirror 1'

```
Sw2(config)# policy qos-mirror
Sw2(config-policy)# class ip bulk-data action mirror 1
Sw2(config-policy)# exit
```

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').

```
Sw2(config)# int 1/1/27
Sw2(config-if)# apply policy qos-mirror out
Sw2(config-if)# exit
```

16. Review the mirror 1, repeat the command after about 10 seconds.

```
Sw2# show mirror 1

Mirror Session: 1
Admin Status: enable
Operation Status: enabled
Source: vlan rx none
Source: vlan tx none
Source: interface rx none
Source: interface tx none
Destination: cpu
Output Packets: 9616
Output Bytes: 14537960
```

...

```
Sw2# show mirror 1

Mirror Session: 1
Admin Status: enable
Operation Status: enabled
Source: vlan rx none
Source: vlan tx none
Source: interface rx none
Source: interface tx none
Destination: cpu
Output Packets: 17395
Output Bytes: 26299294
```

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.

```
Sw2# diagnostics
Sw2# diag utilities tshark
```

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.

```
----- Output Omitted -----
Ethernet II, Src: 90:20:c2:c0:cd:00, Dst: 00:50:56:b1:d1:a5
  Destination: 00:50:56:b1:d1:a5
    Address: 00:50:56:b1:d1:a5
```

```

.....0. .... = LG bit: Globally unique address (factory default)
.....0 .... = IG bit: Individual address (unicast)
Source: 90:20:c2:c0:cd:00
Address: 90:20:c2:c0:cd:00
.....0. .... = LG bit: Globally unique address (factory default)
.....0 .... = IG bit: Individual address (unicast)
Type: IPv4 (0x0800)
Internet Protocol Version 4, Src: 10.x.12.104, Dst: 10.x.11.103
0100 .... = Version: 4
.... 0101 = Header Length: 20 bytes (5)
Differentiated Services Field: 0x20 (DSCP: CS1, ECN: Not-ECT)
0010 00.. = Differentiated Services Codepoint: Class Selector 1 (8)
.... ..00 = Explicit Congestion Notification: Not ECN-Capable Transport (0)
Total Length: 1500
Identification: 0x5fcb (24523)
Flags: 0x4000, Don't fragment
 0... .... = Reserved bit: Not set
.1.. .... = Don't fragment: Set
..0. .... = More fragments: Not set
...0 0000 0000 0000 = Fragment offset: 0
Time to live: 127
Protocol: TCP (6)
Header checksum: 0x6a60 [validation disabled]
[Header checksum status: Unverified]
Source: 10.x.12.104
Destination: 10.x.11.103
Transmission Control Protocol, Src Port: 49749 (49749), Dst Port: 8000 (8000), Seq: 11222881,
Ack: 1, Len: 1460
----- Output Omitted -----

```

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
# show interface 1/1/2 queues
Interface 1/1/2 is up
Admin state is up

```

	Tx Bytes	Tx Packets	Tx Drops
Q0	0	0	0
Q1	855908	565	0
Q2	0	0	0

Q3	0	0	0
Q4	0	0	0
Q5	0	0	0
Q6	0	0	0
Q7	230	2	0

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.

```
# show qos trust
qos trust none
# show interface lag 2 | i qos
qos trust none
```

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.

```
Agg1# qos trust cos
```

```
Agg1# clear interface 1/1/2 statistics
Agg1# show interface 1/1/2 queues
Agg1# repeat count 5

*****
Iteration : 2 Command : show interface 1/1/2 queues
*****
Interface 1/1/2 is up
Admin state is up
      Tx Bytes      Tx Packets      Tx Drops
Q0              0              0              0
Q1          8196306          48828              0
Q2          617728866         406837              0
Q3              0              0              0
Q4              0              0              0
Q5              0              0              0
Q6           484              6              0
Q7          1163             10              0
```

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

*****
Iteration : 7 Command : show interface 1/1/2 queues
*****
Interface 1/1/2 is up
Admin state is up
```

	Tx Bytes	Tx Packets	Tx Drops
Q0	549067736	361620	0
Q1	7226388	42489	0
Q2	0	0	0
Q3	0	0	0
Q4	0	0	0
Q5	0	0	0
Q6	398	5	0
Q7	933	8	0

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.

```
Agg1# qos trust dscp
```

On Sw2, you will now review the local traffic handling.

24. On Sw2, clear the current session statistics of LAG255.

```
Sw2# clear interface lag255 statistics
```

25. Review the queue statistics of the LAG 255 and the member ports.

```
Sw2# show interface 1/1/25 queues

Interface 1/1/25 is up
Admin state is up
```

	Tx Bytes	Tx Packets	Tx Drops

Q0	0	0	0
Q1	86	1	0
Q2	0	0	0
Q3	0	0	0
Q4	0	0	0
Q5	0	0	0
Q6	0	0	0
Q7	0	0	0

Sw2# show interface 1/1/26 queues

Interface 1/1/26 is up

Admin state is up

	Tx Bytes	Tx Packets	Tx Drops
Q0	0	0	0
Q1	289298508	210425	0
Q2	0	0	0
Q3	0	0	0
Q4	0	0	0
Q5	0	0	0
Q6	0	0	0
Q7	1107	9	0

Sw2# show interface lag255 queues

Aggregate-name lag255

Aggregated-interfaces :

1/1/25 1/1/26

Speed 20000 Mb/s

	Tx Bytes	Tx Packets	Tx Drops
Q0	0	0	0
Q1	414149714	301473	0
Q2	0	0	0
Q3	0	0	0
Q4	0	0	0
Q5	0	0	0
Q6	0	0	0
Q7	1599	13	0

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.

Sw2# conf t

```
Sw2 (config)# policy asts-qos
Sw2 (config-policy)# 10 class ip bulk-data action dscp CS1 action local-priority 0
Sw2 (config-policy)# exit
```

27. Verify the policy configuration

```
Sw2# show policy asts-qos
```

```

      Name
      Additional Policy Parameters
Sequence Comment
      Class Type
              action
-----
      asts-qos
10
      bulk-data ipv4
              local-priority 0
              dscp CS1
-----
```

28. Verify that the traffic is now handled by the correct queue based on your local priority configuration.

```
Sw2# clear interface lag255 statistics
Sw2# show interface lag255 queues
```

```
Aggregate-name lag255
Aggregated-interfaces :
1/1/25 1/1/26
Speed 20000 Mb/s
```

	Tx Bytes	Tx Packets	Tx Drops
Q0	169683540	111690	0
Q1	2276423	13738	0
Q2	0	0	0
Q3	0	0	0
Q4	0	0	0
Q5	0	0	0
Q6	0	0	0
Q7	738	6	0

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 scenario's.

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.

```
Agg1# conf t
Agg1(config)# class ip bulk-data
Agg1(config-class-ip)# 10 match tcp any any eq 8000 count
Agg1(config-class-ip)# exit
```

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.

```
Agg1(config)# policy asts-qos
Agg1(config-policy)# 10 class ip bulk-data action cir kbps 20000 cbs 20000 exceed drop
Agg1(config-policy)# exit
```

4. Apply the policy to the LAG2 in the inbound direction.

```
Agg1(config)# interface lag 2
Agg1(config-lag-if)# apply policy asts-qos in
Agg1(config-lag-if)# exit
```

Troubleshoot the VSX-sync

Review the VSX sync status

5. Review the VSX-sync.

```
Agg1# show vsx status config-sync
Admin state           : Enabled
Operational State     : Operational
Error State           : Missing reference error
```



```
Recommended remediation : A. Execute 'show running-config vsx-sync peer-diff' to determine
                           which lines did not sync correctly on secondary
                           B. Identify the configuration that is missing from secondary and
                           manually fix via CLI.
Current time               : Mon Oct 12 16:46:50 2020
Last sync time             : Mon Oct 12 16:46:15 2020
```

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.

```
Agg1# show run vsx-sync peer-diff
----- Output Omitted -----
- apply policy asts-qos in
  interface lag 5 multi-chassis
  description aruba-gw1
  no routing
```

Q: What line is different?

A: The policy line is applied on Agg1, not on Agg2

7. On Agg2, review the defined policies.

```
Agg2# show policy
No Policy found.
```

Q: What do you observe?

A: There is no policy.

8. On Agg1, update the policy with the vsx-sync command.

```
Agg1(config)# policy asts-qos
Agg1(config-policy)# vsx-sync
Agg1(config-policy)# exit
```

9. On Agg2, review the VSX sync and the current configuration of the interface LAG255.

```
Agg2# show vsx status config-sync
Admin state           : Enabled
Operational State     : Operational
```

```
Error State           : Missing reference error
Recommended remediation : A. Execute 'show running-config vsx-sync peer-diff' to determine
which lines did not sync correctly on secondary
                        B. Identify the configuration that is missing from secondary and
manually fix via CLI.
Current time          : Mon Oct 12 16:49:32 2020
Last sync time        : Mon Oct 12 16:49:04 2020

Agg2# show run vsx-sync peer-diff
----- Output Omitted -----
+ 10 class ip bulk-data action cir kbps 20000 cbs 20000 exceed drop
  vlan 11
    vsx-sync
    ip igmp snooping enable

Agg2# show run int lag2
interface lag 2 multi-chassis
  no shutdown
  description access2
  no routing
  vlan trunk native 1
  vlan trunk allowed 1,11-13,211-212
  lacp mode active
! policy asts-qos user configuration currently being processed
! run 'show policy [commands]' to display active policy configuration.
  apply policy asts-qos in
  exit
```

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.

```
Agg1(config)# class ip bulk-data
Agg1(config-class-ip)# vsx-sync
Agg1(config-class-ip)# exit
```

11. On Agg2, review the VSX sync

```
Agg2# show vsx status config-sync
Admin state           : Enabled
Operational State     : Operational
Error State           : None
Recommended remediation : N/A
Current time          : Mon Oct 12 16:53:12 2020
Last sync time        : Mon Oct 12 16:52:44 2020
```

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.

```
Agg2# show run int lag2
interface lag 2 multi-chassis
  no shutdown
  description access2
  no routing
  vlan trunk native 1
  vlan trunk allowed 1,11-13,211-212
  lacp mode active
! policy asts-qos user configuration currently being processed
! run 'show policy [commands]' to display active policy configuration.
  apply policy asts-qos in
  exit

Agg2# show run vsx-sync peer-diff
----- Output Omitted -----
-! policy asts-qos user configuration does not match active configuration.
-! run 'policy NAME reset' to reset policy to match active configuration.
policy asts-qos
  vsx-sync
  !
```

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.

```
Agg2(config)# policy asts-qos reset
```

14. On Agg2, verify that the messages have disappeared from the running configuration.

```
Agg2# show run int lag2
interface lag 2 multi-chassis
  no shutdown
  description access2
  no routing
  vlan trunk native 1
  vlan trunk allowed 1,11-13,211-212
  lacp mode active
  apply policy asts-qos in
```

```

exit
Agg2(config)# show run vsx-sync peer-diff
qos trust dscp
ip helper-address 10.252.1.21
ip helper-address 10.252.1.21
qos trust dscp
ip helper-address 10.252.1.21
ip helper-address 10.252.1.21
No difference in configs.

```

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).

```
[SUM] 2248.0-2250.0 sec  1.75 MBytes  7.34 Mbits/sec
```

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.

```

Sw2(config)# interface lag 255
Sw2(config-lag-if)# qos shape 20000
Sw2(config-lag-if)# exit

```

17. On PC4, verify the impact. The iperf should now run close to 20Mbps again.

```
[SUM] 2374.0-2376.0 sec  4.13 MBytes  17.3 Mbits/sec
```

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

```
Agg1(config)# policy asts-qos  
Agg1(config-policy)# 10 class ip bulk-data action cir kbps 10000 cbs 20000 exceed drop
```

19. On PC4, review the impact.

```
[SUM] 2504.0-2506.0 sec   512 KBytes  2.10 Mbits/sec
```

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.

```
Agg1(config-policy)# 10 class ip bulk-data action cir kbps 10000 cbs 2000000 exceed drop  
Agg1(config-policy)# exit
```

21. On PC4, verify the impact. Traffic should now be around 10Mbps again.

```
[SUM] 2644.0-2646.0 sec  2.38 MBytes  9.96 Mbits/sec
```

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.

```
Sw2# ping 10.252.1.23
```

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.

```
Sw2(config)# radius-server host 10.252.1.23 key plaintext aruba123
Sw2(config)# aaa group server radius cppm
Sw2(config-sg)# server 10.252.1.23
Sw2(config-sg)# exit
```

2. On Sw2, configure 802.1X to authenticate using the server-group 'cppm'

```
Sw2(config)# aaa authentication port-access dot1x authenticator radius server-group cppm
```

3. On Sw2, enable 802.1X authentication globally

```
Sw2(config)# aaa authentication port-access dot1x authenticator enable
```

4. On Sw2, on the port 1/1/4, enable 802.1X and adjust the 802.1X timers.

```
Sw2(config)# interface 1/1/4
Sw2(config-if)# aaa authentication port-access dot1x authenticator
Sw2(config-if-dot1x-auth)# enable
Sw2(config-if-dot1x-auth)# eapol-timeout 30
Sw2(config-if-dot1x-auth)# max-eapol-requests 1
Sw2(config-if-dot1x-auth)# max-retries 1
Sw2(config-if-dot1x-auth)# exit
```

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.


```
Sw2(config-if)# vlan 19
Sw2(config-vlan-19)# description aaa-default-pvid
Sw2(config-vlan-19)# interface 1/1/4
Sw2(config-if)# vlan access 19
Sw2(config-if)# exit
```

- 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.

- 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

Sw2# show port-access clients

Port Access Clients

Port	MAC Address	Onboarded Method	Status	Role

```
-----
1/1/4    00:50:56:b1:95:30 dot1x          Success    RADIUS_773420618
```

```
Sw2# show port-access clients detail
```

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.

```
Sw2# clear dot1x authenticator statistics
```

10. Verify that the port statistics and the client-status have been cleared.

```
Sw2# show aaa authentication port-access dot1x authenticator interface 1/1/4 port-statistics
```

```
Port 1/1/4
```

```
=====
```

Client Details

```
-----
```

```
Number of clients           : 1
Number of authenticated clients : 1
Number of unauthenticated clients : 0
Number of authenticating clients : 0
```

Statistics

```
-----
```

```
EAPOL Frames Received           : 0
EAPOL Frames Transmitted         : 0
EAPOL Start Frames Received      : 0
EAPOL Logoff Frames Received     : 0
EAPOL Response ID Frames Received : 0
EAPOL Response Frames Received   : 0
EAPOL Request ID Frames Transmitted : 0
EAPOL Request Frames Transmitted : 0
EAPOL Invalid Frames Received    : 0
EAPOL EAP Length Error Frames Received : 0
EAPOL Last Received Frame Version : 0
EAPOL Last Received Frame Client MAC : 0
```

```
Sw2# show aaa authentication port-access dot1x authenticator interface 1/1/4 client-status
```

```
Client 00:50:56:b1:95:30, asts-accept, 1/1/4
```

```
=====
```

Authentication Details

```
-----
Status                : Authenticated
Type                  : Pass-Through
EAP-Method            : PEAP
Auth Failure reason   :
Time Since Last State Change : 366s
```

Authentication Statistics

```
-----
Authentication                : 0
Authentication Timeout        : 0
EAP-Start While Authenticating : 0
EAP-Logoff While Authenticating : 0
Successful Authentication      : 0
Failed Authentication          : 0
Re-Authentication             : 0
Successful Re-Authentication   : 0
Failed Re-Authentication       : 0
EAP-Start When Authenticated   : 0
EAP-Logoff When Authenticated  : 0
Re-Auths When Authenticated    : 0
Cached Re-Authentication       : 0
```

11. Clear the RADIUS server statistics and review that they have been cleared.

```
Sw2# clear radius-server statistics authentication
Sw2# show radius-server statistics authentication
```

```
Server Name      : 10.252.1.23
Auth-Port        : 1812
Accounting-Port  : 1813
VRF              : default
TLS Enabled      : No
```

Authentication Statistics

```
-----
Round Trip Time      : 0
Pending Requests     : 0
Timeouts             : 0
Bad Authenticators    : 0
Packets Dropped      : 0
Access Requests      : 0
Access challenge      : 0
Access Accepts       : 0
Access Rejects       : 0
Access Response Malformed : 0
Access Retransmits    : 0
Tracking Requests     : 0
Tracking Responses    : 0
Unknown Response Code : 0
```

12. Enable debugging for the radius and 802.1X portaccess modules.


```
Sw2# debug portaccess dot1x radius
Sw2# debug portaccess dot1x protocol
```

13. Clear the debug buffer

```
Sw2# clear debug buffer
```

14. Manually re-authenticate the client

```
Sw2# port-access reauthenticate interface 1/1/4
```

15. Review the port statistics after the clients was re-authenticated.

```
Sw2# show aaa authentication port-access dot1x authenticator interface 1/1/4 port-statistics
Port 1/1/4
=====

Client Details
-----
Number of clients           : 1
Number of authenticated clients : 1
Number of unauthenticated clients : 0
Number of authenticating clients : 0

Statistics
-----
EAPOL Frames Received      : 9
EAPOL Frames Transmitted   : 10
EAPOL Start Frames Received : 0
EAPOL Logoff Frames Received : 0
EAPOL Response ID Frames Received : 1
EAPOL Response Frames Received : 8
EAPOL Request ID Frames Transmitted : 1
EAPOL Request Frames Transmitted : 8
EAPOL Invalid Frames Received : 0
EAPOL EAP Length Error Frames Received : 0
EAPOL Last Received Frame Version : 0
EAPOL Last Received Frame Client MAC : 0
```

```
Sw2# show aaa authentication port-access dot1x authenticator interface 1/1/4 client-status
Client 00:50:56:b1:95:30, asts-accept, 1/1/4
=====

Authentication Details
-----
Status           : Authenticated
Type             : Pass-Through
EAP-Method       : PEAP
Auth Failure reason :
Time Since Last State Change : 51s
```

Authentication Statistics

```
-----
Authentication                : 0
Authentication Timeout        : 0
EAP-Start While Authenticating : 0
EAP-Logoff While Authenticating : 0
Successful Authentication     : 0
Failed Authentication         : 0
Re-Authentication             : 1
Successful Re-Authentication  : 1
Failed Re-Authentication      : 0
EAP-Start When Authenticated  : 0
EAP-Logoff When Authenticated : 0
Re-Auths When Authenticated   : 0
Cached Re-Authentication      : 0
```

16. Review the RADIUS statistics

Sw2# **show radius-server statistics authentication**

```
Server Name      : 10.252.1.23
Auth-Port       : 1812
Accounting-Port  : 1813
VRF             : default
TLS Enabled     : No
```

Authentication Statistics

```
-----
Round Trip Time      : 0
Pending Requests    : 0
Timeouts            : 0
Bad Authenticators   : 0
Packets Dropped     : 0
Access Requests     : 9
Access challenge     : 8
Access Accepts      : 1
Access Rejects      : 0
Access Response Malformed : 0
Access Retransmits  : 0
Tracking Requests    : 0
Tracking Responses   : 0
Unknown Response Code : 0
```

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

Sw2# **show debug buffer**

```
2020-10-15:14:43:12.219820|port-
accessd|LOG_DEBUG|MSTR|1|PORTACCESS|PORTACCESS_DOT1X_PROTOCOL|logID=4988 logID=4988 Event
handler of Dot1XPort '1/1/4' for event 'Re-Authenticate Client On Po
rt' in state 'READY' returned 'OK'
2020-10-15:14:43:13.080037|port-
accessd|LOG_DEBUG|MSTR|1|PORTACCESS|PORTACCESS_DOT1X_PROTOCOL|logID=4989 dot1xpae SM State
transition [AUTHENTICATED] -> [REAUTHENTICATING] for user 'asts-acc
ept' with key '1/1/4, 00:50:56:b1:95:30'
2020-10-15:14:43:13.081454|port-
accessd|LOG_DEBUG|MSTR|1|PORTACCESS|PORTACCESS_DOT1X_PROTOCOL|logID=4989 logID=4989 DB
Operation: Update state for pae with MAC 00:50:56:b1:95:30 on port '1/1
/4': state update for pae with MAC 00:50:56:b1:95:30 on port '1/1/4' successful
2020-10-15:14:43:13.096469|port-
accessd|LOG_DEBUG|MSTR|1|PORTACCESS|PORTACCESS_DOT1X_PROTOCOL|logID=4991 logID=4991 Handling
event 'Dot1X EAP Packet Received on Port' for Dot1XPort '1/1/4' i
n state 'READY'
----- Output Omitted -----
2020-10-15:14:43:13.248066|port-
accessd|LOG_DEBUG|MSTR|1|PORTACCESS|PORTACCESS_DOT1X_PROTOCOL|logID=5016 dot1xpae SM State
transition [REAUTHENTICATING] -> [AUTHENTICATED] for user 'asts-acc
ept' with key '1/1/4, 00:50:56:b1:95:30'
----- Output Omitted -----
```

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

```
Sw2# no debug all  
Sw2# debug portaccess all
```

2. On Sw2, clear the debug buffer

```
Sw2# clear 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.

Screenshot of the ClearPass RADIUS output:

Request Details		
Summary	Input	Output
Enforcement Profiles:	asts-aruba-role-employee	
System Posture Status:	UNKNOWN (100)	
Audit Posture Status:	UNKNOWN (100)	
RADIUS Response		
Radius:Aruba:Aruba-User-Role	employee	

Showing 1 of 1-4 records | Change Status | Show Configuration | Export | Show Logs | Close

4. On Sw2, review the client details

```
Sw2# show port-access client detail
```

Port Access Client Status Details:

Client 00:50:56:b1:95:30, **asts-employee-vsa**

=====

Session Details

Port : 1/1/4
Session Time : 2234s
IPv4 Address :
IPv6 Address :

Authentication Details

Status : dot1x Authenticated
Auth Precedence : dot1x - Authenticated, mac-auth - Not attempted

Authorization Details

Role :
Status : Invalid

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'.

```
Sw2#show debug buffer severity info
Sw2#show debug buffer
```

```
2020-10-15:14:53:46.673681|port-
accessd|LOG_INFO|MSTR|1|PORTACCESS|PORTACCESS_SERVICES|logID=5701 Secure client state changed
for client 00:50:56:b1:95:30 in port 1/1/4 to Authenticated but
failed to apply role.
2020-10-15:14:53:46.673720|port-
accessd|LOG_INFO|MSTR|1|PORTACCESS|PORTACCESS_SERVICES|logID=5701 Port is in access blocked
state
2020-10-15:14:53:46.674160|port-
accessd|LOG_INFO|MSTR|1|PORTACCESS|PORTACCESS_SERVICES|logID=5701 PortClientProfile current
profile type is invalid 0
```

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.

```
Sw2# diagnostics
Sw2# diag utilities tcpdump host-port 1812
```

7. In your first connection to Sw2, trigger re-authentication on port 1/1/4.

```
Sw2# port-access reauthenticate interface 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.

```
12 15:00:03.836122 ethertype IPv4, IP (tos 0x0, ttl 60, id 50436, offset 0, flags [DF],
proto UDP (17), length 281)
10.252.1.23.radius > TX-Sw2.51056: RADIUS, length: 253
Access-Accept (2), id: 0x27, Authenticator: 9f9bddcf3cdef2dc77b2f737571d75e9
Vendor-Specific Attribute (26), length: 16, Value: Vendor: Unknown (14823)
Vendor Attribute: 1, Length: 8, Value: employee
Vendor-Specific Attribute (26), length: 58, Value: Vendor: Microsoft (311)
Vendor Attribute: 17, Length: 50, Value:
.....K.`..l.n.].u...5)._..q....f..i...2.."P..
Vendor-Specific Attribute (26), length: 58, Value: Vendor: Microsoft (311)
Vendor Attribute: 16, Length: 50, Value:
...D.....9N..D...[.....D0.Mn.....|....7..0.."S.
EAP-Message Attribute (79), length: 6, Value: .R
Message-Authenticator Attribute (80), length: 18, Value: .....fH2...}M.o
User-Name Attribute (1), length: 19, Value: asts-employee-vsa
Class Attribute (25), length: 58, Value: ?03.|.B..n..PD5w..
```

9. Close the second session (the TCPDUMP session).
10. On Sw2, configure the local user role 'employee' with access VLAN 11.

```
Sw2# conf t
Sw2(config)# port-access role employee
Sw2(config-pa-role)# vlan access 11
Sw2(config-pa-role)# exit
```

11. On Sw2, trigger re-authentication and verify that the client is now online.

```
Sw2(config)# do port-access reauthenticate interface 1/1/4
Sw2(config)# show port-access clients
```

Port Access Clients

Port	MAC Address	Onboarded Method	Status	Role
1/1/4	00:50:56:b1:95:30	dot1x	Success	employee

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.

- On Sw2, modify the local user role 'employee' with a re-authentication timer of 30 seconds.

```
Sw2(config)# port-access role employee
Sw2(config-pa-role)# reauth-period 30
Sw2(config-pa-role)# exit
```

NOTE: This short re-authentication timer is for this lab purposes only and should not be applied in production environments.

- On Sw2, enable debug and terminal monitor.

```
Sw2(config)# terminal-monitor severity debug filter dot1xpae
Terminal-monitor is enabled successfully
```

- On Sw2, trigger a re-authentication using the command

```
Sw2(config)# do port-access reauthenticate interface 1/1/4
```

- On Sw2, the terminal monitor should show the messages

```
AUTHENTICATED > REAUTHENTICATING
REAUTHENTICATING > AUTHENTICATED
```

```
2020-10-15T15:09:29.995094-0400 port-accesssd[3361] <DEBUG>
debug|LOG_DEBUG|MSTR|1|PORTACCESS|PORTACCESS_DOT1X_PROTOCOL|logID=6854 dot1xpae SM State
transition [AUTHENTICATED] -> [REAUTHENTICATING] for user 'asts-employee-vsa' with key
'1/1/4, 00:50:56:b1:95:30'
2020-10-15T15:09:30.132812-0400 port-accesssd[3361] <DEBUG>
debug|LOG_DEBUG|MSTR|1|PORTACCESS|PORTACCESS_DOT1X_PROTOCOL|logID=6881 dot1xpae SM State
transition [REAUTHENTICATING] -> [AUTHENTICATED] for user 'asts-employee-vsa' with key
'1/1/4, 00:50:56:b1:95:30'
```

16. On Sw2, after about 30 seconds, new messages should confirm that the re-authentication occurs every 30 seconds.

```
2020-10-15T15:10:00.006141-0400 port-accesssd[3361] <DEBUG>
debug|LOG_DEBUG|MSTR|1|PORTACCESS|PORTACCESS_DOT1X_PROTOCOL|logID=6913 dot1xpae SM State
transition [AUTHENTICATED] -> [REAUTHENTICATING] for user 'asts-employee-vsa' with key
'1/1/4, 00:50:56:b1:95:30'
2020-10-15T15:10:00.154680-0400 port-accesssd[3361] <DEBUG>
debug|LOG_DEBUG|MSTR|1|PORTACCESS|PORTACCESS_DOT1X_PROTOCOL|logID=6940 dot1xpae SM State
transition [REAUTHENTICATING] -> [AUTHENTICATED] for user 'asts-employee-vsa' with key
'1/1/4, 00:50:56:b1:95:30'
```

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

```
Agg1# conf t
Agg1(config)# access-list ip radius
Agg1(config-acl-ip)# deny ip any 10.252.1.23/32 count
Agg1(config-acl-ip)# permit ip any any
Agg1(config-acl-ip)# int 1/1/7
Agg1(config-if)# apply access-list ip radius out
Agg1(config-if)# exit
```

18. On Sw2, attempt to ping 10.252.1.23, this should fail now.

```
Sw2(config)# do ping 10.252.1.23 repetitions 2
PING 10.252.1.23 (10.252.1.23) 100(128) bytes of data.

--- 10.252.1.23 ping statistics ---
2 packets transmitted, 0 received, 100% packet loss, time 1026ms
```

19. On Sw2, wait up to 30 seconds and review terminal monitor output.

```
2020-10-15T15:18:25.282516-0400 port-accesssd[3361] <DEBUG>
debug|LOG_DEBUG|MSTR|1|PORTACCESS|PORTACCESS_DOT1X_PROTOCOL|logID=7679 dot1xpae SM State
transition [AUTHENTICATED] -> [REAUTHENTICATING] for user 'asts-employee-vsa' with key
'1/1/4, 00:50:56:b1:95:30'
2020-10-15T15:18:35.293983-0400 port-accesssd[3361] <DEBUG>
debug|LOG_DEBUG|MSTR|1|PORTACCESS|PORTACCESS_DOT1X_PROTOCOL|logID=7692 dot1xpae SM State
```



```
transition [REAUTHENTICATING] -> [UNAUTHENTICATED] for user 'asts-employee-vsa' with key
'1/1/4, 00:50:56:b1:95:30'
2020-10-15T15:18:43.295533-0400 port-accesssd[3361] <DEBUG>
debug|LOG_DEBUG|MSTR|1|PORTACCESS|PORTACCESS_DOT1X_PROTOCOL|logID=7705 dot1xpae SM State
transition [UNAUTHENTICATED] -> [AUTHENTICATING] for user 'asts-employee-vsa' with key
'1/1/4, 00:50:56:b1:95:30'
2020-10-15T15:18:53.302504-0400 port-accesssd[3361] <DEBUG>
debug|LOG_DEBUG|MSTR|1|PORTACCESS|PORTACCESS_DOT1X_PROTOCOL|logID=7717 dot1xpae SM State
transition [AUTHENTICATING] -> [UNAUTHENTICATED] for user 'asts-employee-vsa' with key
'1/1/4, 00:50:56:b1:95:30'
```

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.

```
Sw2(config)# no terminal
Sw2(config)# show debug buffer reverse severity info

2020-10-15:15:19:29.317899|port-
accesssd|LOG_ERR|MSTR|1|PORTACCESS|PORTACCESS_DOT1X_RADIUS|RADIUS request for context
0xfffffad0125e0 client request id 165 radius id 164 timedout error Server timedout. Inform to
client.
2020-10-15:15:19:24.315163|port-
accesssd|LOG_WARN|MSTR|1|PORTACCESS|PORTACCESS_DOT1X_RADIUS|Retrying RADIUS request for
context 0xfffffad0125e0 client request id 165 radius id 164 .
2020-10-15:15:19:19.332939|port-
accesssd|LOG_INFO|MSTR|1|PORTACCESS|PORTACCESS_DOT1X_RADIUS|Received authentication request
Access-Request(New) from 802.1x to send to group cppm
2020-10-15:15:19:11.307368|port-
accesssd|LOG_ERR|MSTR|1|PORTACCESS|PORTACCESS_DOT1X_RADIUS|RADIUS request for context
0xfffffad012520 client request id 164 radius id 163 timedout error Server timedout. Inform to
client.
```

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.

```
Sw2(config)# show radius-server statistics authentication

Server Name       : 10.252.1.23
Auth-Port         : 1812
Accounting-Port   : 1813
VRF               : default
TLS Enabled       : No

Authentication Statistics
```

```

-----
Round Trip Time      : 0
Pending Requests    : 0
Timeouts            : 17
Bad Authenticators   : 0
Packets Dropped     : 0
Access Requests     : 174
Access challenge     : 139
Access Accepts      : 17
Access Rejects      : 1
Access Response Malformed : 0
Access Retransmits  : 9
Tracking Requests   : 0
Tracking Responses  : 0
Unknown Response Code : 0

```

22. Review the port access clients. The client should now have failed authentication.

```
Sw2(config)# show port-access clients
```

Port Access Clients

```

-----
Port      MAC Address      Onboarded      Status      Role
          |               |      Method    |
-----
1/1/4     00:50:56:b1:95:30      |              | Fail

```

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'

```
Sw2(config)# port-access role asts-critical
Sw2(config-pa-role)# vlan access 11
```

24. On Sw2, configure the port 1/1/4 with the critical role 'asts-critical'

```
Sw2(config-pa-role)# interface 1/1/4
Sw2(config-if)# aaa authentication port-access critical-role asts-critical
Sw2(config-if)# exit
```

25. On Sw2 bounce port 1/1/4 and review the authentication status.

```
Sw2(config)# show port-access clients
```

Port Access Clients

```

-----
Port      MAC Address      Onboarded      Status      Role
          |               |      Method    |
-----
1/1/4     00:50:56:b1:95:30      |              | Success    | asts-critical, Critical

```

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.

```
Agg1(config-if)# interface 1/1/7
Agg1(config-if)# no apply access-list ip radius out
```

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.

```
Sw2(config)# interface 1/1/4
Sw2(config-if)# aaa authentication port-access dot1x authenticator
Sw2(config-if-dot1x-auth)# cached-reauth
```

28. On Sw2, bounce port 1/1/4


```
Sw2(config-if-dot1x-auth)# interface 1/1/4
Sw2(config-if)# shutdown
Sw2(config-if)# no shutdown
```

29. On Sw2, verify the PC4 client is authenticated again

```
Sw2(config-if)# show port-access clients
```

Port Access Clients				
Port	MAC Address	Onboarded Method	Status	Role
1/1/4	00:50:56:b1:95:30	dot1x	Success	employee

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.

```
Sw2(config-if)# terminal-monitor severity debug filter dot1xpae
```

31. On Agg1, block the RADIUS communication again.

```
Agg1(config-if)# interface 1/1/7
Agg1(config-if)# apply access-list ip radius out
```

32. On Sw2, review the log output.

```
2020-10-15T15:34:16.837220-0400 port-accesssd[3361] <DEBUG>
debug|LOG_DEBUG|MSTR|1|PORTACCESS|PORTACCESS_DOT1X_PROTOCOL|logID=8921 dot1xpae SM State
transition [AUTHENTICATED] -> [REAUTHENTICATING] for user 'asts-employee-vsa' with key
'1/1/4, 00:50:56:b1:95:30'
2020-10-15T15:34:26.844728-0400 port-accesssd[3361] <DEBUG>
debug|LOG_DEBUG|MSTR|1|PORTACCESS|PORTACCESS_DOT1X_PROTOCOL|logID=8933 dot1xpae SM State
transition [REAUTHENTICATING] -> [CACHED-REAUTHENTICATED] for user 'asts-employee-vsa' with
key '1/1/4, 00:50:56:b1:95:30'
```

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.

```
2020-10-15T15:34:45.850177-0400 port-accesssd[3361] <DEBUG>
debug|LOG_DEBUG|MSTR|1|PORTACCESS|PORTACCESS_DOT1X_PROTOCOL|logID=8954 dot1xpae SM State
transition [CACHED-REAUTHENTICATED] -> [REAUTHENTICATING] for user 'asts-employee-vsa' with
key '1/1/4, 00:50:56:b1:95:30'
2020-10-15T15:34:55.864493-0400 port-accesssd[3361] <DEBUG>
debug|LOG_DEBUG|MSTR|1|PORTACCESS|PORTACCESS_DOT1X_PROTOCOL|logID=8967 dot1xpae SM State
transition [REAUTHENTICATING] -> [CACHED-REAUTHENTICATED] for user 'asts-employee-vsa' with
key '1/1/4, 00:50:56:b1:95:30'
2020-10-15T15:34:56.863524-0400 port-accesssd[3361] <DEBUG>
debug|LOG_DEBUG|MSTR|1|PORTACCESS|PORTACCESS_DOT1X_PROTOCOL|logID=8968 dot1xpae SM State
transition [CACHED-REAUTHENTICATED] -> [REAUTHENTICATING] for user 'asts-employee-vsa' with
key '1/1/4, 00:50:56:b1:95:30'
```

```
2020-10-15T15:35:06.875245-0400 port-accessd[3361] <DEBUG>
debug|LOG_DEBUG|MSTR|1|PORTACCESS|PORTACCESS_DOT1X_PROTOCOL|logID=8981 dot1xpae SM State
transition [REAUTHENTICATING] -> [UNAUTHENTICATED] for user 'asts-employee-vsa' with key
'1/1/4, 00:50:56:b1:95:30'
```

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.

```
Sw2(config-if)# interface 1/1/4
Sw2(config-if)# aaa authentication port-access dot1x authenticator
Sw2(config-if-dot1x-auth)# no cached-reauth
Sw2(config-if-dot1x-auth)# exit
Sw2(config-if)# exit
```

35. Cleanup: On Agg1, remove the ACL on port 1/1/7

```
Agg1(config-if)# interface 1/1/7
Agg1(config-if)# no apply access-list ip radius out
```

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.

```
Sw2(config)# do port-access reauthenticate interface 1/1/4
Sw2(config)# show port-access clients
```

Port Access Clients

Port	MAC Address	Onboarded Method	Status	Role
1/1/4	00:50:56:b1:95:30	dot1x	Success	employee

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).

```
Sw2(config)# radius-server tracking user-name radius-tracking-user password plaintext aruba123
```

38. Configure the tracking interval to 60 seconds

```
Sw2(config)# radius-server tracking interval 60
```

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.

```
Sw2(config)# radius-server host 10.252.1.23 tracking enable tracking-mode 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.

```
Sw2(config)# show port-access clients detail
```

Port Access Client Status Details:

Client 00:50:56:b1:95:30, **asts-employee-vsa**

=====

Session Details

Port : 1/1/4
Session Time : 727s
IPv4 Address :
IPv6 Address :

Authentication Details

Status : dot1x Authenticated
Auth Precedence : dot1x - Authenticated, mac-auth - Not attempted

Authorization Details

Role : employee
Status : Applied

Role Information:

```
Name : employee
Type : local
-----
Reauthentication Period      : 30 secs
Authentication Mode          :
Session Timeout              :
Client Inactivity Timeout    :
Description                  :
Gateway Zone                 :
UBT Gateway Role             :
Access VLAN                  : 11
----- Output Omitted -----
```

41. On Sw2, enable debug for radius tracking

```
Sw2(config)# debug 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.

```
Sw2# diagnostics
Sw2# diag utilities tcpdump host-port 1812
```

44. On Agg1, apply the RADIUS ACL on port 1/1/7 to block the RADIUS communication again.

```
Agg1(config-if)# interface 1/1/7
Agg1(config-if)# apply access-list ip radius out
```

45. Switch to the original Sw2 session, review the RADIUS server details.

```
Sw2(config)# show radius-server detail
***** Global RADIUS Configuration *****

Shared-Secret: None
Timeout: 5
Auth-Type: pap
Retries: 1
TLS Timeout: 5
Tracking Time Interval (seconds): 60
Tracking Retries: 1
Tracking User-name: radius-tracking-user
Tracking Password: AQBapVWcNJavUC1NBQenFaJwwRrR+nWcJUvsQ1HUBuaiOv1DCAAAAMCnYwT2Fu1+
Number of Servers: 1

***** RADIUS Server Information *****
Server-Name      : 10.252.1.23
Auth-Port        : 1812
Accounting-Port  : 1813
VRF              : default
TLS Enabled      : No
```



```

Shared-Secret      : AQBapVWcNJavUC1NBQenFaJwwRrR+nWcJUvsQlHUBuaiOv1DCAAAAMCnYwT2Fu1+
Timeout (default)  : 5
Retries (default)  : 1
Auth-Type (default): pap
Server-Group       : cppm
Group-Priority     : 1
Tracking           : enabled
Tracking-Mode      : dead-only
Reachability-Status: reachable
ClearPass-Username : 
ClearPass-Password : None
  
```

46. Wait about 30 seconds, the re-authentication should occur and timeout.

```
Sw2(config)# show radius-server detail
```

```

----- Output Omitted -----
Tracking           : enabled
Tracking-Mode      : dead-only
Reachability-Status: unreachable
ClearPass-Username : 
ClearPass-Password : None
  
```

47. On Sw2, verify the client now has now failed authentication.

```
Sw2(config)# show port-access clients
```

Port Access Clients

Port	MAC Address	Onboarded Method	Status	Role
1/1/4	00:50:56:b1:95:30		Success	asts-critical, Critical

48. In the 'trace session', review the RADIUS packets that are sent by the switch.

```

14 15:49:32.465001 IP (tos 0x0, ttl 64, id 12320, offset 0, flags [DF], proto UDP (17),
length 121)
  TX-Sw2.55408 > 10.252.1.23.radius: RADIUS, length: 93
    Access-Request (1), id: 0x04, Authenticator: 39f7a4c6773fd44e1edb507a6b3b6f48
      User-Name Attribute (1), length: 22, Value: radius-tracking-user
      User-Password Attribute (2), length: 18, Value:
      Message-Authenticator Attribute (80), length: 18, Value: I.1..y.."2;.-.g8
      NAS-Identifier Attribute (32), length: 9, Value: TX-Sw2
      NAS-IP-Address Attribute (4), length: 6, Value: TX-Sw2
  
```

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.

```
Sw2# show radius-server detail
```


50. On Sw2, review the debug buffer for the radius module. You should notice 'DEAD' messages.

```
Sw2# show debug buffer module radius
```

```
2020-10-15:15:47:16.400082|radius-srv-trkd|LOG_DEBUG|MSTR|1|RADIUS|RADIUS_TRACKING|logID=8605
Handling event 'Radius Server Update Authentication Stats' in state 'DEAD' for Server with
addre
ss '10.252.1.23', VRF Id '0', and udp_authport '1812'
```

51. On Sw2, check the last event log messages that include 'radius'

```
Sw2(config)# show event -r -n 100 | i radius
```

```
2020-10-15T15:46:42.381837-04:00 TX-Sw2 radius-srv-trkd[3165]:
Event|2304|LOG_INFO|MSTR|1|RADIUS Server with Address:10.252.1.23, Authport:1812, VRF_ID:0 is
"unreachable"
2020-10-15T15:43:04.448699-04:00 TX-Sw2 radius-srv-trkd[3165]:
Event|2304|LOG_INFO|MSTR|1|RADIUS Server with Address:10.252.1.23, Authport:1812, VRF_ID:0 is
"reachable"
```

Recover the link and verify the operation

52. On Agg1, remove the ACL application from the port 1/1/7.

```
Agg1(config-if)# interface 1/1/7
```

```
Agg1(config-if)# no apply access-list ip radius out
```

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'.

```
21 15:53:27.608269 IP (tos 0x0, ttl 64, id 41078, offset 0, flags [DF], proto UDP (17),
length 121)
  TX-Sw2.55408 > 10.252.1.23.radius: RADIUS, length: 93
    Access-Request (1), id: 0x08, Authenticator: 5d4ced537d245e6cfcdb97998a62
      User-Name Attribute (1), length: 22, Value: radius-tracking-user
      User-Password Attribute (2), length: 18, Value:
      Message-Authenticator Attribute (80), length: 18, Value: ....."'.....8
      NAS-Identifier Attribute (32), length: 9, Value: TX-Sw2
      NAS-IP-Address Attribute (4), length: 6, Value: TX-Sw2
  22 15:53:27.635235 ethertype IPv4, IP (tos 0x0, ttl 60, id 53107, offset 0, flags [DF],
proto UDP (17), length 106)
    10.252.1.23.radius > TX-Sw2.55408: RADIUS, length: 78
      Access-Accept (2), id: 0x08, Authenticator: dafd0aa868255f48c30f46fa4d69f4a1
      Class Attribute (25), length: 58, Value: ?03.|.B..n..PD5w..
```

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.

```
Sw2# show event -r -n 100 | i radius
```

```
2020-10-15T15:53:27.636818-04:00 TX-Sw2 radius-srv-trkd[3165]:
Event|2304|LOG_INFO|MSTR|1|RADIUS Server with Address:10.252.1.23, Authport:1812, VRF_ID:0 is
"reachable"
2020-10-15T15:46:42.381837-04:00 TX-Sw2 radius-srv-trkd[3165]:
Event|2304|LOG_INFO|MSTR|1|RADIUS Server with Address:10.252.1.23, Authport:1812, VRF_ID:0 is
"unreachable"
2020-10-15T15:43:04.448699-04:00 TX-Sw2 radius-srv-trkd[3165]:
Event|2304|LOG_INFO|MSTR|1|RADIUS Server with Address:10.252.1.23, Authport:1812, VRF_ID:0 is
"reachable"
```

Sw2# show debug buffer module radius

```
2020-10-15:15:56:51.700156|radius-srv-trkd|LOG_INFO|MSTR|1|RADIUS|RADIUS_TRACKING|Handling
over the event 14 to component Radius Server
2020-10-15:15:56:51.700184|radius-srv-trkd|LOG_DEBUG|MSTR|1|RADIUS|RADIUS_TRACKING|logID=9193
Handling event 'Radius Server Update Authentication Stats' in state 'ALIVE' for Server with
address '10.252.1.23', VRF Id '0', and udp_authport '1812'
```

Sw2# show radius-server detail

```
----- Output Omitted -----
Server-Group          : cppm
Group-Priority        : 1
Tracking              : enabled
Tracking-Mode         : dead-only
Reachability-Status   : reachable
ClearPass-Username    :
ClearPass-Password    : None
```

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

Sw2# no debug all

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.

```
Sw2# show port-access client
```

```
Port Access Clients
```

Port	MAC Address	Onboarded Method	Status	Role
1/1/4	00:50:56:b1:95:30	dot1x	Success	employee

2. On Sw2, disable 802.1X on the port 1/1/4.

```
Sw2# conf t
Sw2(config)# interface 1/1/4
Sw2(config-if)# aaa authentication port-access dot1x authenticator disable
Sw2(config-if)# exit
```

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.

```
Sw2(config)# mac-group vm
Sw2(config-mac-group)# seq 10 match mac-oui 00:50:56
```

4. On Sw2, define a new user role 'printer' and assign it to access VLAN 11.

```
Sw2(config-mac-group)# port-access role printer
Sw2(config-pa-role)# vlan access 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.

```
Sw2(config-pa-role)# port-access device-profile dp-printer
Sw2(config-device-profile)# associate role printer
Sw2(config-device-profile)# associate mac-group vm
Sw2(config-device-profile)# enable
Sw2(config-device-profile)# end
```

6. On Sw2, review the currently active port-access clients. There should be no active clients.

```
Sw2# show port-access clients
```

Port Access Clients				
Port	MAC Address	Onboarded Method	Status	Role

7. On Sw2, enable device profile debugging and clear the buffer.

```
Sw2# debug portaccess deviceprofile
Sw2# debug portaccess services
Sw2# clear debug 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)

```
Sw2# conf t
Sw2(config)# interface 1/1/4
Sw2(config-if)# port-access device-profile mode block-until-profile-applied
Sw2(config-if)# end
```

9. On Sw2, verify the device profile was applied.

```
Sw2# show port-access device-profile interface all

Port 1/1/4, Neighbor-Mac 00:50:56:b1:95:30
Profile Name:           : dp-printer
LLDP Group:             :
CDP Group:              :
MAC Group:              : vm
Role:                   : printer
State:                  : applied
Failure Reason:         :
```

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

```
Sw2# show port-access clients

Port Access Clients
-----
Port      MAC Address      Onboarded      Status      Role
          MAC Address      Method
-----
1/1/4     00:50:56:b1:95:30  device-profile Success      printer
```

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'.

```
Sw2# show debug buffer severity info

2020-10-15:16:39:21.350407|port-
accessd|LOG_INFO|MSTR|1|PORTACCESS|PORTACCESS_SERVICES|logID=14785 Bypass role is set in port
1/1/4 for client 00:50:56:b1:95:30 to role 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.

```
Sw2# conf t
Sw2(config)# port-access role asts-fallback
Sw2(config-pa-role)# vlan access 11
```

13. On Sw2, on port 1/1/4, configure the fallback role.

```
Sw2(config-pa-role)# interface 1/1/4
Sw2(config-if)# port-access fallback-role asts-fallback
```

14. On Sw2, remove the PC MAC from the MAC Group.

```
Sw2(config-if)# mac-group vm
Sw2(config-mac-group)# seq 10 ignore mac-oui 00:50:56
```

15. On Sw2, verify the client has now received the fallback role.

```
Sw2(config-mac-group)# show port-access clients
```

Port	MAC Address	Onboarded Method	Status	Role
1/1/4	00:50:56:b1:95:30		Success	asts-fallback, Fallback

16. Review the debug buffer, the client is marked with 'Authentication failed but with role'.

```
Sw2(config-mac-group)# show debug buffer severity info
```

```
2020-10-15:16:43:52.943944|port-
accessd|LOG_INFO|MSTR|1|PORTACCESS|PORTACCESS_SERVICES|logID=15077 Secure client state
changed for client 00:50:56:b1:95:30 in port 1/1/4 to 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.

```
Sw2(config-mac-group)# mac-group vm
Sw2(config-mac-group)# seq 10 match mac-oui 00:50:56
Sw2(config-mac-group)# end
```

```
Sw2# show port-access clients
```

Port	MAC Address	Onboarded Method	Status	Role
1/1/4	00:50:56:b1:95:30	device-profile	Success	printer

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.

```
Sw2# conf t
Sw2(config)# interface 1/1/4
Sw2(config-if)# aaa authentication port-access dot1x authenticator enable
```

19. Since AAA port-access is now enabled on the port, you may remove the local MAC authentication command. (It is now 'automatically' enabled).

```
Sw2(config-if)# interface 1/1/4
Sw2(config-if)# no port-access device-profile mode block-until-profile-applied
Sw2(config-if)# end
```

20. Re-authenticate 802.1X on port 1/1/4

```
Sw2# do port-access reauthenticate interface 1/1/4
```

21. Review the authenticated clients.

```
Sw2# show port-access clients
```

Port Access Clients				
Port	MAC Address	Onboarded Method	Status	Role
1/1/4	00:50:56:b1:95:30	dot1x	Success	employee

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.

```
Sw2# debug portaccess all
Sw2# clear 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.

```
Sw2# show port-access clients
```

Port Access Clients

Port	MAC Address	Onboarded Method	Status	Role
1/1/4	00:50:56:b1:95:30	device-profile	Success	printer

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.

```
Sw2# show debug buffer severity info
```

```
2020-10-15:16:54:07.809961|port-
accessd|LOG_INFO|MSTR|1|PORTACCESS|PORTACCESS_DOT1X_RADIUS|Received Access-Reject response
for the request context 0xfffffad0128e0 client request id 912 radius id 143 sent by 802.1x
2020-10-15:16:54:07.819127|port-
accessd|LOG_INFO|MSTR|1|PORTACCESS|PORTACCESS_SERVICES|logID=16004 Bypass role is set in port
1/1/4 for client 00:50:56:b1:95:30 to role printer.
2020-10-15:16:54:07.821411|port-
accessd|LOG_INFO|MSTR|1|PORTACCESS|PORTACCESS_SERVICES|logID=16004 Secure client state
changed for client 00:50:56:b1:95:30 in port 1/1/4 to Authenticated(with role).
```

26. On Sw2, review the client details

```
Sw2# show port-access clients detail
```


Port Access Client Status Details:

Client 00:50:56:b1:95:30, asts-reject

```
=====
Session Details
-----
Port      : 1/1/4
Session Time : 1028s
IPv4 Address :
IPv6 Address :

Authentication Details
-----
Status      : Authenticated
Auth Precedence : dot1x - Unauthenticated, mac-auth - Not attempted

Authorization Details
-----
Role       : printer
Status    : Applied

Role Information:

Name      : printer
Type     : local
----- Output Omitted -----
```

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.

Sw2# show port-access clients

Port Access Clients

Port	MAC Address	Onboarded Method	Status	Role
1/1/4	00:50:56:b1:95:30	dot1x	Success	employee

29. On Agg1, apply the RADIUS ACL on port 1/1/7.

```
Agg1(config-if)# interface 1/1/7
Agg1(config-if)# apply access-list ip radius out
```

30. On Sw2, clear the debug buffer.

```
Sw2# clear 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).

```
Sw2# show port-access clients
```

Port Access Clients

Port	MAC Address	Onboarded Method	Status	Role
1/1/4	00:50:56:b1:95:30	device-profile	Success	printer

32. On Sw2, review the debug buffer INFO messages. After the RADIUS timeout, the Bypass role will be applied.

```
Sw2# show debug buffer severity info
```

```
2020-10-15:17:01:54.150946|port-
accessd|LOG_INFO|MSTR|1|PORTACCESS|PORTACCESS_DOT1X_RADIUS|Received authentication request
Access-Request(New) from 802.1x to send to group cppm
2020-10-15:17:01:54.150969|port-
accessd|LOG_ERR|MSTR|1|PORTACCESS|PORTACCESS_DOT1X_RADIUS|RADIUS request failed with error
Server timeout
```

33. Cleanup: On Agg1, remove the ACL

```
Agg1(config-if)# interface 1/1/7
Agg1(config-if)# no apply access-list ip radius out
```

34. Cleanup: On Sw2, bounce port 1/1/4 and verify the 'employee' role was applied. Disable all debugging.

```
Sw2# show port-access clients
```

Port Access Clients

Port	MAC Address	Onboarded Method	Status	Role
1/1/4	00:50:56:b1:95:30	dot1x	Success	employee

```
Sw2# no debug all
```

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.

```
Sw2# conf t
Sw2(config)# interface 1/1/4
Sw2(config-if)# shutdown
```

2. Login to the Aruba Gateway using admin/aruba123
3. On the Aruba Gateway, review the current VLAN list

```
Aruba_gw# show vlan
VLAN CONFIGURATION
-----
VLAN  Description  Ports          AAA Profile  Option-82
----  -
1      Default          GE0/0/3 Pc0-7  N/A          Disabled
2      VLAN0002          GE0/0/0        N/A          Disabled
11     VLAN0011          Pc1            N/A          Disabled
12     VLAN0012          Pc1            N/A          Disabled
23     VLAN0023          Pc1            N/A          Disabled
24     VLAN0024          Pc1            N/A          Disabled
40     VLAN0040          Pc1            N/A          Disabled
101    VLAN0101          Pc1            N/A          Disabled
```

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.

```
Sw2(config-if)# ip source-interface ubt interface vlan1
```

5. On Sw2, enable ubt debugging for the control plane and hw categories

```
Sw2# clear debug buffer
Sw2# debug ubt cp
Sw2# debug ubt hw
Sw2# debug portaccess all
```

6. On Sw2, configure the UBT zone.

```
Sw2(config)# ubt zone aruba_gw vrf default
Sw2(config-ubt-aruba_gw)# primary-controller ip 10.x.1.6
Sw2(config-ubt-aruba_gw)# enable
Sw2(config-ubt-aruba_gw)# exit
```

7. On Sw2, verify the state of UBT.

```
Sw2(config)# show ubt state
Local Master Server (LMS) State:

LMS Type      IP Address      State      Role
-----
Primary       : 10.x.1.6

Switch Anchor Controller (SAC) State:
                IP Address      MAC Address      State
-----
Active        :
```

8. On Sw2, review the debug buffer log for UBT.

```
Sw2# show debug buffer module ubt
2020-10-15:17:18:23.804273|tunnelednode|LOG_INFO|MSTR|1|UBT|UBT_CP|TNS Zone Config - INSERT
or MODIFY
2020-10-15:17:18:23.804295|tunnelednode|LOG_ERR|MSTR|1|UBT|UBT_CP|TNS Client VLAN is not
specified
```

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.

```
Sw2(config)# vlan 4000
Sw2(config-vlan-4000)# ubt-client-vlan 4000
Sw2(config)# exit
```

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.

```
Sw2# show ubt state
```

Local Master Server (LMS) State:

LMS Type	IP Address	State	Role
Primary	: 10.x.1.6	ready_for_bootstrap	operational_primary

Switch Anchor Controller (SAC) State:

	IP Address	MAC Address	State
Active	: 10.x.1.6	20:4c:03:b1:ce:ba	Registered

11. On Sw2, review the debug output to see the steps to connect to the Aruba GW

Sw2# show debug buffer module ubt

```
2020-10-15:17:21:12.963327|tunnelednode|LOG_INFO|MSTR|1|UBT|UBT_CP|UserTNode: LMS FW
Discovery - Firmware discovery request sent to Controller: 10.x.1.6
2020-10-15:17:21:12.988117|tunnelednode|LOG_INFO|MSTR|1|UBT|UBT_CP|UserTNode: LMS FW
Discovery - Firmware discovery response received from Controller:10.x.1.6
2020-10-15:17:21:12.999292|tunnelednode|LOG_INFO|MSTR|1|UBT|UBT_CP|SAC state: Bootstrapping,
event: Start

2020-10-15:17:21:13.000332|tunnelednode|LOG_INFO|MSTR|1|UBT|UBT_CP|updating LMS operational
role to: operational_primary
2020-10-15:17:21:13.003735|tunnelednode|LOG_INFO|MSTR|1|UBT|UBT_CP|UserTNode: LMS FW
Discovery - Controller is operational LMS primary: 10.x.1.6
2020-10-15:17:21:13.005424|tunnelednode|LOG_INFO|MSTR|1|UBT|UBT_CP|ASAC IP:          10.x.1.6
2020-10-15:17:21:13.005777|tunnelednode|LOG_INFO|MSTR|1|UBT|UBT_CP|ASAC MAC:
20:4c:03:b1:ce:ba
2020-10-15:17:21:13.010558|tunnelednode|LOG_INFO|MSTR|1|UBT|UBT_CP|Updating route resolution
for single controller.
```

12. On Sw2, review the diag-dump output

Sw2# diagnostic

Sw2# diag-dump ubt basic

----- Output Omitted -----

Active SAC Information

```
SAC Ip           : 10.x.1.6
SAC Mac          : 20:4c:03:b1:ce:ba
SAC Firmware Version : 8.5.0.5
```

----- Output Omitted -----

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.

```
Aruba_gw# show tunneled-node-mgr tunneled-nodes
```

```
----- Output Omitted -----
Name      Tunneled Node Mac  IP Address  Age(d:h:m)  Key    Tunnel Index  SAC IP Address  S-SAC
IP Address A-Users  S-Users  Flags
-----
TX-Sw2    88:3a:30:a4:1f:40  10.x.1.5   00:00:08    deed  tunnel 9      10.x.1.6        0.0.0.0
0          0          AC
```

```
Aruba_gw# show tunneled-node-mgr trace-buf
```

```
TNM Trace Buffer
```

```
-----
Sep 18 21:23:31 *   TNM Process UP
Oct 15 14:20:29 --> SW Bootstrap Req    10.x.1.5  88:3a:30:a4:1f:40  rsvd-vid=1  sacMode=1
sacIP=0.0.0.0 flags=1 mtu=0
Oct 15 14:20:29 *   SW BS Ack not sent  10.x.1.5  Awaiting License Response.
Oct 15 14:20:29 sos  SW hb tun created    10.x.1.5  tunnel 9.
Oct 15 14:20:29 <-- SW Bootstrap Ack    10.x.1.5
```

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.

```
Sw2(config)# port-access role employee
Sw2(config-pa-role)# gateway-zone zone aruba_gw gateway-role authenticated
Sw2(config-pa-role)# no vlan access 11
Sw2(config-pa-role)# no reauth-period
Sw2(config-pa-role)# exit
```

15. On Sw2, bounce port 1/1/4.

```
Sw2(config)# interface 1/1/4
Sw2(config-if)# shut
Sw2(config-if)# no shut
Sw2(config-if)# end
```

16. On Sw2, review the port-access clients

```
Sw2# show port-access clients
```

```
Port Access Clients
```

```
-----
Port      MAC Address      Onboarded      Status      Role
          Address          Method
-----
1/1/4     00:50:56:b1:95:30 dot1x          Success     employee
```

17. On Sw2, review the tunneled users.

```
Sw2# show ubt users all
```

Displaying All UBT Users for Zone: aruba_gw
Downloaded user roles are preceded by *

Port	Mac-Address	Tunnel Status	Secondary-UserRole	Failure Reason
1/1/4	00:50:56:b1:95:30	activated	authenticated	---/---

18. Review the debug buffer to see the normal steps to register the user with the Aruba Gateway.

```
Sw2# show debug buffer module ubt
```

```
----- Output Omitted -----
2020-10-15:17:33:23.549044|tunnelednode|LOG_INFO|MSTR|1|UBT|UBT_CP|User: 00:50:56:b1:95:30
exists with TNS Client state: registered
2020-10-15:17:33:23.549067|tunnelednode|LOG_DEBUG|MSTR|1|UBT|UBT_CP|UpdateUserCountOrDelUser
hw_status 1, state registered
2020-10-15:17:33:23.549089|tunnelednode|LOG_DEBUG|MSTR|1|UBT|UBT_CP|ActiveTunnel: ab030a0c-
ee5d-4c4f-b38c-54851b458eee, UACID 0, MAC 00:50:56:b1:95:30, Port 4, user_count 1
```

19. On the Aruba Gateway, review the trace buffer. You should see a 'User Bootstrap Req' request.

```
Aruba_gw# show tunneled-node-mgr trace-buf
```

```
Oct 15 14:49:45 --> User Unbootstrap Req 10.x.1.5 00:50:56:b1:95:30 reason=5, key=4.
Oct 15 14:49:45 sos User tunnel removed 10.x.1.5 00:50:56:b1:95:30 tunnel 10.
Oct 15 14:49:45 gsm Delete tun user 10.x.1.5 00:50:56:b1:95:30.
Oct 15 14:49:45 <-- User Unbootstrap Ack 10.x.1.5 00:50:56:b1:95:30 key=4
status=1:Success.
Oct 15 14:49:50 --> User bootstrap req 10.x.1.5 00:50:56:b1:95:30 rsvd-vid=1 vlan=4000
key=4 role=authenticated flags=6 mtu=1500 server=0.0.0.0.
Oct 15 14:49:50 sos User tunnel created 10.x.1.5 00:50:56:b1:95:30 dormant=0 tunnel 10.
Oct 15 14:49:50 gsm Publish tun user 10.x.1.5 00:50:56:b1:95:30.
Oct 15 14:49:50 <-- User bootstrap ack 10.x.1.5 00:50:56:b1:95:30 assignedvlan=12 L2=0
S-UAC=N/A idx=-1 status=1:Success.
```

20. On the Aruba Gateway, review the tunneled-users.

```
Aruba_gw# show tunneled-node-mgr tunneled-users
```

```
----- Output Omitted -----
User Tunnel Index Flags Tunneled User Mac Tunneled Node Mac Vlan UAC IP Address Key
-----
asts-employee-vsa 00:50:56:b1:95:30 88:3a:30:a4:1f:40 4000(12) 10.x.1.6 4 tunnel
10 UAC
```

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.

```
Aruba_gw# show user
```

IP	MAC	Name	Role	Age(d:h:m)	Auth
VPN link	AP name	Roaming	Essid/Bssid/Phy	Profile	
Forward mode	Type	Host Name	User Type		
10.x.12.104	00:50:56:b1:95:30	asts-employee-vsa	authenticated	00:00:03	Tunneled-User-
802.1X	10.x.1.5	Tunneled	tunnel 10/88:3a:30:a4:1f:40/1/1/4	default-tunneled-	user
tunnel		TUNNELED USER			

22. On the Aruba Gateway, review the role derivation for the user

```
Aruba_gw# show user mac 00:50:56:b1:95:30
```

Name: asts-employee-vsa, IP: 10.x.12.104, MAC: 00:50:56:b1:95:30, Age: 00:00:04
 Role: authenticated (how: ROLE_DERIVATION_TUNNELED_USER_ROLE), ACL: 83/0
 Authentication: Yes, status: not started, method: , protocol: , server:
 ----- Output Omitted -----

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

```
Aruba_gw# show rights authenticated
```

Valid = 'Yes'
 CleanedUp = 'No'
 Derived Role = 'authenticated'
 Up BW:No Limit Down BW:No Limit
 L2TP Pool = default-l2tp-pool
 PPTP Pool = default-pptp-pool
 Number of users referencing it = 2
 Assigned VLAN = 12
 ----- Output Omitted -----

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'.

```
Sw2# conf t
Sw2(config)# port-access role employee
Sw2(config-pa-role)# gateway-zone zone aruba_gw gateway-role sales
Sw2(config-pa-role)# end
```

25. On Sw2, review the port-access clients and the detailed output.

```
Sw2# show port-access clients
```

Port Access Clients

Port	MAC Address	Onboarded Method	Status	Role
1/1/4	00:50:56:b1:95:30		Fail	employee

```
Sw2# show port-access clients detail
```

Port Access Client Status Details:

Client 00:50:56:b1:95:30, asts-employee-vsa

=====

Session Details

Port : 1/1/4
Session Time : 351s
IPv4 Address :
IPv6 Address :

Authentication Details

Status : dot1x Authenticated
Auth Precedence : dot1x - Authenticated, mac-auth - Not attempted

Authorization Details

Role : employee
Status : Failed, Failed to setup User Based Tunnel

Role Information:

Name : employee
Type : local

```
-----
Reauthentication Period      :
Authentication Mode          :
Session Timeout              :
Client Inactivity Timeout    :
Description                  :
Gateway Zone                 : aruba_gw
UBT Gateway Role             : sales
-----
```

----- Output Omitted -----

Q: What do you observe?

A: The status is fail.

26. On Sw2, review the debug buffer entries for the UBT module.

Sw2# show debug buffer module UBT

```
2020-10-15:18:12:40.005010|tunnelednode|LOG_DEBUG|MSTR|1|UBT|UBT_CP|Received Bootstrap Ack
for User [00:50:56:b1:95:30]
2020-10-15:18:12:40.005040|tunnelednode|LOG_ERR|MSTR|1|UBT|UBT_CP|User 00:50:56:b1:95:30
bootstrap nack reason:12X
2020-10-15:18:12:40.005059|tunnelednode|LOG_ERR|MSTR|1|UBT|UBT_CP|User 00:50:56:b1:95:30
bootstrap nack reason RC_UB_FAIL_AUTH_ENTRY_CREATE_FAILED :12X
```

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.

Aruba_gw# show log errorlog 10

```
Oct 15 15:11:56 <authmgr 199802> <6699> <ERRS> |authmgr| message.c,
tunneled_user_do_role_and_vlan_derivation:12851: User (00:50:56:b1:95:30) unknown default
role "sales" provided
Oct 15 15:11:56 <authmgr 199802> <6699> <ERRS> |authmgr| message.c,
auth_handle_tunneled_user_add:12685: handle station failed: Mac: 00:50:56:b1:95:30
```

Aruba_gw# show tunneled-node-mgr trace-buf

```
Oct 15 15:18:43 --> User bootstrap req 10.x.1.5 00:50:56:b1:95:30 rsvd-vid=1 vlan=4000
key=4 role=sales flags=6 mtu=1500 server=0.0.0.0.
Oct 15 15:18:43 sos User tunnel created 10.x.1.5 00:50:56:b1:95:30 dormant=0 tunnel 11.
```

```
Oct 15 15:18:43  gsm  Publish tun user      10.x.1.5  00:50:56:b1:95:30.
Oct 15 15:18:43  <--  User bootstrap ack    10.x.1.5  00:50:56:b1:95:30 status=18:User
Bootstrap Failed, Auth Module Could Not Create Entry.
Oct 15 15:18:43  sos  User tunnel removed    10.x.1.5  00:50:56:b1:95:30 tunnel 11.
Oct 15 15:18:43  gsm  Delete tun user      10.x.1.5  00:50:56:b1:95:30.
Oct 15 15:18:43  -->  User bootstrap req    10.x.1.5  00:50:56:b1:95:30 rsvd-vid=1 vlan=4000
key=4 role=sales flags=6 mtu=1500 server=0.0.0.0.
Oct 15 15:18:43  sos  User tunnel created    10.x.1.5  00:50:56:b1:95:30 dormant=0 tunnel 10.
Oct 15 15:18:43  gsm  Publish tun user      10.x.1.5  00:50:56:b1:95:30.
Oct 15 15:18:43  <--  User bootstrap ack    10.x.1.5  00:50:56:b1:95:30 status=18:User
Bootstrap Failed, Auth Module Could Not Create Entry.
Oct 15 15:18:43  sos  User tunnel removed    10.x.1.5  00:50:56:b1:95:30 tunnel 10.
Oct 15 15:18:43  gsm  Delete tun user      10.x.1.5  00:50:56:b1:95:30.
```

Correct the configuration

28. On Sw2, update the user role with the gateway role 'authenticated'.

```
Sw2# conf t
Sw2(config)# port-access role employee
Sw2(config-pa-role)# gateway-zone zone aruba_gw gateway-role authenticated
Sw2(config-pa-role)# exit
```

29. On Sw2, bounce port 1/1/4 and verify the user status is now success.

```
Sw2# show port-access clients

Port Access Clients
-----
Port      MAC Address      Onboarded      Status      Role
          Address          Method
-----
1/1/4     00:50:56:b1:95:30 dot1x          Success     employee
```

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 so set the 'do not fragment' option in the command (-f).

```
ping -f -l 1472 10.x.11.1 -n 2
```

```
Pinging 10.x.11.1 with 1472 bytes of data:
Reply from 10.x.11.1: bytes=1472 time<1ms TTL=64
Reply from 10.x.11.1: bytes=1472 time<1ms TTL=64
```

```
Ping statistics for 10.x.11.1:
    Packets: Sent = 2, Received = 2, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 0ms, Maximum = 0ms, Average = 0ms
```

32. On PC3, attempt to ping with a payload of 1473, you should see an error message.

```
ping -f -l 1473 10.x.11.1 -n 2
```

```
Pinging 10.x.11.1 with 1473 bytes of data:
Packet needs to be fragmented but DF set.
Packet needs to be fragmented but DF set.
```

```
Ping statistics for 10.x.11.1:
    Packets: Sent = 2, Received = 0, Lost = 2 (100% loss),
```

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'

```
Sw2# show port-access clients
```

```
Port Access Clients
```

Port	MAC Address	Onboarded Method	Status	Role
1/1/4	00:50:56:b1:95:30	dot1x	Success	employee

Sw2# show ubt users all

Displaying All UBT Users for Zone: aruba_gw

Downloaded user roles are preceded by *

Port	Mac-Address	Tunnel Status	Secondary-UserRole	Failure Reason

-				
1/1/4	00:50:56:b1:95:30	activated	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.

```
ping 10.x.12.1 -n 2
```

Pinging 10.x.12.1 with 32 bytes of data:

```
Reply from 10.x.12.1: bytes=32 time<1ms TTL=64
Reply from 10.x.12.1: bytes=32 time<1ms TTL=64
```

```
Ping statistics for 10.x.12.1:
    Packets: Sent = 2, Received = 2, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 0ms, Maximum = 0ms, Average = 0ms
```

36. On PC4, ping using the maximum calculated MTU, this should be successful.

```
ping -f -l 1426 10.x.12.1 -n 2
```

```
Pinging 10.x.12.1 with 1426 bytes of data:
Reply from 10.x.12.1: bytes=1426 time<1ms TTL=64
Reply from 10.x.12.1: bytes=1426 time<1ms TTL=64
```

```
Ping statistics for 10.x.12.1:
    Packets: Sent = 2, Received = 2, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 0ms, Maximum = 0ms, Average = 0ms
```

37. On PC4, attempt to ping using an ICMP payload of 1427. This should fail.

```
ping -f -l 1427 10.x.12.1 -n 2
```

```
Pinging 10.x.12.1 with 1427 bytes of data:
Request timed out.
Request timed out.
```

```
Ping statistics for 10.x.12.1:
    Packets: Sent = 2, Received = 0, Lost = 2 (100% loss),
```

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.

```
Sw2# conf t
Sw2(config)# interface 1/1/1-1/1/28
Sw2(config-if-<1/1/1-1/1/28>)# mtu 9198
Sw2(config-if-<1/1/1-1/1/28>)# interface vlan1
Sw2(config-if-vlan)# ip mtu 9198
```

```
Sw2(config-if-vlan)# exit
```

40. On PC4, attempt to ping using maximum expected MTU (ICMP payload of 1472)

```
ping -f -l 1472 10.x.12.1 -n 2
```

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

```
Sw2(config)# ubt zone aruba_gw vrf default
Sw2(config-ubt-aruba_gw)# no enable
Sw2(config-ubt-aruba_gw)# enable
Sw2(config-ubt-aruba_gw)# end
```

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.

```
ping -f -l 1472 10.x.12.1 -n 2

Pinging 10.x.12.1 with 1472 bytes of data:
Reply from 10.x.12.1: bytes=1472 time<1ms TTL=64
Reply from 10.x.12.1: bytes=1472 time<1ms TTL=64

Ping statistics for 10.x.12.1:
    Packets: Sent = 2, Received = 2, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 0ms, Maximum = 0ms, Average = 0ms
```

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.

```
Aruba_gw# show aaa profile default-tunneled-user
AAA Profile "default-tunneled-user" (Predefined (changed))
-----
Parameter                                     Value
-----
Initial role                                wired-default
MAC Authentication Profile                  N/A
MAC Authentication Default Role             guest
MAC Authentication Server Group             default
802.1X Authentication Profile              N/A
802.1X Authentication Default Role          guest
802.1X Authentication Server Group          N/A
Download Role from CPPM                    Disabled
Set username from dhcp option 12           Disabled
L2 Authentication Fail Through              Disabled
Multiple Server Accounting                  Disabled
User idle timeout                          N/A
Max IPv4 for wireless user                  2
RADIUS Accounting Server Group              N/A
RADIUS Roaming Accounting                  Disabled
RADIUS Interim Accounting                  Disabled
RADIUS Acct-Session-Id In Access-Request    Disabled
XML API server                             N/A
RFC 3576 server                             N/A
User derivation rules                       N/A
Wired to Wireless Roaming                  Disabled
Reauthenticate wired user on VLAN change    Disabled
Device Type Classification                  Enabled
Enforce DHCP                               Disabled
PAN Firewall Integration                    Disabled
Open SSID radius accounting                 Disabled
Apply ageout mechanism on bridge mode wireless clients Disabled
```

NOTE: In the lab environment, the 'Logon' 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.

```
Sw2(config)# port-access role employee
```



```
Sw2(config-pa-role)# gateway-zone zone aruba_gw
Sw2(config-pa-role)# end
```

46. On Sw2, review the active tunneled users

```
Sw2# show ubt users all
```

Displaying All UBT Users for Zone: aruba_gw
Downloaded user roles are preceded by *

Port	Mac-Address	Tunnel Status	Secondary-UserRole	Failure Reason
1/1/4	00:50:56:b1:95:30	activated	---/---	---/---

47. On the Aruba Gateway, verify the assigned role of the client.

```
Aruba_gw# show user
```

IP	MAC	Name	Role	Age(d:h:m)	Auth
VPN link	AP name	Roaming	Essid/Bssid/Phy	Profile	
Forward mode	Type	Host Name	User Type		
10.x.12.104	00:50:56:b1:95:30	asts-employee-vsa	wired-default	00:00:03	Tunneled-User-
802.1X	10.x.1.5	Tunneled	tunnel 10/88:3a:30:a4:1f:40/1/1/4	default-tunneled-	
user					
tunnel		TUNNELED USER			

48. On the Aruba Gateway, review the details of the client.

```
Aruba_gw# show user mac 00:50:56:b1:95:30
```

Name: asts-employee-vsa, IP: 10.x.12.104, MAC: 00:50:56:b1:95:30, Age: 00:00:04
Role: wired-default (**how: ROLE_DERIVATION_INITIAL_ROLE**), ACL: 87/0
Authentication: Yes, status: successful, method: Tunneled-User-802.1X, protocol: , server:
----- Output Omitted -----

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.

```
Sw2# conf t
Sw2(config)# class ip cppm
Sw2(config-class-ip)# 10 match any any 10.252.1.23 count
```

2. On Sw2, define an IP class 'any' for any IP traffic.

```
Sw2(config-class-ip)# class ip any
Sw2(config-class-ip)# 10 match any any any
```

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).

```
Sw2(config-class-ip)# port-access policy pac-employee
Sw2(config-pa-policy)# 10 class ip cppm action dscp CS3 action local-priority 3
Sw2(config-pa-policy)# 20 class ip any
```

4. On Sw2, enable debugging for the pacpolicy module.

```
Sw2(config-pa-policy)# debug pacpolicy all
```

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.

```
Sw2# clear interface lag255 statistics
```

```
Sw2# show interface lag255 queues
```

```
Aggregate-name lag255
Aggregated-interfaces :
1/1/25 1/1/26
Speed 20000 Mb/s
```

	Tx Bytes	Tx Packets	Tx Drops
Q0	0	0	0
Q1	683	7	0
Q2	1482	12	0
Q3	0	0	0
Q4	0	0	0
Q5	0	0	0
Q6	0	0	0
Q7	1120	10	0

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.

```
Sw2(config)# port-access role employee
```

```
Sw2(config-pa-role)# associate policy pac-employee
```

7. On Sw2, review the details of the port-access clients. You should notice the Access policy details with the 'cppm' class.

```
Sw2# show port-access clients detail
```

```
----- Output Omitted -----
```

```
Access Policy Details:
```

```
Policy Name   : pac-employee
Policy Type   : Local
Policy Status : Applied
```

SEQUENCE	CLASS	TYPE ACTION
10	cppm	ipv4 dscp CS3 local-priority 3
20	any	ipv4 permit

```
Class Details:
```

```
class ip cppm
  10 match any any 10.252.1.23 count
class ip any
  10 match any any any
```

8. On Sw2, review the debug buffer entries for the pacpolicy module.

```
Sw2# show debug buffer module pacpolicy
```

```
2020-10-15:18:46:34.051905|ops-
switchd|LOG_DEBUG|MSTR|1|PACPOLICY|PACPOLICY|[pac_policy_ofproto_set]: Port 1/1/4 found in
bridge
2020-10-15:18:46:34.052040|ops-
switchd|LOG_DEBUG|MSTR|1|PACPOLICY|PACPOLICY|[pac_policy_apply_need_update]: returns true
2020-10-15:18:46:34.052064|ops-
switchd|LOG_DEBUG|MSTR|1|PACPOLICY|PACPOLICY|[pac_policy_mac_group_client_update]: Received
policy insert/update request for CLIENT:0:50:56:b1:95:30 on port 1/1/4
2020-10-15:18:46:34.052089|ops-
switchd|LOG_DEBUG|MSTR|1|PACPOLICY|PACPOLICY|[pac_policy_mac_group_client_map_update_cfg_inte
rnal]: Call ofproto to apply policy: pac-employee of mac_grp:1 for client:0:50:56:b1:95:30 on
port:1/1/4 direction:IN
----- Output Omitted -----
```

9. On Sw2, review the port-access policy hitcounts

```
Sw2# show port-access policy pac-employee hitcounts client
```

```
----- Output Omitted -----
SEQUENCE      CLASS-ENTRY                                     HIT-COUNT
-----
10             match any any 10.252.1.23 count                 0
```

10. On PC4, ping 10.252.1.23.

11. On Sw2, review the hitcounts again.

```
Sw2# show port-access policy pac-employee hitcounts client
```

```
----- Output Omitted -----
SEQUENCE      CLASS-ENTRY                                     HIT-COUNT
-----
10             match any any 10.252.1.23 count                 4
```

Q: Do you see the hitcount increase?

A: Yes.

12. On Sw2, review the queue statistics of the LAG255 uplink

```
Sw2# show interface lag255 queues
```

```
Aggregate-name lag255
Aggregated-interfaces :
1/1/25 1/1/26
Speed 20000 Mb/s
```

	Tx Bytes	Tx Packets	Tx Drops
Q0	0	0	0
Q1	8740	79	8
Q2	55746	489	0
Q3	312	4	0
Q4	0	0	0
Q5	0	0	0
Q6	0	0	0

Q7

38554

312

0

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

```
ping 10.252.1.23 -t
```

14. On Sw2, review mirror session 1. It should be enabled and have destination set to CPU.

```
Sw2# show mirror 1
Mirror Session: 1
Admin Status: enable
Operation Status: enabled
Source: vlan rx none
Source: vlan tx none
Source: interface rx none
Source: interface tx none
Destination: cpu
Output Packets: 0
Output Bytes: 0
```

```
Sw2# conf t
Sw2(config)# mirror session 1
Sw2(config-mirror-1)# show run current

mirror session 1
    destination cpu
    enable
```

15. On Sw2, define a new policy that will mirror the class 'any' to mirror1.

```
Sw2(config-mirror-1)# policy ubt-mirror
Sw2(config-policy)# class ip any action mirror 1
```

16. On Sw2, apply the policy on the uplink LAG255

```
Sw2(config-policy)# interface lag 255
Sw2(config-lag-if)# apply policy ubt-mirror out
Sw2(config-lag-if)# apply policy ubt-mirror in
Sw2(config-lag-if)# end
```

17. On Sw2, enable Tshark to dump the CPU mirror to a local file

```
Sw2# diagnostics
Sw2# diag utilities tshark 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

```
Sw2# copy tshark-pcap tftp://10.251.x.91/lab11-sw2-ubt.pcap vrf mgmt
```

% Total	% Received	% Xferd	Average Speed	Time	Time	Time	Current
			Dload Upload	Total	Spent	Left	Speed
100 14320	0	0 100 14320	0	335k	--:--:--	--:--:--	1398k
100 14320	0	0 100 14320	0	333k	--:--:--	--:--:--	333k

Copied successfully.

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.

No.	Time	Source	Destination	Protocol	Length	Info
1	0.000000000	88:3a:30:a4:1f:40	ArubaAHe_b1:ce:ba	IEEE80...	106	OUI 08:00:09 (Hewlett Packard), PID 0x0007
2	0.118095055	10.12.12.104	10.252.1.23	ICMP	120	Echo (ping) request id=0x0001, seq=1911/30471, ttl=128 (reply in 3)
3	0.120796628	10.252.1.23	10.12.12.104	ICMP	116	Echo (ping) reply id=0x0001, seq=1911/30471, ttl=60 (request in 2)

> Frame 2: 120 bytes on wire (960 bits), 120 bytes captured (960 bits) on interface 0

> Ethernet II, Src: 88:3a:30:a4:1f:40 (88:3a:30:a4:1f:40), Dst: ArubaAHe_b1:ce:ba (20:4c:03:b1:ce:ba)

Internet Protocol Version 4, Src: 10.12.1.5, Dst: 10.12.1.6

0100 = Version: 4

.... 0101 = Header Length: 20 bytes (5)

> Differentiated Services Field: 0x60 (DSCP: CS3, ECN: Not-ECT)

Total Length: 106

Identification: 0x0000 (0)

> Flags: 0x4000, Don't fragment

Time to live: 100

Protocol: Generic Routing Encapsulation (47)

Header checksum: 0xffe2 [validation disabled]

[Header checksum status: Unverified]

Source: 10.12.1.5

Destination: 10.12.1.6

> Generic Routing Encapsulation (Transparent Ethernet bridging)

> Ethernet II, Src: Vmware_b1:95:30 (00:50:56:b1:95:30), Dst: MS-NLB-PhysServer-02_00 (02:02:00:00:00:00)

> 802.1Q Virtual LAN, PRI: 0, DEI: 0, ID: 4000

Internet Protocol Version 4, Src: 10.12.12.104, Dst: 10.252.1.23

0100 = Version: 4

.... 0101 = Header Length: 20 bytes (5)

> Differentiated Services Field: 0x60 (DSCP: CS3, ECN: Not-ECT)

Total Length: 60

Identification: 0x3754 (14164)

> Flags: 0x0000

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.

Aruba Training-Confidential

Cleanup

23. On Sw2, remove the policy in and out from the LAG255

```
Sw2(config)# interface lag 255
Sw2(config-lag-if)# no apply policy ubt-mirror out
Sw2(config-lag-if)# no apply policy ubt-mirror in
Sw2(config-lag-if)# end
```

24. On Sw2, disable all debugging.

```
Sw2# no debug all
```

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.

```
Sw2# conf t
Sw2(config)# interface 1/1/4
Sw2(config-if)# shutdown
```

2. On Sw2, enable all debugging for portaccess and pacpolicy (port access policy). Clear the debug buffer.

```
Sw2# debug portaccess all
Sw2# debug pacpolicy all
Sw2# clear 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.

```
Sw2(config)# radius-server host 10.252.1.23 clearpass-username admin clearpass-password
plaintext aruba123
```

4. On PC4, configure the 802.1X credentials asts-employee-dur / aruba123.
5. On Sw2, enable port 1/1/4

```
Sw2(config)# interface 1/1/4
Sw2(config-if)# no shutdown
```

6. On Sw2, review the port-access clients

```
Sw2# show port-access clients
```

Port Access Clients

Port	MAC Address	Onboarded Method	Status	Role

1/1/4 00:50:56:b1:95:30

In-Progress

7. On Sw2, review the port-access client details

```
Sw2# show port-access clients detail
----- Output Omitted -----
Authorization Details
-----
Role      :
Status    : Not Ready
```

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.

```
Sw2# show event -r -n 10
2020-10-15T19:11:24.033215-04:00 TX-Sw2 port-accessd[3361]:
Event|7709|LOG_WARN|MSTR|1|Certificate cppm.arubatraining.com rejected due to verification
failure (20)
```

9. On Sw2, review the debug buffer for errors.

```
Sw2# show debug buffer severity err
2020-10-15:19:11:24.032831|port-
accessd|LOG_ERR|MSTR|1|PORTACCESS|PORTACCESS_DOT1X_PROTOCOL|Certificate subject name doesn't
match the expected peer's hostname.
2020-10-15:19:11:24.033786|port-accessd|LOG_ERR|MSTR|1|PORTACCESS|PORTACCESS_ROLE|logID=25297
Error SSL peer certificate or SSH remote key was not OK while downloading role
asts_aruba_dur_employee-3118-6
```

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

```
Sw2# show radius-server
```

----- Output Omitted -----

SERVER NAME	TLS	PORT	VRF
10.252.1.23		1812	default

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.

Sw2# show ip dns

----- Output Omitted -----

VRF Name : default

Name Server(s) : 10.252.1.21

----- Output Omitted -----

Sw2# ping cppm.arubatraining.com repetitions 2

PING cppm.arubatraining.com (10.252.1.23) 100(128) bytes of data.

108 bytes from 10.252.1.23: icmp_seq=1 ttl=60 time=2.21 ms

108 bytes from 10.252.1.23: icmp_seq=2 ttl=60 time=2.49 ms

--- cppm.arubatraining.com ping statistics ---

2 packets transmitted, 2 received, 0% packet loss, time 1000ms

rtt min/avg/max/mdev = 2.217/2.354/2.492/0.145 ms

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.

Sw2# conf t

Sw2(config)# radius-server host cppm.arubatraining.com key plain aruba123

Sw2(config)# radius-server host cppm.arubatraining.com tracking enable tracking-mode dead-only

Sw2(config)# radius-server host cppm.arubatraining.com clearpass-username admin clearpass-password plaintext aruba123

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.

```
Sw2(config)# aaa group server radius cppm
Sw2(config-sg)# server cppm.arubatraining.com
Sw2(config-sg)# no server 10.252.1.23
```

16. On Sw2, bounce port 1/1/4 and wait for the port-access client to show in the output.

```
Sw2# show port-access client detail
```

Port Access Clients				
Port	MAC Address	Onboarded Method	Status	Role
1/1/4	00:50:56:b1:95:30		In-Progress	

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'.

```
Sw2# show debug buffer
----- Output Omitted -----
2020-10-15:19:20:59.899320|port-
accessd|LOG_DEBUG|MSTR|1|PORTACCESS|PORTACCESS_DOT1X_PROTOCOL|Cert chain failed crypto
validation.
----- Output Omitted -----
```

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.

```
Sw2# conf t
Sw2(config)# interface 1/1/4
Sw2(config-if)# shutdown
```

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.

```
Sw2(config)# crypto pki ta-profile cppm
Sw2(config-ta-cppm)# ta-certificate import terminal
```

Paste the certificate in PEM format below, then hit enter and ctrl-D:

22. Paste the copied root cert text and press CTRL-D. Confirm the installation and exit the TA profile.

```
(config-ta-cert)# -----BEGIN CERTIFICATE-----
EzARBgNVBAGMCKNhbgG1mb3JuaWExEjAQBgNVBACMCVn1bm55dmFsZTEYMBYGA1UE
(config-ta-cert)# MIIETCCAwwGgAwIBAgIBejANBgkqhkiG9w0BAQ0FADCBkzELMAkGA1UEBhMCVVMx
(config-ta-cert)# EzARBgNVBAGMCKNhbgG1mb3JuaWExEjAQBgNVBACMCVn1bm55dmFsZTEYMBYGA1UE
(config-ta-cert)# CgwPQXJ1YmEgRWR1Y2F0aW9uMQwwCgYDVQQLDANFZHUxEDA0BgNVBAMMB0F0DU1Ag
(config-ta-cert)# Q0ExITAfBgkqhkiG9w0BCQEWEmFjc3AtY2FAYWNzcC5sb2NhbDAeFw0yMDAxMjc3
(config-ta-cert)# ODE5NDRAfW0zMDAxMjc3ODQ5NDRAIGTMQswCQYDVQGEWJVUzETMBEGA1UECAwK
(config-ta-cert)# Q2FsaWZvcn5pYTESMBAGA1UEBwwJU3Vubn12YWx1MRgwFgYDVQKDA9BcnViYSBF
(config-ta-cert)# ZHVjYXRpb24xDDAKBgNVBASMA0VkdTEQMA4GA1UEAwwHQUNTUCBDQTEhMB8GCSqG
(config-ta-cert)# SIb3DQEFJARYSYWNzcC1jYUBhY3NwLmxvY2F5MIIBIjANBgkqhkiG9w0BAQEFAAO
(config-ta-cert)# AQ8AMIIBCgKCAQEAwwpy0sCEii0f8eSNjHI6WcEQ+EOKHYMoNROGWy2YmhFqcAEf
(config-ta-cert)# V0Df9Mto0tPIJGth3CmQcZELBz+8WpSdaJv1lqj9QbNjZNk1T44a15Sbv59p0yUM
(config-ta-cert)# XsnJj+KwUmLWL/SUVxM8R81pAK36wRAYAeNr+plVj1RpIxU12xHsQv8qMTypShkw
(config-ta-cert)# YfEs+mCl/VKZ8iWtoupRX+iuWUhBXj39+8gM7c+8sSiImWYmkLlQqfXzxN9TQ4hj
j10sC1vzysqjfXk7KbrJU+XE0z4meiNxWmemmQIDAQAB03YwdADBgNVHQ4EFgQU
(config-ta-cert)# Xc6Hk4ovxUri2Yz01ZXtx5zJYKz/Aq1Z1gfUdBKyNFZwRPRt4F0mLKMsQnyG19Cg
LUqHP2YwZFqWIWtNr3agBEievCj0AQNSdFNwKU7BS6B0BaT2RBWyy6x5U71jBKA1
(config-ta-cert)# j10sC1vzysqjfXk7KbrJU+XE0z4meiNxWmemmQIDAQAB03YwdADBgNVHQ4EFgQU
(config-ta-cert)# XYnMTUJ1JxEmOqT8QoA1pYb8KEwDwYDVR0PAQH/BAUDAwEADAPBgNVHRMBAf8E
(config-ta-cert)# BTADAQH/MDEGA1UdJQQqMCgGCCsGAQUFBwMJBggrBgEFBQcDAgYIKwYBBQUHAE
(config-ta-cert)# CCsGAQUFBwMDMA0GCSqGSIb3DQEBDQUAA4IBAQAjLKEREPa/QGVdP5xi1tcMe7LI
(config-ta-cert)# jaG987qoFmCwi9ocDhfMMggNOny9m0uW4GNSHeBSPIcTgN+JYjjBuN4+V5Iz74Z
(config-ta-cert)# LUqHP2YwZFqWIWtNr3agBEievCj0AQNSdFNwKU7BS6B0BaT2RBWyy6x5U71jBKA1
(config-ta-cert)# OQH5xzcNyRpetbTIdf0q5Aa0kC7QmbYYGxBamwbzS5luU+akQy2HD9vr2bi+X5Aq
(config-ta-cert)# jVBKTUKTPuFPQhQroXp2VI3ldIPJYwWz1Gk5NTAQkv/9bImdr5f4NFJrMI/kubjy
(config-ta-cert)# bdXTI+LyPtiLpkbF2dDCN8GBKovNQFvhlXtGT4b/umOzSS2AzHcVj/SldVw6
```

```
(config-ta-cert)# -----END CERTIFICATE-----
(config-ta-cert)#
The certificate you are importing has the following attributes:
Subject: C = US, ST = California, L = Sunnyvale, O = Aruba Education, OU = Edu, CN = ACSP CA,
emailAddress = acsp-ca@acsp.local
Issuer: C = US, ST = California, L = Sunnyvale, O = Aruba Education, OU = Edu, CN = ACSP CA,
emailAddress = acsp-ca@acsp.local
Serial Number: 0x12
TA certificate import is allowed only once for a TA profile
Do you want to accept this certificate (y/n)? y
(config-ta-cppm)# exit
```

IMPORTANT: You must **'exit'** the context to complete the import process!

23. On Sw2, review the ta profiles

```
Sw2# show crypto pki ta-profile
```

TA Profile Name	TA Certificate	Revocation Check
-----	-----	-----
cppm	Installed, valid	disabled

24. On Sw2, review the details of the cppm profile

```
Sw2# show crypto pki ta-profile cppm
```

```

TA Profile Name      : cppm
Revocation Check    : disabled
  OSCP Primary URL: Not Configured
  OSCP Secondary URL: Not Configured
  OSCP Enforcement-level: strict
  OSCP Disable Nonce: false
  OSCP VRF           : mgmt
TA Certificate       : Installed and valid
Certificate:
  Data:
----- Output Omitted -----
```

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.

```
Sw2(config)# clear debug buffer
Sw2(config)# interface 1/1/4
Sw2(config-if)# no shutdown
Sw2(config-if)# port-access reauthenticate interface 1/1/4
```

26. On Sw2, review the port-access client details.

```
Sw2# show port-access clients detail
```


----- Output Omitted -----

Authorization Details

```
-----
Role      : 
Status    : Not Ready
```

Q: Did the authorization succeed?

A: No, the status is 'Not Ready'.

27. On Sw2, review the last 20 events.

```
Sw2# show event -r -n 20
2020-10-15T19:34:17.512462-04:00 TX-Sw2 port-accesssd[3361]:
Event|7708|LOG_INFO|MSTR|1|Certificate cppm.arubatraining.com verified and accepted
2020-10-15T19:34:11.981547-04:00 TX-Sw2 port-accesssd[3361]:
Event|7708|LOG_INFO|MSTR|1|Certificate cppm.arubatraining.com verified and accepted
```

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.

```
Sw2# show debug buffer
2020-10-15:19:34:17.512076|port-
accesssd|LOG_DEBUG|MSTR|1|PORTACCESS|PORTACCESS_DOT1X_PROTOCOL|Cert chain passed crypto
validation.
2020-10-15:19:34:17.523214|port-
accesssd|LOG_DEBUG|MSTR|1|PORTACCESS|PORTACCESS_ROLE|logID=26889 Result of download operation
for role asts_aruba_dur_employee-3118-6 is No error
2020-10-15:19:34:17.523533|port-
accesssd|LOG_DEBUG|MSTR|1|PORTACCESS|PORTACCESS_ROLE|logID=26889 Parsing the XML for the
cprole asts_aruba_dur_employee-3118-6 failed
2020-10-15:19:34:17.523610|port-
accesssd|LOG_DEBUG|MSTR|1|PORTACCESS|PORTACCESS_SERVICES|logID=26890 Handing over the event 1
to component ClearPass Role
2020-10-15:19:34:17.523717|port-
accesssd|LOG_DEBUG|MSTR|1|PORTACCESS|PORTACCESS_ROLE|logID=26890 DB operation Insert is
successfully completed for role asts_aruba_dur_employee-3118-6
```

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.

```
Sw2# start-shell
Sw2:~$ sudo ls -ls /tmp/*.xml
4 -rw-r--r-- 1 root root 381 Oct 15 19:34 /tmp/asts_aruba_dur_employee-3118-6.xml
Sw2:~$ sudo cat /tmp/*.xml
<!DOCTYPE HTML PUBLIC "-//IETF//DTD HTML 2.0//EN">
<html><head>
<title>401 Unauthorized</title>
</head><body>
<h1>Unauthorized</h1>
<p>This server could not verify that you
are authorized to access the document
requested. Either you supplied the wrong
credentials (e.g., bad password), or your
browser doesn't understand how to supply
the credentials required.</p>
</body></html>
Sw2:~$ exit
exit
```

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)

```
Sw2(config)# radius-server host cppm.arubatraining.com clearpass-username duradmin clearpass-
password plaintext aruba123
```

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.

```
Sw2(config)# interface 1/1/4
Sw2(config-if)# shutdown
Sw2(config-if)# no shutdown
Sw2(config-if)# port-access reauthenticate interface 1/1/4
```

33. On Sw2, review the port-access clients

```
Sw2# show port-access clients
```

Port Access Clients				
Port	MAC Address	Onboarded Method	Status	Role
1/1/4	00:50:56:b1:95:30	dot1x	Success	asts_aruba_dur_employee...

34. On Sw2, review the port-access client details

```
Sw2# show port-access clients detail
```

Port Access Client Status Details:

Client 00:50:56:b1:95:30, asts-employee-dur

```
=====
Session Details
-----
Port      : 1/1/4
Session Time : 49s
IPv4 Address :
IPv6 Address :

Authentication Details
-----
Status      : dot1x Authenticated
Auth Precedence : dot1x - Authenticated, mac-auth - Not attempted

Authorization Details
-----
Role       : asts_aruba_dur_employee-3118-6
```

Status : Applied

----- Output Omitted -----

Access Policy Details:

Policy Name : employee_asts_aruba_dur_employee-3118-6
Policy Type : Downloaded
Policy Status : Applied

SEQUENCE	CLASS	TYPE	ACTION
10	any_asts_aruba_dur_employ...	ipv4	permit

Class Details:

class ip any_asts_aruba_dur_employee-3118-6
10 match any any any

35. On Sw2, review the role details that resulted from the XML file download.

Sw2# show port-access role clearpass

Role Information:

Name : asts_aruba_dur_employee-3118-6
Type : clearpass
Status: Completed

Reauthentication Period	:
Authentication Mode	:
Session Timeout	:
Client Inactivity Timeout	:
Description	:
Gateway Zone	:
UBT Gateway Role	:
Access VLAN	: 11
Native VLAN	:
Allowed Trunk VLANs	:
Access VLAN Name	:
Native VLAN Name	:
Allowed Trunk VLAN Names	:
MTU	:
QOS Trust Mode	:
STP Administrative Edge Port	:
PoE Priority	:
Captive Portal Profile	:
Policy	: employee_asts_aruba_dur_employee-3118-6

Sw2# show port-access policy

Access Policy Details:

=====

Policy Name : employee_asts_aruba_dur_employee-3118-6
Policy Type : Downloaded
Policy Status : Applied

SEQUENCE	CLASS	TYPE	ACTION

10	any_asts_aruba_dur_employ...	ipv4	permit

----- Output Omitted -----

Sw2# show class

----- Output Omitted -----

```
IPv4      any_asts_aruba_dur_employee-3118-6
  10
    match
    any
    any
```

----- Output Omitted -----

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'

```
Agg1# show resources
```

Resource Usage:

Mod	Description Resource	Width	Used	Reserved	Free
1/1	Ingress MAC+IPv4 Port Policy				
	Ingress TCAM Entries	3	3	2304	
	Ingress Control Plane Policing				
	Ingress TCAM Entries	3	246	2304	
	Egress Control Plane Policing				
	Egress TCAM Entries	2	84	1024	
	Total				
	Ingress TCAM Entries		249	4608	4608
	Egress TCAM Entries		84	1024	1024
	Policers		1		6143
	Ingress L4 Port Ranges		0		32

Resource data is updated every 10 seconds.

2. On Agg1, define a new access-list 'servers', allow all traffic and apply it outbound on 1/1/7.

```
Agg1# conf t
Agg1(config)# access-list ip servers
Agg1(config-acl-ip)# 10 permit ip any any
```

```
Agg1(config-acl-ip)# int 1/1/7
Agg1(config-if)# apply access-list ip servers out
```

3. Review the resource usage.

```
Agg1# show resources
```

Resource Usage:

Mod	Description Resource	Width	Used	Reserved	Free
1/1	Routed Egress IPv4 Port ACL				
	Egress TCAM Entries	1	2	512	
	Ingress MAC+IPv4 Port Policy				
	Ingress TCAM Entries	3	3	2304	
	Ingress Control Plane Policing				
	Ingress TCAM Entries	3	246	2304	
	Egress Control Plane Policing				
	Egress TCAM Entries	2	84	1024	
	Total				
	Ingress TCAM Entries		249	4608	4608
	Egress TCAM Entries		86	1536	512
	Policers		1		6143
	Ingress L4 Port Ranges		0		32

Resource data is updated every 10 seconds.

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.

```
object-group ip address servers
10 10.252.1.1
20 10.252.1.2
30 10.252.1.3
40 10.252.1.4
50 10.252.1.5
60 10.252.1.6
70 10.252.1.7
80 10.252.1.8
90 10.252.1.9
100 10.252.1.10
```

5. On Agg1, define an object-group of type 'ip' for the 'clients'. Define 10 entries.

```
object-group ip address clients
```

```
10 10.x.11.0/24
20 10.x.12.0/24
30 10.x.13.0/24
40 10.x.14.0/24
50 10.x.15.0/24
60 10.x.16.0/24
70 10.x.17.0/24
80 10.x.18.0/24
90 10.x.19.0/24
100 10.x.20.0/24
```

6. Define an object-group of type 'port' named 'srvtcp'. Define 5 ports.

```
object-group port srvtcp
eq 25
eq 80
eq 3389
eq 4343
eq 8080
```

7. Access the access-list 'servers', add a new line

```
Agg1# conf t
Agg1(config)# access-list ip servers
Agg1(config-acl-ip)# permit tcp clients servers group srvtcp
Agg1(config-acl-ip)# end
```

8. Review the resources. It may take 10-20 seconds to update the counters.

```
Agg1# show resources
```

Resource Usage:

Mod	Description Resource	Width	Used	Reserved	Free
1/1	Routed Egress IPv4 Port ACL				
	Egress TCAM Entries	1	502	512	
	Ingress MAC+IPv4 Port Policy				
	Ingress TCAM Entries	3	3	2304	
	Ingress Control Plane Policing				
	Ingress TCAM Entries	3	246	2304	
	Egress Control Plane Policing				
	Egress TCAM Entries	2	84	1024	
	Total				
	Ingress TCAM Entries		249	4608	4608
	Egress TCAM Entries		586	1536	512
	Policers		1		6143
	Ingress L4 Port Ranges		0		32

Resource data is updated every 10 seconds.

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)

```
Agg1# conf t
Agg1(config)# access-list ip servers
Agg1(config-acl-ip)# permit tcp clients servers group srvtcp
```

10. On Agg1, use 'show resources' to verify 500 entries were added.

```
Agg1# show resources
Resource Usage:
Mod  Description
      Resource                                Width  Used  Reserved  Free
-----
1/1  Routed Egress IPv4 Port ACL
      Egress TCAM Entries                      1    1002    1024
      Ingress MAC+IPv4 Port Policy
      Ingress TCAM Entries                     3         3    2304
      Ingress Control Plane Policing
      Ingress TCAM Entries                     3    246    2304
      Egress Control Plane Policing
      Egress TCAM Entries                     2     84    1024
      Total
      Ingress TCAM Entries                    249    4608    4608
      Egress TCAM Entries                    1086    2048         0
      Policers                               1    6143
      Ingress L4 Port Ranges                   0         32
```

Resource data is updated every 10 seconds.

11. On Agg1, add the line 1 more time. This will exhaust the TCAM resources.

```
Agg1(config-acl-ip)# permit tcp clients servers group srvtcp
```

12. On Agg1, review the event log. You should see a TCAM error.

```
Agg1# show event -r -s err
2020-10-16T06:38:18.311502-04:00 TX-Agg1 ops-switchd[3274]: Event|10202|LOG_ERR|AMM|1/1|TCAM
entry installation has failed in table Routed_Egress_IPv4_Port_ACL
```

13. On Agg1, review the running configuration of port 1/1/7

```
Agg1# show run int 1/1/7
interface 1/1/7
```



```
no shutdown
! access-list ip servers user configuration does not match active configuration.
! run 'access-list TYPE NAME reset' or 'access-list all reset' to reset the ACL to match the
active configuration.
  apply access-list ip servers out
  ip address 10.255.101.2/24
exit
```

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.

```
Agg1# show access-list configuration
----- Output Omitted -----
% Warning: servers user configuration does not match active configuration.
%      run 'access-list TYPE NAME reset' or 'access-list all reset' to reset the ACL to
match the active configuration.
IPv4
  servers
    10
      permit          any
      any
      any
    20
      permit          tcp
      clients          srvtcp
      servers
    30
      permit          tcp
      clients          srvtcp
      servers
    40
      permit          tcp
      clients          srvtcp
      servers          srvtcp
```

15. Remove the ACL from port 1/1/7

```
Agg1(config)# interface 1/1/7
Agg1(config-if)# no apply access-list ip servers out
```

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.

```
Sw2# show class
----- Output Omitted -----

IPv4      cppm
      10
      match          any
      any
      10.252.1.23
      Hit-counts: enabled

----- Output Omitted -----
```

17. On Sw2, define a policy named 'asts-count1' and add the class 'cppm'. No action need to be set.

```
Sw2# conf t
Sw2(config)# policy counter1
Sw2(config-policy)# 10 class ip cppm
```

18. On Sw2, apply the policy on interfaces 1/1/1 and 1/1/27 inbound.

```
Sw2(config-policy)# interface 1/1/4
Sw2(config-if)# apply policy counter1 in
Sw2(config-if)# interface 1/1/27
Sw2(config-if)# apply policy counter1 in
Sw2(config-if)# end
```

19. On Sw2, review the current hitcounts of the policy.

```
Sw2# show policy hitcounts counter1

Statistics for Policy counter1:
Interface 1/1/4* (in):
      Hit Count  Configuration
10 class ip cppm
      0 10 match any any 10.252.1.23 count
Interface 1/1/27* (in):
      Hit Count  Configuration
10 class ip cppm
      0 10 match any any 10.252.1.23 count
```

20. On PC4, send a single ping to 10.252.1.23. This ping should be counted on 1/1/4.

```
ping 10.252.1.23 -n 1
```

```
Pinging 10.252.1.23 with 32 bytes of data:
Reply from 10.252.1.23: bytes=32 time=3ms TTL=60
```

```
Ping statistics for 10.252.1.23:
    Packets: Sent = 1, Received = 1, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 3ms, Maximum = 3ms, Average = 3ms
```

21. On Sw2, review the hitcounts.

```
Sw2# show policy hitcounts counter1
Statistics for Policy counter1:
Interface 1/1/4* (in):
    Hit Count Configuration
10 class ip cppm      1 10 match any any 10.252.1.23 count
Interface 1/1/27* (in):
    Hit Count Configuration
10 class ip cppm      1 10 match any any 10.252.1.23 count
```

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 'cppm' to it.

```
Sw2# conf t
Sw2(config)# policy counter2
Sw2(config-policy)# 10 class ip cppm
```

23. Replace the policy 'counter1' on port 1/1/27 with the 'counter2'.

```
Sw2(config-policy)# interface 1/1/27
Sw2(config-if)# apply policy counter2 in
Sw2(config-if)# end
```

24. Clear the all hitcount statistics

```
Sw2# clear policy hitcount all
```

25. On PC4, repeat the test ping to 10.252.1.23.

26. On Sw2, verify that the counters are now unique.

```
Sw2# show policy hitcounts counter1
Statistics for Policy counter1:
Interface 1/1/4* (in):
    Hit Count Configuration
```

```
10 class ip cppm
```

```
1 10 match any any 10.252.1.23 count
```

```
Sw2# show policy hitcounts counter2
```

```
Statistics for Policy counter2:
```

```
Interface 1/1/27* (in):
```

```
Hit Count Configuration
```

```
10 class ip cppm
```

```
0 10 match any any 10.252.1.23 count
```

Aruba Training-Confidential

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.

```
Agg1# ping 10.251.x.4 vrf mgmt repetitions 2
PING 10.251.x.4 (10.251.x.4) 100(128) bytes of data.
108 bytes from 10.251.x.4: icmp_seq=1 ttl=63 time=0.231 ms
108 bytes from 10.251.x.4: icmp_seq=2 ttl=63 time=0.261 ms

--- 10.251.x.4 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1023ms
rtt min/avg/max/mdev = 0.231/0.246/0.261/0.015 ms
```

```
Agg1# ssh 10.251.x.4 vrf mgmt
The authenticity of host '10.251.x.4 (10.251.x.4)' can't be established.
ECDSA key fingerprint is SHA256:CVo5bxU64XJBlHHS636zmr2Ve+UQGIqiuuooFj+AB0k.
Are you sure you want to continue connecting (yes/no/[fingerprint])? Yes
----- Output Omitted -----
```

NOTE: There is no need to actually login, just verify that the device is reachable.

2. On Sw1, configure an ACL named 'copp-mgmt'.

```
Sw1# conf t
Sw1(config)# access-list ip 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.

```
Sw1(config-acl-ip)# permit ip 10.251.x.91/32 any count
Sw1(config-acl-ip)# permit ip 10.251.x.200/32 any count
```

4. Add a rule to deny traffic from any to any with the count action.

```
Sw1(config-acl-ip)# deny any any any count
```

5. Apply the ACL to the control-plane on the VRF mgmt.

```
Sw1(config-acl-ip)# apply access-list ip copp-mgmt control-plane 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.

```
Agg1# ssh 10.251.x.4 vrf mgmt
(you can use control-C to abort the connection) ^C
```

Q: Did this succeed?

A: The traffic should have been blocked by the Control Plane ACL.

8. Review the hitcounts on Sw1

```
Sw1# show access-list hitcounts control-plane vrf mgmt

Statistics for ACL copp-mgmt (ipv4):
VRF mgmt* (control-plane):
    Hit Count  Configuration
          49  10 permit any 10.251.x.91 any count
          97  20 permit any 10.251.x.200 any count
          14  30 deny any any any count
```

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.

```
Sw1(config)# end
Sw1# checkpoint auto 1
Copying configuration: [Success]
Auto checkpoint mode expires in 1 minute(s)
```

10. On Sw1, introduce an error in the mgmt ACL as the first rule (you should be able to do it during one minute)

```
Sw1# conf t
Sw1(config)# access-list ip copp-mgmt
Sw1(config-acl-ip)# show run current
access-list ip copp-mgmt
 10 permit any 10.251.x.91 any count
 20 permit any 10.251.x.200 any count
 30 deny any any any count
Sw1(config-acl-ip)# 1 deny ip any any
```

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.

```
Sw1# show event -r -d hpe-config
2020-10-16T07:20:42.417327-04:00 TX-Sw1 hpe-config[19592]: Event|6801|LOG_INFO|AMM|-|Copying
configs from: checkpoint TEMPAUTOCHECK to: running-config
2020-10-16T07:20:42.416933-04:00 TX-Sw1 hpe-config[19592]: Event|6803|LOG_INFO|AMM|-|Config-
Validation:WARNING: Restoring configuration. Do NOT add any new configuration.
2020-10-16T07:19:38.774198-04:00 TX-Sw1 hpe-config[19536]: Event|6801|LOG_INFO|AMM|-|Copying
configs from: running-config to: checkpoint TEMPAUTOCHECK
```

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.

```
Sw1# show ip ospf neighbors
```

```
OSPF Process ID 1 VRF default
```

```
=====
```

```
Total Number of Neighbors: 4
```

Neighbor ID	Priority	State	Nbr Address	Interface
10.x.200.2	n/a	FULL	10.x.241.2	vlan241
10.x.200.3	n/a	FULL	10.x.242.3	vlan242
10.x.0.2	n/a	FULL	10.x.41.2	vlan41
10.x.0.3	n/a	FULL	10.x.42.3	vlan42

```
Sw1# show ip pim neighbor brief
```

VRF : default				Total number of neighbors : 4	
Interface Time (HH:MM:SS)	Neighbor Secondary Address (IPV4)	Uptime Address (HH:MM:SS) (IPV4)	Expires (HH:MM:SS)	DR Priority	Hold
vlan241 Nill	10.x.241.2	17:53:37	00:01:19	1	00:01:45
vlan42 Nill	10.x.42.3	17:56:48	00:01:33	1	00:01:45
vlan242 Nill	10.x.242.3	17:56:50	00:01:31	1	00:01:45
vlan41 Nill	10.x.41.2	17:53:38	00:01:20	1	00:01:45

13. Define a new ACL named 'copp-vrf-default' on Sw1.

```
Sw1# conf t
Sw1(config)# access-list ip copp-vrf-default
```

14. Add a rule for any SSH and any ICMP traffic.

```
Sw1(config-acl-ip)# permit tcp any any eq 22 count
Sw1(config-acl-ip)# permit icmp any any count
```

15. Apply the policy to the Control Plane of the VRF 'default' on Sw1.

```
Sw1(config-acl-ip)# apply access-list ip copp-vrf-default control-plane vrf default
```

16. On Agg1, attempt to ping and SSH to the Sw1(10.x.41.4) using the default VRF.

```
Agg1# ping 10.x.1.4 repetitions 2
PING 10.x.1.4 (10.x.1.4) 100(128) bytes of data.
108 bytes from 10.x.1.4: icmp_seq=1 ttl=64 time=0.208 ms
108 bytes from 10.x.1.4: icmp_seq=2 ttl=64 time=0.201 ms

--- 10.x.1.4 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1028ms
rtt min/avg/max/mdev = 0.201/0.204/0.208/0.014 ms
```

```
Agg1# ssh 10.x.1.4
The authenticity of host '10.x.1.4 (10.x.1.4)' can't be established.
ECDSA key fingerprint is SHA256:ih+1RT68Jc12DB2+ZuXnbz6vUVP8CwrPXtVwocw0aGo.
Are you sure you want to continue connecting (yes/no/[fingerprint])? yes
```

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.

```
Sw1# show access-list hitcounts control-plane vrf default
Statistics for ACL copp-vrf-default (ipv4):
VRF default* (control-plane):
    Hit Count  Configuration
          14  10 permit tcp any any eq 22 count
           2  20 permit icmp any any count
```

18. Now review your OSPF neighbors and PIM neighbors

```
Sw1# show ip ospf neighbors
No OSPF neighbor found on VRF default.
Sw1# show ip pim neighbor brief
```

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

```
Sw1# clear ip ospf statistics
Sw1# show ip ospf statistics int vlan41

OSPF Process ID 1 VRF default, interface vlan41 statistics (cleared 0h0m24s ago)
=====

Tx Hello Packets      : 2          Rx Hello Packets      : 0
Tx Hello Bytes        : 128        Rx Hello Bytes        : 0
Tx DD Packets         : 0          Rx DD Packets         : 0
Tx DD Bytes           : 0          Rx DD Bytes           : 0
Tx LS Request Packets : 0          Rx LS Request Packets : 0
Tx LS Request Bytes   : 0          Rx LS Request Bytes   : 0
Tx LS Update Packets  : 0          Rx LS Update Packets  : 0
Tx LS Update Bytes    : 0          Rx LS Update Bytes    : 0
Tx LS Ack Packets     : 0          Rx LS Ack Packets     : 0
Tx LS Ack Bytes       : 0          Rx LS Ack Bytes       : 0
----- Output Omitted -----
```

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.

```
Sw1# conf t
Sw1(config)# access-list ip copp-vrf-default
Sw1(config-acl-ip)# permit ospf any any
Sw1(config-acl-ip)# permit pim any any
Sw1(config-acl-ip)# end
```

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.

```
Sw1# show copp-policy
applied copp_policy_name
-----
applied default
```

```
Sw1# show copp-policy default configuration
class                drop priority rate pps burst pkts applied
-----
acl-logging          0           25      25      yes
arp-broadcast         2          1250    1250    yes
arp-protect           2          2075    2075    yes
arp-unicast           3           825     825     yes
bfd-control           5           850     850     yes
bgp                   5           750     750     yes
captive-portal        2          2075    259     yes
dhcp                  2           750     750     yes
erps                  6           225     225     yes
icmp-broadcast-ipv4    2           325     325     yes
icmp-multicast-ipv6    2           475     475     yes
icmp-security-ipv6     2           325     325     yes
icmp-unicast-ipv4      3           225     225     yes
icmp-unicast-ipv6      3           400     400     yes
----- Output Omitted -----
```

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'

Agg1# show copp-policy default configuration

class	drop	priority	rate pps	burst	pkts	applied
acl-logging	0		50	50		yes
arp-broadcast	3		7000	7000		yes
arp-unicast	4		2500	2500		yes
bfd	7		4000	4000		yes
bgp-ipv4	6		1500	1500		yes
bgp-ipv6	6		1500	1500		yes
dhcp-ipv4	1		1000	1000		yes
dhcp-ipv6	1		1000	1000		yes
erps	7		1000	1000		yes
hypertext	5		1000	500		yes
icmp-broadcast-ipv4	3		2000	2000		yes
icmp-multicast-ipv6	3		2000	2000		yes
icmp-unicast-ipv4	4		1000	1000		yes
icmp-unicast-ipv6	4		1000	1000		yes

----- Output Omitted -----

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.

Sw1# show copp-policy statistics

Statistics for CoPP policy 'default':

Totals:

packets passed	: 194248	packets dropped	: 0
Class: default			
packets passed	: 699	packets dropped	: 0
Class: acl-logging			
packets passed	: 0	packets dropped	: 0
Class: arp-broadcast			
packets passed	: 20563	packets dropped	: 0
Class: arp-protect			
packets passed	: 0	packets dropped	: 0
Class: arp-unicast			
packets passed	: 319	packets dropped	: 0

----- Output Omitted -----

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).

```
C:\Users\student>fping 10.x.1.4 -t 1 -w 1 -n 1000
Fast pinger version 3.00
(c) Wouter Dhondt (http://www.kwakkelflap.com)

Pinging 10.x.1.4 with 32 bytes of data every 1 ms:

10.x.1.4: request timed out
10.x.1.4: request timed out
10.x.1.4: request timed out
10.x.1.4: request timed out
10.x.1.4: request timed out
10.x.1.4: request timed out
10.x.1.4: request timed out
10.x.1.4: request timed out
10.x.1.4: request timed out
10.x.1.4: request timed out
10.x.1.4: request timed out
Reply[11] from 10.x.1.4: bytes=32 time=1.6 ms TTL=64
Reply[12] from 10.x.1.4: bytes=32 time=0.3 ms TTL=64
Reply[13] from 10.x.1.4: bytes=32 time=0.3 ms TTL=64
Reply[14] from 10.x.1.4: bytes=32 time=0.3 ms TTL=64
Reply[15] from 10.x.1.4: bytes=32 time=0.2 ms TTL=64
----- Output Omitted -----
Ping statistics for 10.x.1.4:
    Packets: Sent = 1000, Received = 764, Lost = 236 (23% loss)
Approximate round trip times in milli-seconds:
    Minimum = 0.2 ms, Maximum = 4.4 ms, Average = 0.2 ms
```

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.

```
show copp-policy statistics
Statistics for CoPP policy 'default':
```

```
Totals:
  packets passed   : 195880      packets dropped   : 226
----- Output Omitted -----
Class: icmp-unicast-ipv4
  packets passed   : 796          packets dropped   : 226
```

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).

```
C:\Users\student>fping 10.251.x.4 -t 1 -w 1 -n 1000

Fast pinger version 3.00
(c) Wouter Dhondt (http://www.kwakkelflap.com)

Pinging 10.251.x.4 with 32 bytes of data every 1 ms:

Reply[1] from 10.251.x.4: bytes=32 time=0.1 ms TTL=64
Reply[2] from 10.251.x.4: bytes=32 time=0.3 ms TTL=64
----- Output Omitted -----
Ping statistics for 10.251.x.4:
    Packets: Sent = 1000, Received = 1000, Lost = 0 (0% loss)
Approximate round trip times in milli-seconds:
    Minimum = 0.1 ms, Maximum = 9.7 ms, Average = 0.3 ms
```

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.

The screenshot shows the Aruba Analytics dashboard. The breadcrumb navigation is **Analytics > Dashboard > copp-agent > copp-agent.graph_0**. The top right corner displays the IP address **10.251.12.4(T12-Sw1)**. On the left sidebar, the **Analytics** section is selected. The main area shows a table of monitors with the following columns: **Show**, **Monitor**, **Aggregation**, **Time Series**, and **Alert Total**. The table lists 12 monitors, with the first 10 being disabled (checkbox unchecked) and the last two, **icmp_unicast_ipv4 traffic dropped** and **icmp_unicast_ipv4 traffic passed**, being enabled (checkbox checked). All monitors are set to **RATE** aggregation and **Rate_System_copp_statistics** time series, with an **Alert Total** of 0.

Show	Monitor	Aggregation	Time Series	Alert Total
☐	icmp_broadcast_ipv4 traffic dropped	RATE	Rate_System_copp_statistics	0
☐	icmp_broadcast_ipv4 traffic passed	RATE	Rate_System_copp_statistics	0
☐	icmp_multicast_ipv6 traffic dropped	RATE	Rate_System_copp_statistics	0
☐	icmp_multicast_ipv6 traffic passed	RATE	Rate_System_copp_statistics	0
☐	icmp_security_ipv6 traffic dropped	RATE	Rate_System_copp_statistics	0
☐	icmp_security_ipv6 traffic passed	RATE	Rate_System_copp_statistics	0
☒	icmp_unicast_ipv4 traffic dropped	RATE	Rate_System_copp_statistics	0
☒	icmp_unicast_ipv4 traffic passed	RATE	Rate_System_copp_statistics	0
☐	icmp_unicast_ipv6 traffic dropped	RATE	Rate_System_copp_statistics	0
☐	icmp_unicast_ipv6 traffic passed	RATE	Rate_System_copp_statistics	0
☐	ieee_8021x traffic dropped	RATE	Rate_System_copp_statistics	0
☐	ieee_8021x traffic passed	RATE	Rate_System_copp_statistics	0

8. On PC1, run the fping to the VLAN 1 IP address of Sw1 (10.x.1.4)

```
C:\Users\student>fping 10.x.1.4 -t 1 -w 1 -n 1000
```

```
Fast pinger version 3.00
(c) Wouter Dhondt (http://www.kwakkelflap.com)
```

```
Pinging 10.x.1.4 with 32 bytes of data every 1 ms:
```

```
Reply[1] from 10.x.1.4: bytes=32 time=0.4 ms TTL=64
Reply[2] from 10.x.1.4: bytes=32 time=0.3 ms TTL=64
Reply[3] from 10.x.1.4: bytes=32 time=0.2 ms TTL=64
```

```
----- Output Omitted -----
```


```
Ping statistics for 10.x.1.4:
```

```
    Packets: Sent = 1000, Received = 774, Lost = 226 (22% loss)
```

```
Approximate round trip times in milli-seconds:
```

```
    Minimum = 0.2 ms, Maximum = 13.9 ms, Average = 0.3 ms
```

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.

```
Sw2# show dhcpv4-snooping
DHCPv4-Snooping Information

  DHCPv4-Snooping      : No      Verify MAC Address  : Yes
  Allow Overwrite Binding : No      Enabled VLANs      :

Option 82 Configurations

  Untrusted Policy      : drop      Insertion           : Yes
  Option 82 Remote-id   : mac

External Storage Information

  Volume Name           : --
  File Name              : --
  Inactive Since         : --
  Error                  : --

Authorized Server Configurations
VRF                      Authorized Servers
-----
Port Information

Port    Trust    Max    Static    Dynamic
-----  ----    -
Port    Trust    Bindings  Bindings  Bindings
-----  ----    -
```

3. On Sw2, enable DHCPv4 snooping globally and on VLAN 11

```
Sw2# conf t
```

```
Sw2(config)# dhcpv4-snooping
Sw2(config)# vlan 11
Sw2(config-vlan-11)# dhcpv4-snooping
Sw2(config-vlan-11)# end
```

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.

```
Sw2# show dhcpv4-snooping statistics
```

Packet-Type	Action	Reason	Count
server	forward	from trusted port	0
client	forward	to trusted port	0
server	drop	received on untrusted port	0
server	drop	unauthorized server	0
client	drop	destination on untrusted port	5
client	drop	untrusted option 82 field	0
client	drop	bad DHCP release request	0
client	drop	failed verify MAC check	0
client	drop	failed on max-binding limit	0

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.

```
Sw2# no debug all
Sw2# debug dhcpv4snooping all
Sw2# clear debug buffer
```

8. On Sw2, make the LAG255 uplink trusted for DHCP snooping.

```
Sw2# conf t
Sw2(config)# int lag 255
```

```
Sw2(config-lag-if)# dhcpv4-snooping trust
Sw2(config-lag-if)# end
```

9. On PC4, attempt to renew your IP. This should now succeed.

10. On Sw2, review the DHCPv4 snooping statistics.

```
Sw2# show dhcpv4-snooping statistics
```

Packet-Type	Action	Reason	Count
server	forward	from trusted port	2
client	forward	to trusted port	2
server	drop	received on untrusted port	0
server	drop	unauthorized server	0
client	drop	destination on untrusted port	9
client	drop	untrusted option 82 field	0
client	drop	bad DHCP release request	0
client	drop	failed verify MAC check	0
client	drop	failed on max-binding limit	0

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.

```
Sw2# show debug buffer
```

```
2020-10-
16:08:44:02.683891|ipsavd|LOG_DEBUG|MSTR|1|DHCPV4SNOOPING|DHCPV4SNOOPING_PACKET|DHCPV4_DISCOVER
packet received on port:1/1/4 vlan:11
----- Output Omitted -----
2020-10-
16:08:44:02.688044|ipsavd|LOG_DEBUG|MSTR|1|DHCPV4SNOOPING|DHCPV4SNOOPING_PACKET|DHCPV4_OFFER
packet received on port:lag255 len:322 vlan:11
----- Output Omitted -----
2020-10-
16:08:44:02.689629|ipsavd|LOG_DEBUG|MSTR|1|DHCPV4SNOOPING|DHCPV4SNOOPING_PACKET|DHCPV4_CLIENT
REQUEST packet received on port:1/1/4 vlan:11 from client:00:50:56:b1:95:30 to
server:10.252.1.21
----- Output Omitted -----
2020-10-
16:08:44:02.692610|ipsavd|LOG_DEBUG|MSTR|1|DHCPV4SNOOPING|DHCPV4SNOOPING_PACKET|DHCPV4_ACK
packet received on port:lag255 vlan:11 from server:10.252.1.21 to client:00:50:56:b1:95:30
assigned ip:10.x.11.104, lease time:432000
----- Output Omitted -----
```

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.

```
Sw2(config)# dhcpv4-snooping authorized-server 10.252.1.22
```

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.

```
Sw2# show dhcpv4-snooping statistics
```

Packet-Type	Action	Reason	Count
server	forward	from trusted port	2
client	forward	to trusted port	9
server	drop	received on untrusted port	0
server	drop	unauthorized server	6
client	drop	destination on untrusted port	9
client	drop	untrusted option 82 field	0
client	drop	bad DHCP release request	0
client	drop	failed verify MAC check	0
client	drop	failed on max-binding limit	0

15. On Sw2, review the event log. A WARN packet should be observed.

```
Sw2# show events -r -n 10
```

```
Event logs from current boot
```

```
2020-10-16T08:52:50.913690-04:00 TX-Sw2 ipsavd[3222]: Event|8203|LOG_WARN|MSTR|1|Packet
received from unauthorized server 10.252.1.21 on port lag255.
```

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'

```
Sw2# show debug buffer
```

```
2020-10-16:08:51:18.787181|ipsavd|LOG_DEBUG|MSTR|1|DHCPV4SNOOPING|DHCPV4SNOOPING_PACKET|DHCPV4_BOOTP_
REPLY packet was dropped at port:lag255.Reason - Reply packet came from unauthorized
server:10.252.1.21 (Server's IPv4 addr inside DHCP pkt: 10.252.1.21)
```

17. On Sw2, add the correct DHCP server with IP 10.252.1.21.

```
Sw2(config)# dhcpv4-snooping authorized-server 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

```
Sw2# show dhcpv4-snooping binding
```

MacAddress	IP	VLAN	Interface	Time-Left
00:50:56:b1:95:30	10.x.11.104	11	1/1/4	431983

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

```
Sw2# no debug all
```

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.

```
C:\Users\student>ipconfig

Windows IP Configuration

Ethernet adapter Lab NIC 6300-B Port 4:

    Connection-specific DNS Suffix  . : arubatraining.com
    IPv4 Address. . . . . : 10.x.11.104
    Subnet Mask . . . . . : 255.255.255.0
    Default Gateway . . . . . : 10.x.11.1

Ethernet adapter Do NOT Touch!:

    Connection-specific DNS Suffix  . :
    IPv4 Address. . . . . : 172.16.15.82
    Subnet Mask . . . . . : 255.255.255.0
    Default Gateway . . . . . :
```

2. On PC4, verify that you can ping to the default gateway on 10.x.11.1.

```
C:\Users\student>ping 10.x.11.1 -n 2
```

```
Pinging 10.x.11.1 with 32 bytes of data:
Reply from 10.x.11.1: bytes=32 time<1ms TTL=64
Reply from 10.x.11.1: bytes=32 time<1ms TTL=64
```

```
Ping statistics for 10.x.11.1:
    Packets: Sent = 2, Received = 2, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 0ms, Maximum = 0ms, Average = 0ms
```

3. On PC4, review the arp table, check the entry for 10.x.11.1

```
C:\Users\student>arp -a
```

Internet Address	Physical Address	Type
10.x.11.1	02-00-00-00-00-01	dynamic

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.

```
Sw2# show dhcpv4-snooping binding
```

MacAddress	IP	VLAN	Interface	Time-Left
00:50:56:b1:95:30	10.x.11.104	11	1/1/4	431073

5. On Sw2, enable ARP Inspect on VLAN 11.

```
Sw2(config)# vlan 11
Sw2(config-vlan-11)# arp inspection
Sw2(config-vlan-11)# end
```

6. On Sw2, enable arp inspection debugging

```
Sw2# debug arpsecurity all
```

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.

```
C:\WINDOWS\system32>arp -d *
C:\WINDOWS\system32>ping 10.x.11.1 -n 2

Pinging 10.x.11.1 with 32 bytes of data:
Request timed out.
Request timed out.

Ping statistics for 10.x.11.1:
    Packets: Sent = 2, Received = 0, Lost = 2 (100% loss),
```

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.

```
C:\WINDOWS\system32>arp -a
Interface: 172.16.15.82 --- 0x8
    Internet Address      Physical Address      Type
    172.16.15.196         00-50-56-b1-e2-23     dynamic
    224.0.0.22            01-00-5e-00-00-16     static
Interface: 10.x.11.104 --- 0xb
    Internet Address      Physical Address      Type
    224.0.0.22            01-00-5e-00-00-16     static
```

9. On Sw2, investigate the ARP Inspect output, review the arp inspect statistics for VLAN11.

```
Sw2# show arp inspect statistics vlan
```

VLAN	Name	Forwarded	Dropped
1	DEFAULT_VLAN_1	0	0
11	VLAN11	79	183
12	VLAN12	0	0
13	VLAN13	0	0

----- Output Omitted -----

Q: What do you observe?

A: There are several dropped entries.

10. On Sw2, review the debug buffer

```
Sw2# show debug buffer severity warn
```

```
2020-10-16:09:11:14.992884|arpsecured|LOG_WARN|MSTR|1|ARPSECURITY|ARPSECURITY_INSPECTION|DAI
Validation failed for arp pkt on vlan 11, untrusted port lag255 and ip 10.x.11.2
2020-10-16:09:11:15.304925|arpsecured|LOG_WARN|MSTR|1|ARPSECURITY|ARPSECURITY_INSPECTION|DAI
Validation failed for arp pkt on vlan 11, untrusted port lag255 and ip 10.x.11.1
2020-10-16:09:11:23.394107|arpsecured|LOG_WARN|MSTR|1|ARPSECURITY|ARPSECURITY_INSPECTION|DAI
Validation failed for arp pkt on vlan 11, untrusted port lag255 and ip 10.x.11.3
```

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.

```
Sw2# show arp inspect interface
```

Interface	Trust-State
1/1/27	Untrusted

----- Output Omitted -----

1/1/28	Untrusted
lag255	Untrusted
ubt	Untrusted

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.

```
Sw2# conf t
Sw2(config)# interface lag 255
Sw2(config-lag-if)# arp inspection trust
Sw2(config-lag-if)# end
```

13. On PC4, you should now be able to reach the network again.

```
C:\WINDOWS\system32>ping 10.x.11.1 -n 2

Pinging 10.x.11.1 with 32 bytes of data:
Reply from 10.x.11.1: bytes=32 time<1ms TTL=64
Reply from 10.x.11.1: bytes=32 time<1ms TTL=64

Ping statistics for 10.x.11.1:
    Packets: Sent = 2, Received = 2, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 0ms, Maximum = 0ms, Average = 0ms
```

14. On Sw2, review the debug buffer again to see the normal operation.

```
Sw2# show debug buffer module dhcpv4snooping

----- Output Omitted -----
2020-10-
16:08:54:42.332568|ipsavd|LOG_DEBUG|MSTR|1|DHCPV4SNOOPING|DHCPV4SNOOPING_PACKET|DHCPV4_ACK
packet received on port:lag255 vlan:11 from server:10.252.1.21 to client:00:50:56:b1:95:30
assigned ip:10.x.11.104, lease time:432000
----- Output Omitted -----
```

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.

```
Sw2# show dhcpv4-snooping binding
```

MacAddress	IP	VLAN	Interface	Time-Left
00:50:56:b1:95:30	10.x.11.104	11	1/1/4	430284

16. On Sw2, disable DHCP snooping for VLAN 11. Verify the IP Binding database is empty.

```
Sw2# conf t
Sw2(config)# vlan 11
Sw2(config-vlan-11)# no dhcpv4-snooping
Sw2(config-vlan-11)# end
```

```
Sw2# show dhcpv4-snooping binding
```

MacAddress	IP	VLAN	Interface	Time-Left
-----	-----	----	-----	-----

17. On Agg1 and on Agg2, clear the arp cache, since the PC4 entry is still present in this cache.

```
Agg1# clear arp
```

```
Agg2# clear arp
```

18. On Sw2, clear the debug buffer.

```
Sw2# clear debug buffer
```

19. On PC4, clear the arp cache and ping to the gateway.

```
C:\WINDOWS\system32>arp -d *
C:\WINDOWS\system32>ping 10.x.11.1 -n 2

Pinging 10.x.11.1 with 32 bytes of data:
Request timed out.
Request timed out.

Ping statistics for 10.x.11.1:
    Packets: Sent = 2, Received = 0, Lost = 2 (100% loss),
```

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.

```
Sw2# show debug buffer severity warn
```

```
2020-10-16:09:25:33.783288|arpsecured|LOG_WARN|MSTR|1|ARPSECURITY|ARPSECURITY_INSPECTION|DAI
Validation failed for arp pkt on vlan 11, untrusted port 1/1/4 and ip 10.x.11.104
```

This demonstrates the dependency on the IP Binding database.

21. On Sw2, enable DHCP snooping again for VLAN 11

```
Sw2# conf t
Sw2(config)# vlan 11
Sw2(config-vlan-11)# dhcpv4-snooping
Sw2(config-vlan-11)# end
```

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.

```
C:\WINDOWS\system32>ipconfig

Windows IP Configuration

Ethernet adapter Lab NIC 6300-B Port 4:

    Connection-specific DNS Suffix  . : 
    IPv4 Address. . . . . : 10.x.11.200
    Subnet Mask . . . . . : 255.255.255.0
    Default Gateway . . . . . : 10.x.11.1
```

24. On PC4, attempt to ping 10.x.11.1. This should fail.

```
C:\WINDOWS\system32>ping 10.x.11.1 -n 2

Pinging 10.x.11.1 with 32 bytes of data:
Reply from 10.x.11.200: Destination host unreachable.
Reply from 10.x.11.200: Destination host unreachable.

Ping statistics for 10.x.11.1:
    Packets: Sent = 2, Received = 0 (0% loss),
```

25. On Sw2, review the debug buffer.

```
Sw2# show debug buffer sev warn
2020-10-16:09:32:20.771677|arpsecured|LOG_WARN|MSTR|1|ARPSECURITY|ARPSECURITY_INSPECTION|DAI
Validation failed for arp pkt on vlan 11, untrusted port 1/1/4 and ip 10.x.11.200
```

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.

```
Sw2# conf t
Sw2(config)# interface 1/1/4
Sw2(config-if)# arp inspection trust
```

28. On PC4, you should now be able to access the network.

```
C:\WINDOWS\system32>ping 10.x.11.1 -n 2

Pinging 10.x.11.1 with 32 bytes of data:
Reply from 10.x.11.1: bytes=32 time<1ms TTL=64
Reply from 10.x.11.1: bytes=32 time<1ms TTL=64

Ping statistics for 10.x.11.1:
    Packets: Sent = 2, Received = 2, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 0ms, Maximum = 0ms, Average = 0ms
```

29. Cleanup: Disable debugging, make port 1/1/4 untrusted again.

```
Sw2(config-if)# no debug all
Sw2(config)# interface 1/1/4
Sw2(config-if)# no arp inspection trust
Sw2(config-if)# end
```

30. On PC4, make the Lab NIC a DHCP client again and verify you can ping the gateway.

```
C:\WINDOWS\system32>ipconfig
Windows IP Configuration
Ethernet adapter Lab NIC 6300-B Port 4:

    Connection-specific DNS Suffix  . : arubatraining.com
    IPv4 Address. . . . . : 10.x.11.104
    Subnet Mask . . . . . : 255.255.255.0
    Default Gateway . . . . . : 10.x.11.1

----- Output Omitted -----

C:\WINDOWS\system32>ping 10.x.11.1 -n 2

Pinging 10.x.11.1 with 32 bytes of data:
```

```
Reply from 10.x.11.1: bytes=32 time<1ms TTL=64  
Reply from 10.x.11.1: bytes=32 time<1ms TTL=64
```

```
Ping statistics for 10.x.11.1:  
    Packets: Sent = 2, Received = 2, Lost = 0 (0% loss),  
Approximate round trip times in milli-seconds:  
    Minimum = 0ms, Maximum = 0ms, Average = 0ms
```

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.

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.

```
Agg1# conf t
Agg1(config)# interface loopback 0
Agg1(config-loopback-if)# ipv6 address 2001:db8:x:0::2/128
```

2. Configure SVI 1 with IPv6 address 2001:db8:x:1::2/64.

```
Agg1(config-loopback-if)# interface vlan 1
Agg1(config-if-vlan)# ipv6 address 2001:db8:x:1::2/64
```

3. Review your IPv6 interfaces.

```
Agg1# show ipv6 interface brief | include up
1/1/7
loopback0          fe80::b8d4:e7b0:41:3800/64      up/ Admin state is up
vlan1              fe80::b8d4:e780:141:3800/64      up/ Admin state is up
vlan11
vlan12
vlan41              up/ Admin state is up
```

4. Review the IPv6 interface VLAN1.

```
Agg1# show ipv6 interface vlan1
Interface vlan1 is up
Admin state is up
IPv6 address:
  2001:db8:x:1::2/64 [VALID]
IPv6 link-local address: fe80::b8d4:e780:141:3800/64 [VALID]
IPv6 virtual address configured: none
IPv6 multicast routing: disable
IPv6 Forwarding feature: enabled
IPv6 multicast groups locally joined:
  ff02::1 ff02::1:ff00:2 ff02::1:ff00:0 ff02::1:ff41:3800
```

```
ff02::2
IPv6 multicast (S,G) entries joined: none
IPv6 MTU 1500
IPv6 unicast reverse path forwarding: none
IPv6 load sharing: none
L3 Counters: Rx Enabled, Tx Enabled
```

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.

```
Agg2# conf t
Agg2(config)# interface loopback 0
Agg2(config-loopback-if)# ipv6 address 2001:db8:x:0::3/128
```

6. On Agg2, configure SVI 1, use 2001:db8:x:1::3/64 as the host IP.

```
Agg2(config-loopback-if)# interface vlan 1
Agg2(config-if-vlan)# ipv6 address 2001:db8:x:1::3/64
```

7. On Agg1, verify connectivity using ping6 to Agg2 SVI1 IPv6 address. This should be successful.

```
Agg1# ping6 2001:db8:x:1::3 repetitions 2
PING 2001:db8:x:1::3(2001:db8:x:1::3) 100 data bytes
108 bytes from 2001:db8:x:1::3: icmp_seq=1 ttl=64 time=0.173 ms
108 bytes from 2001:db8:x:1::3: icmp_seq=2 ttl=64 time=0.223 ms

--- 2001:db8:x:1::3 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1051ms
rtt min/avg/max/mdev = 0.173/0.198/0.223/0.025 ms
```

8. On Agg1, review the IPv6 neighbor list. This should reveal the Layer2 address of the neighbor.

```
Agg1# show ipv6 neighbors
```

IPv6 Address State	MAC	Port	Physical Port
2001:db8:x:1::3 reachable	b8:d4:e7:41:68:00	vlan1	lag256

Total Number Of IPv6 Neighbors Entries Listed- 1.

9. On Agg1, clear the neighbor cache.

```
Agg1# clear arp
```

NOTE: This command applies to both IPv4 and IPv6 neighbor entries.

10. On Agg1, verify the neighbor list is empty.

```
Agg1# show ipv6 neighbors
IPv6 Address          MAC          Port          Physical Port
State
-----
Total Number Of IPv6 Neighbors Entries Listed- 0.
```

11. On Agg1, repeat the ping and verify it was learned in the IPv6 neighbor list.

```
Agg1# ping6 2001:db8:x:1::3 repetitions 2
PING 2001:db8:x:1::3(2001:db8:x:1::3) 100 data bytes
108 bytes from 2001:db8:x:1::3: icmp_seq=1 ttl=64 time=20.4 ms
108 bytes from 2001:db8:x:1::3: icmp_seq=2 ttl=64 time=0.292 ms

--- 2001:db8:x:1::3 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1001ms
rtt min/avg/max/mdev = 0.292/10.365/20.439/10.074 ms
```

```
Agg1# show ipv6 neighbors
IPv6 Address          MAC          Port          Physical Port
State
-----
2001:db8:x:1::3      b8:d4:e7:41:68:00  vlan1        lag256
reachable
Total Number Of IPv6 Neighbors Entries Listed- 1.
```

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.

```
Agg1# conf t
Agg1(config)# router ospfv3 1
Agg1(config-ospfv3-1)# router-id 10.x.0.2
Agg1(config-ospfv3-1)# area 0.0.0.0
```

13. On Agg1, enable Loopback0 and SVI1 for OSPFv3.

```
Agg1(config-ospfv3-1)# interface loopback 0
Agg1(config-loopback-if)# ipv6 ospfv3 1 area 0
Agg1(config-loopback-if)# interface vlan 1
Agg1(config-if-vlan)# ipv6 ospfv3 1 area 0
```

14. On Agg2, configure OSPFv3 with area 0. Assign the router-id to the IPv4 loopback0 address.

```
Agg2# conf t
Agg2(config)# router ospfv3 1
Agg2(config-ospfv3-1)# router-id 10.x.0.3
Agg2(config-ospfv3-1)# area 0.0.0.0
```

15. On Agg2, enable Loopback0 and SVI1 for OSPFv3.

```
Agg2(config-ospfv3-1)# interface loopback 0
Agg2(config-loopback-if)# ipv6 ospfv3 1 area 0
Agg2(config-loopback-if)# interface vlan 1
Agg2(config-if-vlan)# ipv6 ospfv3 1 area 0
```

16. On Agg1, review the OSPFv3 neighbors.

```
Agg1# show ipv6 ospfv3 neighbors

OSPFv3 Process ID 1 VRF default
=====

Total Number of Neighbors: 1

Neighbor ID      Priority   State          Interface
-----
10.x.0.3         1         FULL/BDR       vlan1
Neighbor address fe80::b8d4:e780:141:6800
```

Q: What is the IPv6 neighbor address ? Does it match the configure 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.

```
Agg1# show ipv6 ospfv3 1 routes

Codes: i - Intra-area route, I - Inter-area route
       E1 - External type-1, E2 - External type-2

OSPFv3 Process ID 1 VRF default, Routing Table
-----
Total Number of OSPFv3 Routes : 2

2001:db8:x::3/128 (i) area:0.0.0.0
    via fe80::b8d4:e780:141:6800 interface vlan1, cost 100 distance 110
2001:db8:x:1::/64 (i) area:0.0.0.0
    directly attached to interface vlan1, cost 100 distance 110
```

18. On Agg1, review the IPv6 routing table.

```
Agg1# show ipv6 route

Displaying ipv6 routes selected for forwarding
'[x/y]' denotes [distance/metric]
```

```
2001:db8:x::2/128, vrf default
    via loopback0, [0/0], local
2001:db8:x::3/128, vrf default
    via fe80::b8d4:e780:141:6800%vlan1, [110/100], ospf
2001:db8:x:1::/64, vrf default
    via vlan1, [0/0], connected
2001:db8:x:1::2/128, vrf default
    via vlan1, [0/0], local
```

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.

```
Agg1# show event -r -d hpe-routing
2020-10-16T11:33:36.001801-04:00 TX-Agg1 hpe-routing[8418]:
Event|4904|LOG_INFO|AMM|1/1|AdjChg: Nbr rtr ID 10.x.0.3 on interface IP addr
FE80::B8D4:E780:141:3800 on vlan1( area ID 0.0.0.0): Loading -> Full
2020-10-16T11:33:36.00319-04:00 TX-Agg1 hpe-routing[8418]:
Event|4904|LOG_INFO|AMM|1/1|AdjChg: Nbr rtr ID 10.x.0.3 on interface IP addr
FE80::B8D4:E780:141:3800 on vlan1( area ID 0.0.0.0): Exchange -> Loading
2020-10-16T11:33:35.998907-04:00 TX-Agg1 hpe-routing[8418]:
Event|4904|LOG_INFO|AMM|1/1|AdjChg: Nbr rtr ID 10.x.0.3 on interface IP addr
FE80::B8D4:E780:141:3800 on vlan1( area ID 0.0.0.0): Exstart -> Exchange
2020-10-16T11:33:35.998388-04:00 TX-Agg1 hpe-routing[8418]:
Event|4904|LOG_INFO|AMM|1/1|AdjChg: Nbr rtr ID 10.x.0.3 on interface IP addr
FE80::B8D4:E780:141:3800 on vlan1( area ID 0.0.0.0): Two-way -> Exstart
2020-10-16T11:33:35.998119-04:00 TX-Agg1 hpe-routing[8418]:
Event|4904|LOG_INFO|AMM|1/1|AdjChg: Nbr rtr ID 10.x.0.3 on interface IP addr
FE80::B8D4:E780:141:3800 on vlan1( area ID 0.0.0.0): Init -> Two-way
2020-10-16T11:33:28.971513-04:00 TX-Agg1 hpe-routing[8418]:
Event|4904|LOG_INFO|AMM|1/1|AdjChg: Nbr rtr ID 10.x.0.3 on interface IP addr
FE80::B8D4:E780:141:3800 on vlan1( area ID 0.0.0.0): Down -> Init
```

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.

```
Sw1# conf t
Sw1(config)# interface loopback 0
Sw1(config-loopback-if)# ipv6 address 2001:db8:x:0::4/128
```

21. On Sw1, enable SVI 41 and SVI 42 for IPv6.

```
Sw1(config-loopback-if)# interface vlan 41
Sw1(config-if-vlan)# ipv6 address link-local
```

```
Sw1(config-if-vlan)# interface vlan 42
Sw1(config-if-vlan)# ipv6 address link-local
```

22. On Sw1, configure OSPFv3 with area 0. Assign the router-id to the IPv4 loopback0 address.

```
Sw1(config-if-vlan)# router ospfv3 1
Sw1(config-ospfv3-1)# router-id 10.x.0.4
Sw1(config-ospfv3-1)# area 0.0.0.0
```

23. On Sw1, enable Loopback0 for OSPFv3.

```
Sw1(config-ospfv3-1)# interface loopback 0
Sw1(config-loopback-if)# ipv6 ospfv3 1 area 0
```

24. On Sw1, enable SVI 41 and SVI 42 for OSPFv3 and set the network type to point to point.

```
Sw1(config-loopback-if)# interface vlan 41
Sw1(config-if-vlan)# ipv6 ospfv3 1 area 0
Sw1(config-if-vlan)# ipv6 ospfv3 network point-to-point
Sw1(config-if-vlan)# interface vlan 42
Sw1(config-if-vlan)# ipv6 ospfv3 1 area 0
Sw1(config-if-vlan)# ipv6 ospfv3 network point-to-point
Sw1(config-if-vlan)# end
```

25. On Sw1, review the OSPFv3 interfaces.

```
Sw1# show ipv6 ospfv3 interface brief
```

```
OSPFv3 Process ID 1 VRF default
=====
```

```
Total Number of Interfaces: 3
```

Interface	Area	Cost	State	Status
loopback0	0.0.0.0	0	Loopback	up
vlan41	0.0.0.0	100	Point-to-point	up
vlan42	0.0.0.0	100	Point-to-point	up

26. On Agg1, configure the link to Sw1: enable SVI 41 for IPv6 and OSPFv3.

```
Agg1# conf t
Agg1(config)# interface vlan 41
Agg1(config-if-vlan)# ipv6 address link-local
Agg1(config-if-vlan)# ipv6 ospfv3 1 area 0
Agg1(config-if-vlan)# ipv6 ospfv3 network point-to-point
```

27. On Agg2, configure the link to Sw1: enable SVI 42 for IPv6 and OSPFv3.

```
Agg2(config)# interface vlan 42
Agg2(config-if-vlan)# ipv6 address link-local
Agg2(config-if-vlan)# ipv6 ospfv3 1 area 0
Agg2(config-if-vlan)# ipv6 ospfv3 network point-to-point
```

28. On Sw1, review the OSPFv3 neighbors, both Agg1 and Agg2 should be listed.


```
Sw1# show ipv6 ospfv3 neighbors
```

```
OSPFv3 Process ID 1 VRF default
```

```
=====
```

```
Total Number of Neighbors: 2
```

Neighbor ID	Priority	State	Interface
10.x.0.2	n/a	FULL	vlan41
Neighbor address fe80::b8d4:e780:2941:3800			
10.x.0.3	n/a	FULL	vlan42
Neighbor address fe80::b8d4:e780:2a41:6800			

29. On Sw1, review the IPv6 routing table, you should see the Loopback IPv6 addresses of Agg1 and Agg2.

```
Sw1# show ipv6 route
```

```
Displaying ipv6 routes selected for forwarding
```

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

```
2001:db8:x::2/128, vrf default
    via fe80::b8d4:e780:2941:3800%vlan41, [110/100], ospf
2001:db8:x::3/128, vrf default
    via fe80::b8d4:e780:2a41:6800%vlan42, [110/100], ospf
2001:db8:x::4/128, vrf default
    via loopback0, [0/0], local
2001:db8:x:1::/64, vrf default
    via fe80::b8d4:e780:2941:3800%vlan41, [110/200], ospf
    via fe80::b8d4:e780:2a41:6800%vlan42, [110/200], ospf
```

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.

```
Sw1# ping6 2001:db8:x::2 source 2001:db8:x::4 repetitions 2
```

```
PING 2001:db8:x::2(2001:db8:x::2) from 2001:db8:x::4 : 100 data bytes
```

```
108 bytes from 2001:db8:x::2: icmp_seq=1 ttl=64 time=0.231 ms
```

```
108 bytes from 2001:db8:x::2: icmp_seq=2 ttl=64 time=0.212 ms
```

```
--- 2001:db8:x::2 ping statistics ---
```

```
2 packets transmitted, 2 received, 0% packet loss, time 1020ms
```

```
rtt min/avg/max/mdev = 0.212/0.221/0.231/0.017 ms
```

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.

```
Agg1(config)# interface vlan 11
Agg1(config-if-vlan)# ipv6 address 2001:db8:x:11::2/64
```

2. On Agg1 SVI11, enable OSPFv3 area 0 and set it as OSPFv3 passive interface.

```
Agg1(config-if-vlan)# ipv6 ospfv3 1 area 0
Agg1(config-if-vlan)# ipv6 ospfv3 passive
```

3. On Agg1 SVI11, configure VSX Active Gateway for SVI 11.

```
Agg1(config-if-vlan)# active-gateway ipv6 2001:db8:x:11::1 mac 02:00:00:00:00:01
```

4. On Agg1 SVI11, enable the IPv6 ND Router Advertisements (no ipv6 nd suppress-ra)

```
Agg1(config-if-vlan)# no ipv6 nd suppress-ra
```

5. On Agg2, configure SVI 11 with IPv6 address 2001:db8:x:11::3/64.


```
Agg2(config)# interface vlan 11
Agg2(config-if-vlan)# ipv6 address 2001:db8:x:11::3/64
```

6. On Agg2, complete the SVI 11 configuration: OSPFv3 and RA.

```
Agg2(config-if-vlan)# ipv6 ospfv3 1 area 0
Agg2(config-if-vlan)# ipv6 ospfv3 passive
Agg2(config-if-vlan)# no ipv6 nd suppress-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.



```
C:\Users\student>ipconfig /all
```

```
----- Output Omitted -----
```

```
Ethernet adapter Lab NIC - 6300-A Port 3:
```

```

Connection-specific DNS Suffix  . : arubatraining.com
Description . . . . . : Intel(R) 82574L Gigabit Network Connection #3
Physical Address. . . . . : 00-50-56-B1-04-25
DHCP Enabled. . . . . : Yes
Autoconfiguration Enabled . . . . : Yes
IPv6 Address. . . . . : 2001:db8:x:11:603c:2ce1:331c:8358(Preferred)
Temporary IPv6 Address. . . . . : 2001:db8:x:11:2c1c:4164:31ff:9ca7(Preferred)
Link-local IPv6 Address . . . . . : fe80::603c:2ce1:331c:8358%15(Preferred)
IPv4 Address. . . . . : 10.x.11.103(Preferred)
Subnet Mask . . . . . : 255.255.255.0
Lease Obtained. . . . . : Friday, October 16, 2020 11:50:41 AM
Lease Expires . . . . . : Wednesday, October 21, 2020 11:54:39 AM
Default Gateway . . . . . : fe80::b8d4:e780:b41:6800%15
                             fe80::b8d4:e780:b41:3800%15
                             10.x.11.1
DHCP Server . . . . . : 10.252.1.21
DHCPv6 IAID . . . . . : 553668694
DHCPv6 Client DUID. . . . . : 00-01-00-01-27-1B-7A-97-00-50-56-B1-F6-6B
DNS Servers . . . . . : 10.252.1.21
NetBIOS over Tcpip. . . . . : Enabled

```

```
----- Output Omitted -----
```

8. On the PC3, verify IPv6 connectivity with the Agg1 and Agg2 IPv6 loopback addresses using ping -6.

```
C:\Users\student> ping -6 2001:db8:x::2 -n 2
```

```
Pinging 2001:db8:x::2 with 32 bytes of data:
Reply from 2001:db8:x::2: time=1ms
Reply from 2001:db8:x::2: time<1ms
```

```
Ping statistics for 2001:db8:x::2:
    Packets: Sent = 2, Received = 2, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 0ms, Maximum = 1ms, Average = 0ms
```

```
C:\Users\student> ping -6 2001:db8:x::3 -n 2
```

```
Pinging 2001:db8:x::3 with 32 bytes of data:
Reply from 2001:db8:x::3: time<1ms
Reply from 2001:db8:x::3: time<1ms
```

```
Ping statistics for 2001:db8:x::3:
    Packets: Sent = 2, Received = 2, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 0ms, Maximum = 0ms, Average = 0ms
```

9. On PC3, use tracert -d -6 to the Sw1 IPv6 Loopback (2001:db8:x::4)

```
C:\Users\student> tracert -d -6 2001:db8:x::4
```

```
Tracing route to 2001:db8:x::4 over a maximum of 30 hops
```

```
  1    <1 ms    <1 ms    <1 ms    2001:db8:x:11::3
  2    <1 ms    <1 ms    <1 ms    2001:db8:x::4
```

10. On Agg1, review the IPv6 neighbors. You should observe the PC3 in VLAN11 on LAG1.

```
Agg1# show ipv6 neighbors
```

IPv6 Address State	MAC	Port	Physical Port

fe80::883a:3080:29a5:5280 reachable	88:3a:30:a5:52:80	vlan41	1/1/1
fe80::b8d4:e780:141:6800 reachable	b8:d4:e7:41:68:00	vlan1	lag256
2001:db8:x:11:2c1c:4164:31ff:9ca7 reachable	00:50:56:b1:04:25	vlan11	lag2
fe80::603c:2ce1:331c:8358 reachable	00:50:56:b1:04:25	vlan11	lag2

2001:db8:x:1::3
reachable

b8:d4:e7:41:68:00 vlan1

lag256

Total Number Of IPv6 Neighbors Entries Listed- 5.

Aruba Training-Confidential

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

```
Sw2(config)# interface 1/1/4
Sw2(config-if)# vlan access 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.

```
Sw2# show mac-address port 1/1/4
MAC age-time           : 300 seconds
```

Number of MAC addresses : 2

MAC Address	VLAN	Type	Port
00:50:56:b1:95:30	12	port-access-security	1/1/4

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

```
Sw1# conf t
Sw1(config)# dhcpv6-server vrf default
Sw1(config-dhcpv6-server)# authoritative
Sw1(config-dhcpv6-server)# pool vlan12
```

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.

```
Sw1(config-dhcpv6-server-pool)# range 2001:db8:x:12::100 2001:db8:x:12::200 prefix-len 64
Sw1(config-dhcpv6-server-pool)# dns-server 2001:db8:252:1::21
```

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.

```
Sw1(config-dhcpv6-server-pool)# exit
Sw1(config-dhcpv6-server)# enable
```

18. On Sw1, review the settings.

```
Sw1# show dhcpv6-server

VRF Name          : default
DHCPv6 Server     : enabled
Operational State : operational
Authoritative Mode : true

Pool Name         : vlan12
Lease Duration    : 00:01:00

DHCPv6 dynamic IP allocation
-----
Start-IPv6-Address      End-IPv6-Address      Prefix-Length
```

```

-----
2001:db8:x:12::100          2001:db8:x:12::200          64
DHCPv6 Server options
-----
*****
Note:DHCPv6 Server well known option number mapping:
dns-server(23), vendor-class(16), vendor-data(17)
*****
Option-Number      Option-Type      Option-Value
-----
23                 ipv6             2001:db8:252:1::21

DHCPv6 Server static IP allocation
-----
DHCPv6 Server static host is not configured.

```

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

```
Agg1(config)# dhcpv6-relay
```

20. On Agg1, set the IPv6 SVI address, and enable OSPFv3

```

Agg1(config)# interface vlan 12
Agg1(config-if-vlan)# ipv6 address 2001:db8:x:12::2/64
Agg1(config-if-vlan)# ipv6 ospfv3 1 area 0
Agg1(config-if-vlan)# ipv6 ospfv3 passive

```

21. On Agg1 SVI 12, enable DHCP relay. The destination address is the Sw1 Loopback IP address.

```
Agg1 (config-if-vlan)# ipv6 helper-address unicast 2001:db8:x::4
```

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.

```

Agg1(config-if-vlan)# no ipv6 nd suppress-ra
Agg1(config-if-vlan)# ipv6 nd ra other

```

23. Review the DHCPv6 status.

```

Agg1# show dhcpv6-relay
DHCPV6 Relay Agent : enabled
Option 79           : disabled

```

24. On Agg2, repeat this configuration

```
Agg2(config)# dhcpv6-relay
```



```
Agg2(config)# interface vlan 12
Agg2(config-if-vlan)# ipv6 address 2001:db8:x:12::3/64
Agg2(config-if-vlan)# ipv6 ospfv3 1 area 0
Agg2(config-if-vlan)# ipv6 ospfv3 passive
Agg2(config-if-vlan)# no ipv6 nd suppress-ra
Agg2(config-if-vlan)# ipv6 nd ra other
Agg2(config-if-vlan)# ipv6 helper-address unicast 2001:db8:x::4
```

Verify Stateless DHCPv6

25. On Agg1, enable dhcp-relay debugging and clear the buffer.

```
Agg1# no debug all
Agg1# debug dhcpv6relay all
Agg1# clear debug 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.

```
C:\WINDOWS\system32> ipconfig /all
----- Output Omitted -----
Ethernet adapter Lab NIC 6300-B Port 4:

    Connection-specific DNS Suffix  . : arubatraining.com
    Description . . . . . : Intel(R) 82574L Gigabit Network Connection #2
    Physical Address. . . . . : 00-50-56-B1-95-30
    DHCP Enabled. . . . . : Yes
    Autoconfiguration Enabled . . . . : Yes
    IPv6 Address. . . . . : 2001:db8:x:12:45a5:1d00:9c5d:db4(Preferred)
    Temporary IPv6 Address. . . . . : 2001:db8:x:12:a85b:a099:cb0c:318d(Preferred)
    Link-local IPv6 Address . . . . . : fe80::45a5:1d00:9c5d:db4%11(Preferred)
    IPv4 Address. . . . . : 10.x.12.104(Preferred)
    Subnet Mask . . . . . : 255.255.255.0
    Lease Obtained. . . . . : Friday, October 16, 2020 1:45:05 PM
    Lease Expires . . . . . : Wednesday, October 21, 2020 1:45:06 PM
    Default Gateway . . . . . : fe80::b8d4:e780:c41:6800%11
                                fe80::b8d4:e780:c41:3800%11
                                10.x.12.1
    DHCP Server . . . . . : 10.252.1.21
    DHCPv6 IAID . . . . . : 453005398
    DHCPv6 Client DUID. . . . . : 00-01-00-01-27-1B-8F-D3-00-50-56-B1-D8-00
    DNS Servers . . . . . : 2001:db8:252:1::21
                                10.252.1.21
    NetBIOS over Tcpip. . . . . : Enabled
----- Output Omitted -----
```

28. On Agg1, review the debug buffer. You should see the DHCPv6 relay flow.

```
Agg1# show debug buffer
```

```
2020-10-16:13:45:43.142461|hpe-relay|LOG_DEBUG|AMM|-|DHCPV6RELAY|DHCPV6RELAY|Packet received
on interface vlan12 from client fe80::45a5:1d00:9c5d:db4.
2020-10-16:13:45:43.142536|hpe-relay|LOG_DEBUG|AMM|-|DHCPV6RELAY|DHCPV6RELAY|Packet sent to
server 2001:db8:x::4 from client fe80::45a5:1d00:9c5d:db4 successfully.
2020-10-16:13:45:43.157088|hpe-relay|LOG_DEBUG|AMM|-|DHCPV6RELAY|DHCPV6RELAY|Packet received
on interface vlan12 from Server 2001:db8:x::4.
2020-10-16:13:45:43.157128|hpe-relay|LOG_DEBUG|AMM|-|DHCPV6RELAY|DHCPV6RELAY|Packet received
on interface vlan12 from Server 2001:db8:x::4 to client fe80::45a5:1d00:9c5d:db4.
2020-10-16:13:45:43.157194|hpe-relay|LOG_DEBUG|AMM|-|DHCPV6RELAY|DHCPV6RELAY|Packet sent to
client fe80::45a5:1d00:9c5d:db4 from server 2001:db8:x::4 successfully via interface vlan12.
2020-10-16:13:45:45.545293|ndmd|LOG_ERR|AMM|-
|NDM|NDM_NBRTABLE|[nd_nbr_mgr_process_icmp6_rcv_na_event(2038)] [Port vlan12] Nbr not found
for 2001:db8:x:12:45a5:1d00:9c5d:db4
2020-10-16:13:45:46.194674|hpe-relay|LOG_DEBUG|AMM|-|DHCPV6RELAY|DHCPV6RELAY|Packet received
on interface vlan12 from client fe80::45a5:1d00:9c5d:db4.
2020-10-16:13:45:46.194740|hpe-relay|LOG_DEBUG|AMM|-|DHCPV6RELAY|DHCPV6RELAY|Packet sent to
server 2001:db8:x::4 from client fe80::45a5:1d00:9c5d:db4 successfully.
2020-10-16:13:45:46.195281|hpe-relay|LOG_DEBUG|AMM|-|DHCPV6RELAY|DHCPV6RELAY|Packet received
on interface vlan12 from Server 2001:db8:x::4.
2020-10-16:13:45:46.195319|hpe-relay|LOG_DEBUG|AMM|-|DHCPV6RELAY|DHCPV6RELAY|Packet received
on interface vlan12 from Server 2001:db8:x::4 to client fe80::45a5:1d00:9c5d:db4.
2020-10-16:13:45:46.195384|hpe-relay|LOG_DEBUG|AMM|-|DHCPV6RELAY|DHCPV6RELAY|Packet sent to
client fe80::45a5:1d00:9c5d:db4 from server 2001:db8:x::4 successfully via interface vlan12.
2020-10-16:13:45:46.195414|hpe-relay|LOG_DEBUG|AMM|-|DHCPV6RELAY|DHCPV6RELAY|Packet received
on interface vlan12 from Server 2001:db8:x::4.
2020-10-16:13:45:46.195441|hpe-relay|LOG_DEBUG|AMM|-|DHCPV6RELAY|DHCPV6RELAY|Packet received
on interface vlan12 from Server 2001:db8:x::4 to client fe80::45a5:1d00:9c5d:db4.
2020-10-16:13:45:46.195488|hpe-relay|LOG_DEBUG|AMM|-|DHCPV6RELAY|DHCPV6RELAY|Packet sent to
client fe80::45a5:1d00:9c5d:db4 from server 2001:db8:x::4 successfully via interface vlan12.
```

29. On Sw1, review the DHCPv6 lease database.

```
Sw1# show dhcpv6-server leases
No DHCP leases in the 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.

```
Sw1# diagnostics
Sw1# diag-dump dhcpv6-server basic
----- Output Omitted -----
Oct 16 13:41:03 TX-Sw1 ip[28486]: dnsmasq-dhcp: DHCPINFORMATION-REQUEST(vlan41)
00:01:00:01:27:1b:8f:d3:00:50:56:b1:d8:00
----- Output Omitted -----
```

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.

```
Agg1(config)# interface vlan 12
Agg1(config-if-vlan)# ipv6 nd prefix default no-advertise
```

32. On Agg1 SVI12, enable the stateful DHCP with the 'MANAGED' flag in the RA options.

```
Agg1(config-if-vlan)# ipv6 nd ra managed-config-flag
```

33. On Agg2, repeat this configuration on SVI 12.

```
Agg2(config)# int vlan 12
Agg2(config-if-vlan)# ipv6 nd prefix default no-advertise
Agg2(config-if-vlan)# ipv6 nd ra managed-config-flag
```

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.

```
C:\WINDOWS\system32>ipconfig /all

Ethernet adapter Lab NIC 6300-B Port 4:

    Connection-specific DNS Suffix  . : arubatraining.com
    Description . . . . . : Intel(R) 82574L Gigabit Network Connection #2
    Physical Address. . . . . : 00-50-56-B1-95-30
    DHCP Enabled. . . . . : Yes
    Autoconfiguration Enabled . . . . : Yes
    IPv6 Address. . . . . : 2001:db8:x:12::1e5(Preferred)
    Lease Obtained. . . . . : Friday, October 16, 2020 1:54:59 PM
    Lease Expires . . . . . : Friday, October 16, 2020 2:54:59 PM
    Link-local IPv6 Address . . . . . : fe80::45a5:1d00:9c5d:db4%11(Preferred)
    IPv4 Address. . . . . : 10.x.12.104(Preferred)
    Subnet Mask . . . . . : 255.255.255.0
    Lease Obtained. . . . . : Friday, October 16, 2020 1:54:57 PM
    Lease Expires . . . . . : Wednesday, October 21, 2020 1:54:58 PM
    Default Gateway . . . . . : fe80::b8d4:e780:c41:3800%11
```

```

                                fe80::b8d4:e780:c41:6800%11
                                10.x.12.1
DHCP Server . . . . . : 10.252.1.21
DHCPv6 IAID . . . . . : 453005398
DHCPv6 Client DUID. . . . . : 00-01-00-01-27-1B-8F-D3-00-50-56-B1-D8-00
DNS Servers . . . . . : 2001:db8:252:1::21
                                10.252.1.21
NetBIOS over Tcip. . . . . : Enabled

```

36. On Sw1, review the DHCPv6 lease database

Sw1# show dhcpv6-server leases

IPv6-Address Hostname -----	Client-Id VRF-Name -----	Expiry-Time Link-Address -----	Client-
2001:db8:x:12::1e5	*	14:55:36 16/10/2020	P15-TX-PC4
default	00:01:00:01:27:1b:8f:d3:00:50:56:b1:d8:00		

37. On Sw1, review the diag-dump of the DHCPv6 server. There should be 4 packets.

Sw1# diag-dump dhcpv6-server basic

```

Oct 16 13:55:35 TX-Sw1 ip[28486]: dnsmasq-dhcp: DHCP SOLICIT(vlan41)
00:01:00:01:27:1b:8f:d3:00:50:56:b1:d8:00
Oct 16 13:55:35 TX-Sw1 ip[28486]: dnsmasq-dhcp: DHCP ADVERTISE(vlan41) 2001:db8:x:12::1e5
00:01:00:01:27:1b:8f:d3:00:50:56:b1:d8:00 P15-TX-PC4
Oct 16 13:55:36 TX-Sw1 ip[28486]: dnsmasq-dhcp: DHCP REQUEST(vlan41)
00:01:00:01:27:1b:8f:d3:00:50:56:b1:d8:00
Oct 16 13:55:36 TX-Sw1 ip[28486]: dnsmasq-dhcp: DHCP REPLY(vlan41) 2001:db8:x:12::1e5
00:01:00:01:27:1b:8f:d3:00:50:56:b1:d8:00 P15-TX-PC4

```

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.

Agg1# show debug buffer

```

----- Output Omitted -----
2020-10-16:13:55:36.756518|hpe-relay|LOG_DEBUG|AMM|-|DHCPV6RELAY|DHCPV6RELAY|Packet received
on interface vlan12 from Server 2001:db8:x::4.
2020-10-16:13:55:36.756556|hpe-relay|LOG_DEBUG|AMM|-|DHCPV6RELAY|DHCPV6RELAY|Packet received
on interface vlan12 from Server 2001:db8:x::4 to client fe80::45a5:1d00:9c5d:db4.
2020-10-16:13:55:36.756615|hpe-relay|LOG_DEBUG|AMM|-|DHCPV6RELAY|DHCPV6RELAY|Packet sent to
client fe80::45a5:1d00:9c5d:db4 from server 2001:db8:x::4 successfully via interface vlan12.
2020-10-16:13:55:36.756631|hpe-relay|LOG_DEBUG|AMM|-|DHCPV6RELAY|DHCPV6RELAY|Packet received
on interface vlan12 from Server 2001:db8:x::4.
2020-10-16:13:55:36.756652|hpe-relay|LOG_DEBUG|AMM|-|DHCPV6RELAY|DHCPV6RELAY|Packet received
on interface vlan12 from Server 2001:db8:x::4 to client fe80::45a5:1d00:9c5d:db4.
2020-10-16:13:55:36.756695|hpe-relay|LOG_DEBUG|AMM|-|DHCPV6RELAY|DHCPV6RELAY|Packet sent to
client fe80::45a5:1d00:9c5d:db4 from server 2001:db8:x::4 successfully via interface vlan12.

```

39. On Agg1, disable the debugging.

```
Agg1# no debug all
```

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 ran 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, but 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

Spanning tree was not active on Agg2.
Sw2 has STP loop protection active on the uplink LAG.
When the uplink to Agg1 goes down, it no longer receives STP BPDUs, so it disables the port.
This is not observed on the VSX LAG under normal conditions since Agg1 is responsible for the BPDU generation.

```
# agg2
spanning-tree
```

```
T1-Sw2-lab14-tt01(config)# show interface lag255 brief
```

Port	Native	Mode	Type	Enabled	Status	Reason	Speed
Description	VLAN						(Mb/s)
lag255	1	trunk	--	yes	up	--	10000

```
T1-Sw2-lab14-tt01(config)# show lacp interfaces
```

State abbreviations :

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

Actor details of all interfaces:

Intf	Aggr Name	Port Id	Port Pri	State	System-ID	System Pri	Aggr Key	Forwarding State
1/1/25	lag255							down
1/1/26	lag255	27	1	ALFNCD	88:3a:30:a5:26:80	65534	255	up

Partner details of all interfaces:

Intf	Aggr Name	Port Id	Port Pri	State	System-ID	System Pri	Aggr Key
1/1/25	lag255						
1/1/26	lag255	1002	1	ALFNCD	02:01:00:00:00:00	65534	2

T1-Sw2-lab14-tt01(config)# **show spanning-tree**
Spanning tree status : Enabled Protocol: MSTP

MST0

Root ID Priority : 32768
MAC-Address: 88:3a:30:a5:26:80
This bridge is the root
Hello time(in seconds):2 Max Age(in seconds):20
Forward Delay(in seconds):15

Bridge ID Priority : 32768
MAC-Address: 88:3a:30:a5:26:80
Hello time(in seconds):2 Max Age(in seconds):20
Forward Delay(in seconds):15

Port BPDU-Rx	Role TCN-Tx	State TCN-Rx	Cost	Priority	Type	BPDU-Tx
1/1/1	Disabled	Blocking	20000	128	P2P	0
0	0	0				
1/1/2	Disabled	Blocking	20000	128	P2P	0
0	0	0				
1/1/3	Disabled	Blocking	20000	128	P2P	0
0	0	0				
1/1/4	Designated	Forwarding	20000	128	P2P	32257
0	0	0				
1/1/5	Disabled	Blocking	20000	128	P2P	0
0	0	0				
1/1/6	Disabled	Blocking	20000	128	P2P	0
0	0	0				
1/1/7	Disabled	Blocking	20000	128	P2P	0
0	0	0				
1/1/8	Disabled	Blocking	20000	128	P2P	0
0	0	0				
1/1/9	Disabled	Blocking	20000	128	P2P	0
0	0	0				
1/1/10	Disabled	Blocking	20000	128	P2P	0
0	0	0				

1/1/11	Disabled	Blocking	20000	128	P2P	0
0	0					
1/1/12	Disabled	Blocking	20000	128	P2P	0
0	0					
1/1/13	Disabled	Blocking	20000	128	P2P	0
0	0					
1/1/14	Disabled	Blocking	20000	128	P2P	0
0	0					
1/1/15	Disabled	Blocking	20000	128	P2P	0
0	0					
1/1/16	Disabled	Blocking	20000	128	P2P	0
0	0					
1/1/17	Disabled	Blocking	20000	128	P2P	0
0	0					
1/1/18	Disabled	Blocking	20000	128	P2P	0
0	0					
1/1/19	Disabled	Blocking	20000	128	P2P	0
0	0					
1/1/20	Disabled	Blocking	20000	128	P2P	0
0	0					
1/1/21	Disabled	Blocking	20000	128	P2P	0
0	0					
1/1/22	Disabled	Blocking	20000	128	P2P	0
0	0					
1/1/23	Disabled	Blocking	20000	128	P2P	0
0	0					
1/1/24	Disabled	Blocking	20000	128	P2P	0
0	0					
1/1/27	Designated	Forwarding	2000	128	P2P	32416
0	0					
1/1/28	Disabled	Blocking	20000	128	P2P	0
0	0					
lag255	Alternate	Loop-Inc	2000	64	P2P	5
32370	2	4				

Number of topology changes : 1
Last topology change occurred : 12096086 seconds ago

T1-Sw2-lab14-tt01(config)# **show spanning-tree inconsistent-ports**

Instance ID	Blocked Port	Reason
0	lag255	Loop Guard

T1-Agg2-lab14-tt01(config)# **show spanning-tree**
Spanning-tree is disabled

T1-Agg2-lab14-tt01(config)#

Trouble Ticket 01.02

Hints

Use PC4 with `tracert -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

```
# PC4
tracert -d 10.252.1.23
```

Aruba Training-Confidential

```
C:\Users\student>tracert -d 10.252.1.23

Tracing route to 10.252.1.23 over a maximum of 30 hops

  1  <1 ms    <1 ms    <1 ms    10.14.11.3
  2  *         <1 ms    *        10.14.11.2
  3  ^C

C:\Users\student>tracert -d 10.252.1.23

Tracing route to 10.252.1.23 over a maximum of 30 hops

  1  <1 ms    <1 ms    <1 ms    10.14.11.3
  2  *         *        <1 ms    10.14.11.2
  3  *         *        *        Request timed out.
  4  *         *        *        Request timed out.
  5  *         *        *        Request timed out.
  6  *         *        *        Request timed out.
  7  *         4 ms    *        10.252.1.23
  8  *         *        *        Request timed out.
  9  _

10.252.1.23: request timed out
Reply[20401] from 10.252.1.23: bytes=32 time=2.8 ms TTL=60
10.252.1.23: request timed out
10.252.1.23: request timed out
10.252.1.23: request timed out
Reply[20405] from 10.252.1.23: bytes=32 time=2.4 ms TTL=60
10.252.1.23: request timed out
Reply[20407] from 10.252.1.23: bytes=32 time=3.9 ms TTL=60
10.252.1.23: request timed out
10.252.1.23: request timed out
Reply[20410] from 10.252.1.23: bytes=32 time=2.4 ms TTL=60
10.252.1.23: request timed out
10.252.1.23: request timed out
Reply[20413] from 10.252.1.23: bytes=32 time=4.5 ms TTL=60
10.252.1.23: request timed out
10.252.1.23: request timed out
Reply[20416] from 10.252.1.23: bytes=32 time=2.4 ms TTL=60
10.252.1.23: request timed out
Reply[20418] from 10.252.1.23: bytes=32 time=2.5 ms TTL=60
10.252.1.23: request timed out
10.252.1.23: request timed out
10.252.1.23: request timed out
Reply[20422] from 10.252.1.23: bytes=32 time=2.8 ms TTL=60
10.252.1.23: request timed out
```

```
Agg2 seems ok, from Agg1 there are responses missing
> Force traffic over Agg1
# sw2
int 1/1/26
shutdown
```

PC4 Test trace again

```
C:\Users\student>tracert -d 10.252.1.23

Tracing route to 10.252.1.23 over a maximum of 30 hops

  1    <1 ms    <1 ms    <1 ms    10.14.11.2
  2      *      *      *      Request timed out.
  3      *      *

10.252.1.23: request timed out
Reply[24792] from 10.252.1.23: bytes=32 time=2.6 ms TTL=60
10.252.1.23: request timed out
10.252.1.23: request timed out
Reply[24795] from 10.252.1.23: bytes=32 time=2.0 ms TTL=60
10.252.1.23: request timed out
Reply[24797] from 10.252.1.23: bytes=32 time=3.7 ms TTL=60
10.252.1.23: request timed out
Reply[24799] from 10.252.1.23: bytes=32 time=2.4 ms TTL=60
10.252.1.23: request timed out
10.252.1.23: request timed out
Reply[24802] from 10.252.1.23: bytes=32 time=2.5 ms TTL=60
10.252.1.23: request timed out
Reply[24804] from 10.252.1.23: bytes=32 time=2.1 ms TTL=60
10.252.1.23: request timed out
10.252.1.23: request timed out
Reply[24807] from 10.252.1.23: bytes=32 time=2.9 ms TTL=60
10.252.1.23: request timed out
10.252.1.23: request timed out
10.252.1.23: request timed out
Reply[24811] from 10.252.1.23: bytes=32 time=2.7 ms TTL=60
10.252.1.23: request timed out
```

This confirms Agg1 responds ok, but there is no more response after Agg1
> Suspect issue on link between Agg1 and RouterA or packet does not reach RouterA

> Investigate Agg1 inbound interface (LAG2 - 1/1/2)
> Investigate Agg1 outbound interface (1/1/7)

T1-Agg1-lab13(config)# show int 1/1/2 extended

Interface 1/1/2

Statistics	Value
...	
RX Bytes	359272
RX Dropped	0
RX Filtered	2144
RX Packets	4165
TX Bytes	201231
TX Dropped	0
TX Filtered	18
TX Packets	2262

```
# Repeat this command to verify that the 'RX Filter' value increases
# Stop the fping on PC4 to verify the 'RX Filter' value is now stable
> This confirms some configuration is active that Filters(drops) packets.
```

```
> Review the port configuration
```

```
T1-Agg1-lab13(config-queue)# show run int lag2
interface lag 2 multi-chassis
  no shutdown
  description sw2
  no routing
  vlan trunk native 1
  vlan trunk allowed 1,11-12,212
  lacp mode active
  spanning-tree root-guard
  apply policy asts-qos in
exit
```

```
T1-Agg1-lab13(config-queue)# show policy asts-qos
```

Sequence	Name	Additional Policy Parameters	Comment	Class Type
				action

```
-----
      asts-qos
      VSX synchronization: Enabled
10    bulk-data ipv4
      cir kbps 10000
      cbs 2000000
      exceed drop

20    icmp    ipv4
      cir kbps 5
      cbs 0
      exceed drop
```

```
>> Invalid rate limiter applied.
```

```
>> Either remove or adjust the CIR value so the fping has no more drops.
```

```
# agg1
policy asts-qos
  20 class ip icmp action cir kbps 5000 cbs 0 exceed drop
```

Aruba Training-Confidential

Trouble Ticket 02.01

Hints

```
On Sw2
show interface brief
show lacp interfaces
```

Solution

Customer had defined LAG2 with LACP on Agg2, but as a local LAG, not a VSX LAG (no multi-chassis).

```
T1-Sw2-lab14-tt02(config-if)# show lacp interfaces
```

State abbreviations :

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

Actor details of all interfaces:

Intf	Aggr Name	Port Id	Port Pri	State	System-ID	System Pri	Aggr Key	Forwarding State
1/1/25	lag255	26	1	ALFNCD	88:3a:30:a5:26:80	65534	255	up
1/1/26	lag255	27	1	ALFO	88:3a:30:a5:26:80	65534	255	lacp-block

Partner details of all interfaces:

Intf	Aggr Name	Port Id	Port Pri	State	System-ID	System Pri	Aggr Key
1/1/25	lag255	2	1	ALFNCD	02:01:00:00:00:00	65534	2
1/1/26	lag255	3	1	ALFO	b8:d4:e7:41:b8:00	65534	2

On Agg1, review the lacp multi-chassis. Remote Actor does not have lag2 listed.

```
T1-Agg1-lab14-tt02(config-policy)# show lacp interfaces multi-chassis
```

State abbreviations :

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

Actor details of all interfaces:

Intf	Aggregate name	Port id	Port Priority	State	System-ID	System Priority	Aggr Key
1/1/2	lag2(mc)	2	1	ALFNCD	02:01:00:00:00:00	65534	2
1/1/5	lag5(mc)	5	1	ALFNCD	02:01:00:00:00:00	65534	5

Partner details of all interfaces:

Intf	Aggregate name	Partner Port-id	Port Priority	State	System-ID	System Priority	Aggr Key
1/1/2	lag2(mc)	26	1	ALFNCD	88:3a:30:a5:26:80	65534	255
1/1/5	lag5(mc)	2	255	ASFNCD	20:4c:03:b1:cd:e2	32768	2

Remote Actor details of all interfaces:

Intf	Aggregate name	Port id	Port Priority	State	System-ID	System Priority	Aggr Key
1/1/5	lag5(mc)	1005	1	ALFNCD	02:01:00:00:00:00	65534	5

Remote Partner details of all interfaces:

Intf	Aggregate name	Partner Port-id	Port Priority	State	System-ID	System Priority	Aggr Key
1/1/5	lag5(mc)	3	255	ASFNCD	20:4c:03:b1:cd:e2	32768	2

Agg1 VSX LAG, compare the LAG2 to LAG5

T1-Agg1-lab14-tt02(config-policy)# show vsx lacp aggregates lag2
Local-peer Remote-peer

```
-----
Aggregate name : lag2 (multi-chassis) lag2
Interfaces      : 1/1/2                1/1/2
Peer interfaces :
Heartbeat rate  : Slow                  Slow
Hash            : 13-src-dst            13-src-dst
Aggregate mode  : Active                 Active
```

```
T1-Agg1-lab14-tt02(config-policy)#
T1-Agg1-lab14-tt02(config-policy)#
T1-Agg1-lab14-tt02(config-policy)# show vsx lacp aggregates lag5
Local-peer Remote-peer
-----
```


Aggregate name	: lag5 (multi-chassis)	lag5 (multi-chassis)
Interfaces	: 1/1/5	1/1/5
Peer interfaces	: 1/1/5	1/1/5
Heartbeat rate	: Slow	Slow
Hash	: l3-src-dst	l3-src-dst
Aggregate mode	: Active	Active

correct: agg2: remove LAG2 and define as LAG2 multi-chassis, assign port 1/1/2

Trouble Ticket 03.01

Hints

```
# On Agg1
show interface brief
show lacp interfaces multi-chassis
```

Solution

Customer had defined LAG255 as static on Sw2 and they had the LACP fallback enabled on the VSX side to support new switches for ZTP.
On the VSX side, this is reported as 'Individual'. This causes the LAG to function even without LACP on the Partner side.
The Partner is not listed with a 'System-ID', this indicates there is no LACP on the partner side.
A VSX LAG in Individual mode only accepts the native VLAN, this causes the tagged VLAN11 traffic to be dropped.

T1-Agg1-lab14-tt03(config-policy)# **show lacp interfaces multi-chassis**

State abbreviations :

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

Actor details of all interfaces:

Intf	Aggregate name	Port id	Port Priority	State	System-ID	System Priority	Aggr Key
1/1/2	lag2(mc)	2	1	IE	02:01:00:00:00:00	65534	2
1/1/5	lag5(mc)	5	1	ALFNCD	02:01:00:00:00:00	65534	5

Partner details of all interfaces:

Intf	Aggregate name	Partner Port-id	Port Priority	State	System-ID	System Priority	Aggr Key
1/1/2	lag2(mc)	0	65534	IE	00:00:00:00:00:00	0	0
1/1/5	lag5(mc)	2	255	ASFNCD	20:4c:03:b1:cd:e2	32768	2

Remote Actor details of all interfaces:

Intf	Aggregate name	Port id	Port Priority	State	System-ID	System Priority	Aggr Key
1/1/2	lag2(mc)	1002	1	IE	02:01:00:00:00:00	65534	2

```
1/1/5      lag5(mc)    1005    1      ALFNCD  02:01:00:00:00:00 65534    5
```

Remote Partner details of all interfaces:

Intf	Aggregate name	Partner Port-id	Port Priority	State	System-ID	System Priority	Aggr Key
1/1/2	lag2(mc)	0	65534	IE	00:00:00:00:00:00	0	0
1/1/5	lag5(mc)	3	255	ASFNCD	20:4c:03:b1:cd:e2	32768	2

T1-Agg1-lab14-tt03(config-policy)# **show lacp interfaces multi-chassis vsx-peer**

State abbreviations :

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

Actor details of all interfaces:

Intf	Aggregate name	Port id	Port Priority	State	System-ID	System Priority	Aggr Key
1/1/2	lag2(mc)	1002	1	IE	02:01:00:00:00:00	65534	2
1/1/5	lag5(mc)	1005	1	ALFNCD	02:01:00:00:00:00	65534	5

Partner details of all interfaces:

Intf	Aggregate name	Partner Port-id	Port Priority	State	System-ID	System Priority	Aggr Key
1/1/2	lag2(mc)	0	65534	IE	00:00:00:00:00:00	0	0
1/1/5	lag5(mc)	3	255	ASFNCD	20:4c:03:b1:cd:e2	32768	2

Remote Actor details of all interfaces:

Intf	Aggregate name	Port id	Port Priority	State	System-ID	System Priority	Aggr Key
1/1/2	lag2(mc)	2	1	IE	02:01:00:00:00:00	65534	2
1/1/5	lag5(mc)	5	1	ALFNCD	02:01:00:00:00:00	65534	5

Remote Partner details of all interfaces:

Intf	Aggregate name	Partner Port-id	Port Priority	State	System-ID	System Priority	Aggr Key

```
1/1/2    lag2(mc)  0      65534    IE      00:00:00:00:00:00  0
1/1/5    lag5(mc)  2      255      ASFNCD  20:4c:03:b1:cd:e2  32768  2
```

```
# sw2 correction:
interface lag 255
  lacp mode active
```

Trouble Ticket 04.01

Hints

```
# On Agg2
show interface brief
show ip interface brief
show vlan
show vlan <id>
```

Solution

Customer has made a mistake and enabled VLAN 42 on the VSX LAG 5. When the ISL fails, the split brain disables the VSX LAG interfaces on the VSX Secondary. Any VLANs that are enabled on these VSX LAG interfaces are also disabled. Customer did not follow best practice, since the vsx-sync for multi-chassis was not enabled.

```
T1-Agg2-lab14-tt04(config)# show int brief
```

```
-----
Port      Native Mode  Type      Enabled Status Reason      Speed
Description
VLAN      (Mb/s)
-----
1/1/1     1      trunk SFP+DAC1  yes      up      --          10000  --
1/1/2     1      trunk SFP+DAC1  yes      down    Disabled by VSX  --      sw2
1/1/3     --     routed --        no       down    No XCVR installed  --      --
1/1/4     --     routed --        no       down    No XCVR installed  --      --
1/1/5     1      trunk SFP-BT    yes      down    Disabled by VSX  --      --
1/1/6     --     routed --        no       down    No XCVR installed  --      --
1/1/7     --     routed SFP-BT    yes      up      --          1000
routerB
1/1/8     --     routed --        no       down    No XCVR installed  --      --
1/1/9     --     routed --        no       down    No XCVR installed  --      --
1/1/10    --     routed --        no       down    No XCVR installed  --      --
1/1/11    --     routed --        no       down    No XCVR installed  --      --
1/1/12    --     routed --        no       down    No XCVR installed  --      --
1/1/13    --     routed --        no       down    No XCVR installed  --      --
1/1/14    --     routed --        no       down    No XCVR installed  --      --
1/1/15    --     routed --        no       down    No XCVR installed  --      --
1/1/16    --     routed --        no       down    No XCVR installed  --      --
1/1/17    --     routed --        no       down    No XCVR installed  --      --
1/1/18    --     routed --        no       down    No XCVR installed  --      --
```

1/1/19	--	routed	--	no	down	No XCVR installed	--	--
1/1/20	--	routed	--	no	down	No XCVR installed	--	--
1/1/21	--	routed	--	no	down	No XCVR installed	--	--
1/1/22	--	routed	--	no	down	No XCVR installed	--	--
1/1/23	--	routed	--	no	down	No XCVR installed	--	--
1/1/24	--	routed	--	no	down	No XCVR installed	--	--
1/1/25	--	routed	--	no	down	No XCVR installed	--	--
1/1/26	--	routed	--	no	down	No XCVR installed	--	--
1/1/27	--	routed	--	no	down	No XCVR installed	--	--
1/1/28	--	routed	--	no	down	No XCVR installed	--	--
1/1/29	--	routed	--	no	down	No XCVR installed	--	--
1/1/30	--	routed	--	no	down	No XCVR installed	--	--
1/1/31	--	routed	--	no	down	No XCVR installed	--	--
1/1/32	--	routed	--	no	down	No XCVR installed	--	--
1/1/33	--	routed	--	no	down	No XCVR installed	--	--
1/1/34	--	routed	--	no	down	No XCVR installed	--	--
1/1/35	--	routed	--	no	down	No XCVR installed	--	--
1/1/36	--	routed	--	no	down	No XCVR installed	--	--
1/1/37	--	routed	--	no	down	No XCVR installed	--	--
1/1/38	--	routed	--	no	down	No XCVR installed	--	--
1/1/39	--	routed	--	no	down	No XCVR installed	--	--
1/1/40	--	routed	--	no	down	No XCVR installed	--	--
1/1/41	--	routed	--	no	down	No XCVR installed	--	--
1/1/42	--	routed	--	no	down	No XCVR installed	--	--
1/1/43	--	routed	--	no	down	No XCVR installed	--	--
1/1/44	--	routed	--	no	down	No XCVR installed	--	--
1/1/45	--	routed	SFP28DAC0.65	yes	up		25000	VSX
keepalive								
1/1/46	1	trunk	SFP28DAC0.65	yes	up		25000	agg1
1/1/47	1	trunk	SFP28DAC0.65	yes	up		25000	--
1/1/48	--	routed	--	no	down	No XCVR installed	--	--
1/1/49	--	routed	--	no	down	No XCVR installed	--	--
1/1/50	--	routed	--	no	down	No XCVR installed	--	--
1/1/51	--	routed	--	no	down	No XCVR installed	--	--
1/1/52	--	routed	--	no	down	No XCVR installed	--	--
1/1/53	--	routed	--	no	down	No XCVR installed	--	--
1/1/54	--	routed	--	no	down	No XCVR installed	--	--
1/1/55	--	routed	--	no	down	No XCVR installed	--	--
1/1/56	--	routed	--	no	down	No XCVR installed	--	--
loopback0	--	routed	--	yes	up		--	--
loopback2	--	routed	--	yes	up		--	--
vlan1	--	--	--	yes	down	Disabled by VSX	--	--
vlan11	--	--	--	yes	down	Disabled by VSX	--	--
vlan12	--	--	--	yes	down	Disabled by VSX	--	--
vlan40	--	--	--	yes	up		--	--
vlan42	--	--	--	yes	down	Disabled by VSX	--	--
vlan212	--	--	--	yes	up		--	--
vlan242	--	--	--	yes	up		--	--
lag2	1	trunk	--	yes	down	--	auto	
lag5	1	trunk	--	yes	down	--	auto	
lag256	1	trunk	--	yes	up	--	50000	

```
T1-Agg2-lab14-tt04(config)# show vlan 42
```

VLAN	Name	Status	Reason	Type	Interfaces
42	VLAN42	up	ok	static	1/1/1, lag5, lag256

```
T1-Agg2-lab14-tt04(config-if)# show run int lag5
```

```
interface lag 5 multi-chassis
  no shutdown
  description aruba-gw1
  no routing
  vlan trunk native 1
  vlan trunk allowed 1,11-12,40,42
  lacp mode active
  apply policy mirror in
  exit
```

```
# Agg2: Remove VLAN 42 from the LAG 5
```

```
int lag 5
  no VLAN trunk allow 42
```

```
or
```

```
# Agg1: Apply VSX Best Practice to sync mclag-interfaces configuration
```

```
vsx
  vsx-sync mclag-interfaces
```

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

Customer initially had a working OSPF configuration. This configuration did have differences in the port IDs (SVI41 on Agg1, SVI42 on Agg2), but that works fine since each node has its local configuration.

The customer later decided to apply the VSX OSPF sync, as this is one of the best practices. This requires both nodes to have the same port IDs, since the primary will sync its OSPF ports to the secondary.

Since the primary does not have SVI42, any OSPF process configuration on SVI42 on Agg2 is automatically removed by the sync process.

```
T1-Agg2-lab13(config-if-vlan)# show ip ospf interface brief
OSPF Process ID 1 VRF default
=====
```

Total Number of Interfaces: 5

Interface	Area	IP Address/Mask	Cost	State	Status
loopback0	0.0.0.0	10.14.0.3/32	0	Loopback	up
vlan1	0.0.0.0	10.14.1.3/24	100	Backup-dr	up
vlan11	0.0.0.0	10.14.11.3/24	100	Dr-other	up
vlan12	0.0.0.0	10.14.12.3/24	100	Dr-other	up
vlan40	0.0.0.0	10.14.40.3/24	100	Dr-other	up

```
T1-Agg2-lab13(config-if-vlan)# show event -r -n 10
```

Event logs from current boot

```
2020-11-13T05:56:31.832431-05:00 T1-Agg2-lab13 vsx-syncd[6253]: Event|7601|LOG_ERR|AMM|-
|Configuration sync error : Missing reference in database while syncing configuration. Use
"show running-config vsx-sync peer-diff" to help identify the missing reference.
```

2020-11-13T05:56:31.720992-05:00 T1-Agg2-lab13 vsx-syncd[6253]: Event|7602|LOG_INFO|AMM| -
|Configuration sync update : Missing Reference Error: Could not find Port vlan41 on secondary
VSX database.

```
T1-Agg2-lab13(config-if-vlan)# show running-config vsx-sync peer-diff
qos trust dscp
    ip helper-address 10.252.1.21
    ip helper-address 10.252.1.21
qos trust dscp
    ip helper-address 10.252.1.21
    ip helper-address 10.252.1.21
--- /tmp/running-config-vsx.4094      2020-11-13 05:57:29.464659146 -0500
+++ /tmp/peer-running-config-vsx.4094  2020-11-13 05:57:29.454659210 -0500
@@ -1,7 +1,7 @@
!
!Version ArubaOS-CX GL.10.05.0001
!export-password: default
-lldp management-ipv4-address 10.251.14.3
+lldp management-ipv4-address 10.251.14.2
access-list ip igmpfilter
    vsx-sync
!
@@ -77,8 +77,14 @@
    ipv6 ospfv3 passive
interface vlan 40
    ip ospf 1 area 0.0.0.0
-interface vlan 42
+interface vlan 41
+    ip ospf 1 area 0.0.0.0
+    ip ospf dead-interval 15
+    ip ospf hello-interval 5
+    ip ospf network point-to-point
+    ip ospf authentication keychain
+    ip ospf keychain asts
+    ipv6 ospfv3 1 area 0.0.0.0
+    ipv6 ospfv3 network point-to-point
interface vlan 212
    vsx-sync active-gateways policies
@@ -86,7 +92,8 @@
    active-gateway ip 10.14.212.1
    ip ospf 2 area 0.0.0.0
    ip ospf passive
-interface vlan 242
+interface vlan 241
+    ip ospf 2 area 0.0.0.0
+    ip ospf network point-to-point
vsx
    system-mac 02:01:00:00:00:00
```



```
# On Agg1: disable vsx-sync for ospf.  
# On Agg2, enable SVI42 for OSPF process 1 area 0
```

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

There is a filter on the core network (sw1) that prevents switch loopback IP traffic communication, this affects the Keepalive traffic when it passes over the Sw1 device.

With the ISL active, OSPF takes the path over the ISL, this is why the Keepalive seems ok under normal conditions.

When the ISL breaks, the Keepalive must be sent over the OSPF network via Sw1 and it is hit with the filter.

With the ISL UP (make sure to wait for link up delay timer to pass)

```
T1-Agg2-lab14-tt06(config-vsx)# show ip route 10.14.0.2
```

Displaying ipv4 routes selected for forwarding

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

```
10.14.0.2/32, vrf default
    via 10.14.40.2, [110/100], ospf
```

```
T1-Agg2-lab14-tt06(config-vsx)# traceroute 10.14.0.2
```

```
traceroute to 10.14.0.2 (10.14.0.2), 1 hops min, 30 hops max, 3 sec. timeout, 3 probes
 1  10.14.0.2  0.169ms  0.090ms  0.076ms
```

```
T1-Agg2-lab14-tt06(config-vsx)# show vsx status keepalive
```

```
Keepalive State      : Keepalive-Established
Last Established      : Fri Nov 13 06:35:16 2020
Last Failed          : Fri Nov 13 06:27:47 2020
Peer System Id       : 02:01:00:00:00:00
Peer Device Role     : primary
```

Keepalive Counters

```
Keepalive Packets Tx      : 738
Keepalive Packets Rx      : 287
Keepalive Timeouts        : 0
Keepalive Packets Dropped : 0
```

with the ISL down

```
T1-Agg2-lab14-tt06(config-vsx)# show ip route 10.14.0.2
```

Displaying ipv4 routes selected for forwarding

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

```
10.14.0.2/32, vrf default
    via 10.14.42.4, [110/200], ospf
```

```
T1-Agg2-lab14-tt06(config-vsx)# traceroute 10.14.0.2
```

traceroute to 10.14.0.2 (10.14.0.2), 1 hops min, 30 hops max, 3 sec. timeout, 3 probes

```
 1  10.14.42.4  0.113ms  0.078ms  0.074ms
 2  * * *
```

On Sw1

```
T1-Sw1-lab14-tt06(config)# show run int vlan 41
```

```
interface vlan41
    apply access-list ip core routed-in
```

...

```
interface vlan42
    apply access-list ip core routed-in
```

```
T1-Sw1-lab14-tt06(config)# access-list ip core
```

```
T1-Sw1-lab14-tt06(config-acl-ip)# show run cur
```

```
access-list ip core
```

```
 10 deny any any 10.0.0.0/255.0.255.0
 20 permit any any any
```

```
T1-Sw1-lab14-tt06(config-acl-ip)# no 10
```

Note:

It will take a few moments longer for the split-brain detection due to the routing convergence that needs to be completed before KeepAlive link can be re-established.

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

The customer made several mistakes.

They had set the device-profile as more preferred than the aaa

T1-Sw2-lab14-tt07# show port-access clients

Port Access Clients

Port	MAC Address	Onboarded Method	Status	Role
1/1/4	00:50:56:b1:9b:b7	device-profile	Success	printer

T1-Sw2-lab14-tt07# show run int 1/1/4

```
interface 1/1/4
  no shutdown
  mtu 9198
  apply policy counter1 in
  description PC4
  no routing
  vlan access 12
  spanning-tree bpduguard
  port-access fallback-role asts-fallback
  port-access onboarding-method precedence device-profile aaa
  aaa authentication port-access critical-role asts-critical
  aaa authentication port-access dot1x authenticator
    eapol-timeout 30
    max-eapol-requests 1
    max-retries 1
  enable
exit
T1-Sw2-lab14-tt07#
```

```
# on Sw2
int 1/1/4
port-access onboarding-method precedence aaa device-profile
```

```
T1-Sw2-lab14-tt07(config)# show port-access clients
```

Port Access Clients

Port	MAC Address	Onboarded Method	Status	Role
1/1/4	00:50:56:b1:9b:b7		In-Progress	

```
T1-Sw2-lab14-tt07(config)#
```

```
### Sw2 debug
debug portaccess all
```

```
T1-Sw2-lab14-tt07(config)# show debug buffer sev err
```

```
show debug buffer
```

```
2020-11-13:07:07:57.910565|port-
accessd|LOG_ERR|MSTR|1|PORTACCESS|PORTACCESS_DOT1X_RADIUS|RADIUS request failed with error
Server is not found
```

Review of server groups shows groups 'aaa' and 'cppm'

```
T1-Sw2-lab14-tt07(config)# show aaa server-groups radius
```

```
***** AAA Mechanism RADIUS *****
```

GROUP NAME	SERVER NAME	TLS	PORT
VRF	PRIORITY		
aaa			
cppm	cppm1.arubatraining.com		1812
default	1		
cppm	cppm2.arubatraining.com		1812
default	2		

```
-----
radius      (default)      | 10.252.1.23                |      | 1812
| default      | 1
radius      (default)      | cppm.arubatraining.com    |      | 1812
| default      | 2
radius      (default)      | cppm1.arubatraining.com   |      | 1812
| default      | 3
radius      (default)      | cppm2.arubatraining.com   |      | 1812
| default      | 4
-----
```

```
-----
T1-Sw2-lab14-tt07(config)#
```

```
# Lookup which server group is used by the 802.1X authenticator
```

```
T1-Sw2-lab14-tt07(config)# show run aaa authentication port-access
aaa authentication port-access dot1x authenticator
    radius server-group aaa
    enable
```

```
# Need to correct the group aaa to use the server cppm.arubatraining.com
```

```
aaa group server radius aaa
    server cppm.arubatraining.com
```

```
# Attempt authentication on PC4 again (restart the 'Wired AutoConfig' service)
```

```
# Review event log
```

```
T1-Sw2-lab14-tt07(config-if)# show event -r -n 10
```

```
-----
Event logs from current boot
-----
```

```
2020-11-13T07:15:19.717894-05:00 T1-Sw2-lab14-tt07 radius-srv-trkd[3052]:
Event|2304|LOG_INFO|MSTR|1|RADIUS Server with Address:cppm.arubatraining.com, Authport:1812,
VRF_ID:0 is "unreachable"
```

```
# Verify reachability of the host >> this reveals an issue with the name resolution
```

```
T1-Sw2-lab14-tt07(config-if)# ping cppm.arubatraining.com
PING cppm.arubatraining.com (10.252.1.21) 100(128) bytes of data.
108 bytes from 10.252.1.21: icmp_seq=1 ttl=124 time=3.08 ms
108 bytes from 10.252.1.21: icmp_seq=2 ttl=124 time=4.06 ms
108 bytes from 10.252.1.21: icmp_seq=3 ttl=124 time=3.87 ms
108 bytes from 10.252.1.21: icmp_seq=4 ttl=124 time=4.45 ms
```

```
# T1-Sw2-lab14-tt07(config-if)# show ip dns
```

```
VRF Name : branch
```

Host Name	Address

VRF Name : default
Name Server(s) : 10.252.1.23

Host Name	Address

cppm.arubatraining.com	10.252.1.21
cppm1.arubatraining.com	10.252.1.21
cppm2.arubatraining.com	10.252.1.22

VRF Name : iot

Host Name	Address

VRF Name : mgmt

Host Name	Address

>> This reveals the customer has defined local hostnames on the switch, but they made a mistake.

>> Remove the local host entry

```
T1-Sw2-lab14-tt07(config)# no ip dns host cppm.arubatraining.com 10.252.1.21
```

>> Attempt to ping again

```
T1-Sw2-lab14-tt07(config)# ping cppm.arubatraining.com
ping: unknown host cppm.arubatraining.com
T1-Sw2-lab14-tt07(config)#
```

>> There is still an issue with the name resolution. Check 'show ip dns' output again.

>> Customer has used wrong IP address for the DNS host.

```
T1-Sw2-lab14-tt07(config)# no ip dns server-address 10.252.1.23
T1-Sw2-lab14-tt07(config)# ip dns server-address 10.252.1.21
```

>> Test again, should now be successful

```
T1-Sw2-lab14-tt07(config)# ping cppm.arubatraining.com
PING cppm.arubatraining.com (10.252.1.23) 100(128) bytes of data.
108 bytes from 10.252.1.23: icmp_seq=1 ttl=60 time=4.54 ms
```

```
108 bytes from 10.252.1.23: icmp_seq=2 ttl=60 time=3.22 ms
108 bytes from 10.252.1.23: icmp_seq=3 ttl=60 time=4.77 ms
108 bytes from 10.252.1.23: icmp_seq=4 ttl=60 time=4.23 ms
108 bytes from 10.252.1.23: icmp_seq=5 ttl=60 time=2.55 ms
```

```
--- cpm.arubatraining.com ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4006ms
rtt min/avg/max/mdev = 2.555/3.869/4.779/0.846 ms
```

>> Clear logs and Attempt to authenticate with the client again (restart service)

```
T1-Sw2-lab14-tt07(config)# clear debug buffer
T1-Sw2-lab14-tt07(config)# int 1/1/4
T1-Sw2-lab14-tt07(config-if)# shut
T1-Sw2-lab14-tt07(config-if)# no shut
```

>> Review the logs

```
T1-Sw2-lab14-tt07(config)# show debug buffer sev err
```

```
-----
show debug buffer
-----
```

```
2020-11-13:07:22:35.973518|port-
accessd|LOG_ERR|MSTR|1|PORTACCESS|PORTACCESS_DOT1X_PROTOCOL|logID=15414 logID=15414 Event
handler of Dot1XPort '1/1/4' for event 'Profile Status For Client On Port' in state
'DISCOVERING' returned 'Error'
2020-11-13:07:22:37.927547|radius-srv-trkd|LOG_ERR|MSTR|1|RADIUS|RADIUS_TRACKING|RADIUS
request for context 0xffff84018f20 client request id 40 radius id 39 timedout error Group is
not found. Inform to client.
2020-11-13:07:22:41.463793|port-
accessd|LOG_ERR|MSTR|1|PORTACCESS|PORTACCESS_DOT1X_RADIUS|RADIUS request failed with error
Server timedout
2020-11-13:07:22:44.966713|port-
accessd|LOG_ERR|MSTR|1|PORTACCESS|PORTACCESS_DOT1X_RADIUS|RADIUS request failed with error
Server timedout
2020-11-13:07:23:02.986251|port-
accessd|LOG_ERR|MSTR|1|PORTACCESS|PORTACCESS_DOT1X_RADIUS|RADIUS request failed with error
Server timedout
2020-11-13:07:23:37.961179|radius-srv-trkd|LOG_ERR|MSTR|1|RADIUS|RADIUS_TRACKING|RADIUS
request for context 0xffff84002bb0 client request id 41 radius id 40 timedout error Group is
not found. Inform to client.
```

>> Ping worked, but RADIUS Timedout shown. This requires correct RADIUS settings, such as radius secret.

>> Reconfigure RADIUS host with secret aruba123

```
T1-Sw2-lab14-tt07(config)# radius-server host cpm.arubatraining.com key plaintext aruba123
```


>> Test again on the client, still fails.

>> On Sw2, review clients and logs

T1-Sw2-lab14-tt07(config-if)# **show port-access clients detail**

Port Access Client Status Details:

Client 00:50:56:b1:9b:b7, asts-employee-dur

=====

Session Details

Port : 1/1/4
Session Time : 88s
IPv4 Address :
IPv6 Address :

Authentication Details

Status : dot1x Authenticated
Auth Precedence : dot1x - Authenticated, mac-auth - Not attempted

Authorization Details

Role :
Status : Not Ready

>> Authenticated indicates success RADIUS

>> Not Ready indicates issue with the role.

T1-Sw2-lab14-tt07(config-if)# **show event -r -n 10**

Event logs from current boot

2020-11-13T07:26:40.905287-05:00 T1-Sw2-lab14-tt07 port-accessd[3327]: Event|7708|LOG_WARN|Throttled 1 Messages
2020-11-13T07:26:35.981779-05:00 T1-Sw2-lab14-tt07 port-accessd[3327]: Event|10502|LOG_INFO|MSTR|1|Port 1/1/4 is blocked by port-access
2020-11-13T07:26:35.958783-05:00 T1-Sw2-lab14-tt07 intfd[772]: Event|403|LOG_INFO|||Link status for interface 1/1/4 is up
2020-11-13T07:26:35.932815-05:00 T1-Sw2-lab14-tt07 lldpd[3305]: Event|110|LOG_INFO|MSTR|1|Configured LLDP reinit-delay to 2
2020-11-13T07:26:29.307014-05:00 T1-Sw2-lab14-tt07 port-accessd[3327]: Event|10503|LOG_INFO|MSTR|1|Port 1/1/4 is unblocked by port-access

```
2020-11-13T07:26:29.286421-05:00 T1-Sw2-lab14-tt07 intfd[772]: Event|404|LOG_INFO||Link
status for interface 1/1/4 is down
2020-11-13T07:26:28.803669-05:00 T1-Sw2-lab14-tt07 port-accesssd[3327]:
Event|7708|LOG_INFO|MSTR|1|Certificate cppm.arubatraining.com verified and accepted
2020-11-13T07:26:28.790415-05:00 T1-Sw2-lab14-tt07 ops-switchd[717]:
Event|2110|LOG_INFO|MSTR|1|Deleted Mac based VLAN entry for 00:50:56:b1:9b:b7 with VLAN 11 on
port 1/1/4
2020-11-13T07:26:28.766961-05:00 T1-Sw2-lab14-tt07 port-accesssd[3327]:
Event|10502|LOG_INFO|MSTR|1|Port 1/1/4 is blocked by port-access
2020-11-13T07:26:28.293133-05:00 T1-Sw2-lab14-tt07 radius-srv-trkd[3052]:
Event|2304|LOG_INFO|MSTR|1|RADIUS Server with Address:cppm.arubatraining.com, Authport:1812,
VRF_ID:0 is "reachable"
```

```
>> Switch attempted to get role from CPPM (Certificate validation was done indicates HTTPS
connection was made).
```

```
>> Review debug buffer
```

```
T1-Sw2-lab14-tt07(config-if)# show debug buffer
```

```
2020-11-13:07:26:40.762878|port-
accesssd|LOG_DEBUG|MSTR|1|PORTACCESS|PORTACCESS_SERVICES|logID=15782 Handling db write(Update)
success for client 00:50:56:b1:9b:b7 in port 1/1/4
2020-11-13:07:26:40.797961|port-accesssd|LOG_INFO|MSTR|1|PORTACCESS|PORTACCESS_DOT1X|Checking
chain of 3 certificates against C = US, ST = California, L = Santa Clara - Santa Clara -
Santa Clara - Santa Clara - Santa Clara, O = Aruba Networks - Aruba Networks - Aruba Networks
- Aruba Networks, OU = Aruba Training - Aruba Training - Aruba Training - Aruba Training -
Aruba Training, CN = ASTS - ClearPass Onboard Local Certificate Authority - Troubleshooting
and Solutions, emailAddress = aefd20a3-61a8-4e40-9508-ec418725b514@example.com
2020-11-13:07:26:40.798224|port-accesssd|LOG_INFO|MSTR|1|PORTACCESS|PORTACCESS_DOT1X|Checking
certificate at depth 0 with subject: C = US, ST = California, L = Santa Clara - Santa Clara -
Santa Clara - Santa Clara - Santa Clara, O = Aruba Networks - Aruba Networks - Aruba Networks
- Aruba Networks, OU = Aruba Training - Aruba Training - Aruba Training - Aruba Training -
Aruba Training, CN = cppm.arubatraining.com, emailAddress = cppm@arubatraining.com
2020-11-13:07:26:40.798285|port-accesssd|LOG_INFO|MSTR|1|PORTACCESS|PORTACCESS_DOT1X|Checking
certificate at depth 1 with subject: C = US, ST = California, L = Santa Clara - Santa Clara -
Santa Clara - Santa Clara - Santa Clara, O = Aruba Networks - Aruba Networks - Aruba Networks
- Aruba Networks, OU = Aruba Training - Aruba Training - Aruba Training - Aruba Training -
Aruba Training, CN = ClearPass Onboard Local Certificate Authority (Signing) -
Troubleshooting and Solutions, emailAddress = aefd20a3-61a8-4e40-9508-
ec418725b514@example.com
2020-11-13:07:26:40.798391|port-accesssd|LOG_INFO|MSTR|1|PORTACCESS|PORTACCESS_DOT1X|Checking
certificate at depth 2 with subject: C = US, ST = California, L = Santa Clara - Santa Clara -
Santa Clara - Santa Clara - Santa Clara, O = Aruba Networks - Aruba Networks - Aruba Networks
- Aruba Networks, OU = Aruba Training - Aruba Training - Aruba Training - Aruba Training -
Aruba Training, CN = ASTS - ClearPass Onboard Local Certificate Authority - Troubleshooting
and Solutions, emailAddress = aefd20a3-61a8-4e40-9508-ec418725b514@example.com
2020-11-13:07:26:40.799894|port-
accesssd|LOG_INFO|MSTR|1|PORTACCESS|PORTACCESS_DOT1X|Certificate with subject: C = US, ST =
```

```
California, L = Santa Clara - Santa Clara - Santa Clara - Santa Clara - Santa Clara, O =
Aruba Networks - Aruba Networks - Aruba Networks - Aruba Networks, OU = Aruba Training -
Aruba Training - Aruba Training - Aruba Training - Aruba Training, CN =
cppm.arubatraining.com, emailAddress = cppm@arubatraining.com successfully validated against
TA with subject: C = US, ST = California, L = Santa Clara - Santa Clara - Santa Clara - Santa
Clara - Santa Clara, O = Aruba Networks - Aruba Networks - Aruba Networks - Aruba Networks,
OU = Aruba Training - Aruba Training - Aruba Training - Aruba Training - Aruba Training, CN =
ASTS - ClearPass Onboard Local Certificate Authority - Troubleshooting and Solutions,
emailAddress = aefd20a3-61a8-4e40-9508-ec418725b514@example.com
2020-11-13:07:26:40.799920|port-
accessd|LOG_DEBUG|MSTR|1|PORTACCESS|PORTACCESS_DOT1X_PROTOCOL|Cert chain passed crypto
validation.
2020-11-13:07:26:40.806809|port-
accessd|LOG_DEBUG|MSTR|1|PORTACCESS|PORTACCESS_ROLE|logID=15782 Result of download operation
for role asts_aruba_dur_employee-3118-6 is No error
2020-11-13:07:26:40.807196|port-
accessd|LOG_DEBUG|MSTR|1|PORTACCESS|PORTACCESS_ROLE|logID=15782 Parsing the XML for the
cprole asts_aruba_dur_employee-3118-6 failed
```

>> This indicates either error in the syntax or error with the DUR user account

>> update the radius configuration clearpass username and password

```
T1-Sw2-lab14-tt07(config)# radius-server host cppm.arubatraining.com clearpass-username
duradmin clearpass-password plaintext aruba123
```

>> test again, the authentication should now be successful

```
T1-Sw2-lab14-tt07(config-if)# show port-access clients
```

Port Access Clients

Port	MAC Address	Onboarded Method	Status	Role
1/1/4	00:50:56:b1:9b:b7	dot1x	Success	asts_aruba_dur_employee...

Trouble Ticket 08.01

Hints

```
# On Sw2
show port-access clients
show port-access clients detail
```

Solution

The customer has assigned a VLAN in the DUR that does not exist on the switch.

```
T1-Sw2-lab14-tt08(config-if)# show port-access clients
```

Port Access Clients

Port	MAC Address	Onboarded Method	Status	Role
1/1/4	00:50:56:b1:9b:b7		Fail	asts_aruba_dur_error1-3...

```
T1-Sw2-lab14-tt08(config-if)# show port-access clients detail
```

Port Access Client Status Details:

Client 00:50:56:b1:9b:b7, asts-error1

=====

Session Details

Port : 1/1/4
Session Time : 64s
IPv4 Address :
IPv6 Address :

Authentication Details

Status : dot1x Authenticated
Auth Precedence : dot1x - Authenticated, mac-auth - Not attempted

Authorization Details

Role : asts_aruba_dur_error1-3125-7
Status : Failed, Failed to assign VLAN

Role Information:

Name : asts_aruba_dur_error1-3125-7
Type : clearpass
Status: Completed

```

Reauthentication Period      :
Authentication Mode          :
Session Timeout              :
Client Inactivity Timeout    :
Description                  :
Gateway Zone                 :
UBT Gateway Role             :
Access VLAN                  : 111
Native VLAN                  :
Allowed Trunk VLANs          :
Access VLAN Name             :
Native VLAN Name             :
Allowed Trunk VLAN Names     :
MTU                           :
QOS Trust Mode               :
STP Administrative Edge Port :
PoE Priority                  :
Captive Portal Profile       :
Policy                       : employee_asts_aruba_dur_error1-3125-7
  
```

Access Policy Details:

```

Policy Name   : employee_asts_aruba_dur_error1-3125-7
Policy Type   : Downloaded
Policy Status : Applied
  
```

SEQUENCE	CLASS	TYPE	ACTION
10	any_asts_aruba_dur_error1...	ipv4	permit

Class Details:

```

class ip any_asts_aruba_dur_error1-3125-7
  10 match any any any
  
```

```

>> Define VLAN 111 on the switch
T1-Sw2-lab14-tt08(config)# vlan 111
T1-Sw2-lab14-tt08(config-vlan-111)# exit

T1-Sw2-lab14-tt08(config)# do port-access reauthenticate interface 1/1/4

T1-Sw2-lab14-tt08(config)# show port-access clients
  
```

Port Access Clients

Port	MAC Address	Onboarded Method	Status	Role

1/1/4 00:50:56:b1:9b:b7 dot1x Success asts_aruba_dur_error1-3...

Aruba Training-Confidential

Trouble Ticket 08.02

Hints

```
# On Sw2
show port-access clients
show port-access clients detail
```

Solution

The customer has assigned a different gateway zone name in the DUR that does not exist on the switch.

```
T1-Sw2-lab14-tt08(config)# show port-access clients
```

Port Access Clients

Port	MAC Address	Onboarded Method	Status	Role
1/1/4	00:50:56:b1:9b:b7		In-Progress	asts_aruba_dur_error2-3...

```
T1-Sw2-lab14-tt08(config)# show port-access clients detail
```

Port Access Client Status Details:

Client 00:50:56:b1:9b:b7, asts-error2

=====

Session Details

Port : 1/1/4
Session Time : 299s
IPv4 Address :
IPv6 Address :

Authentication Details

Status : dot1x Authenticated
Auth Precedence : dot1x - Authenticated, mac-auth - Not attempted

Authorization Details

Role : asts_aruba_dur_error2-3126-8
Status : In Progress

Role Information:

Name : asts_aruba_dur_error2-3126-8
Type : clearpass

Status: Completed

```

-----
Reauthentication Period      :
Authentication Mode          :
Session Timeout              :
Client Inactivity Timeout    :
Description                  :
Gateway Zone                  : gw1
UBT Gateway Role             :
Access VLAN                  : 11
Native VLAN                  :
Allowed Trunk VLANs          :
Access VLAN Name             :
Native VLAN Name             :
Allowed Trunk VLAN Names     :
MTU                           : 9198
QOS Trust Mode               :
STP Administrative Edge Port :
PoE Priority                  :
Captive Portal Profile       :
Policy                        : employee_asts_aruba_dur_error2-3126-8
  
```

UBT Zone Details:

Access Policy Details:

```

Policy Name   : employee_asts_aruba_dur_error2-3126-8
Policy Type   : Downloaded
Policy Status : Applied
  
```

SEQUENCE	CLASS	TYPE	ACTION
10	any_asts_aruba_dur_error2...	ipv4	permit

Class Details:

```

class ip any_asts_aruba_dur_error2-3126-8
  10 match any any any
  
```

```

# Correct the zone configuration on Sw2 to match the zone name returned by ClearPass
no ubt zone aruba_gw vrf default
ubt zone gw1 vrf default
  primary-controller ip 10.14.1.6
enable
  
```

T1-Sw2-lab14-tt08(config)# show port-access clients

Port Access Clients

Port	MAC Address	Onboarded Method	Status	Role
1/1/4	00:50:56:b1:9b:b7	dot1x	Success	asts_aruba_dur_error2-3...

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

The ClearPass administrator has made a typing error in the role name.

```
Sw2-lab14-tt08(config-if)# show event -r -n 10
```

```
-----
Event logs from current boot
-----
```

```
2020-12-04T11:37:40.369148-05:00 Sw2-lab14-tt08 hpe-restd[1089]: Event|4605|LOG_INFO|AMM| -
|Session ended for user admin, session 1BkeTvi9A0GfEXbQtm3aw==
```

```
2020-12-04T11:37:40.367761-05:00 Sw2-lab14-tt08 hpe-restd[1089]: Event|4608|LOG_INFO|AMM| -
|Authorization allowed for user admin, for resource SessionMgmt, with action POST
```

```
2020-12-04T11:37:36.216576-05:00 Sw2-lab14-tt08 port-accesssd[3516]:
```

```
Event|9301|LOG_ERR|MSTR|1|Failed to apply ClearPass role - asts_aruba_dur_error3_adv-3128-2 -
Error at Line 10 - vlan access
```

```
2020-12-04T11:37:36.216362-05:00 Sw2-lab14-tt08 port-accesssd[3516]:
```

```
Event|9301|LOG_ERR|MSTR|1|Failed to apply ClearPass role - asts_aruba_dur_error3_adv-3128-2 -
Invalid input:
```

```
# Role details: check CPPM configuration for the enforcement profile
```

Request Details

Summary

Input

Output

Enforcement Profiles:	asts-aruba-dur-error3-adv
System Posture Status:	UNKNOWN (100)
Audit Posture Status:	UNKNOWN (100)

RADIUS Response

Radius:Aruba:Aruba-CPPM-Role

asts_aruba_dur_error3_adv-3128-2

class ip any
10 match any any any
exit
port-access policy employee
10 class ip any
exit
port-access role employee
associate policy employee
mtu 9198
vlan access pc
exit

Showing 3 of 1-100 records

Change Status

Show Configuration

Export

Show Logs

Close

```

class ip any
10 match any any any
exit
port-access policy employee
10 class ip any
exit
port-access role employee
associate policy employee
mtu 9198
vlan access pc
exit

```

```

vlan access pc
>> should be
vlan access name pc

```

>> This demonstrates why it is highly recommended to use the ClearPass role UI instead of entering the role commands manually in the ClearPass enforcement profile.

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

```
Sw2-lab14-tt08(config-if)# show port-access clients

Port Access Clients
-----
Port      MAC Address      Onboarded      Status      Role
          Method
-----
1/1/4     00:50:56:b1:9b:b7      Fail          asts_aruba_dur_error4_a...

Sw2-lab14-tt08(config-if)# show port-access clients detail

Port Access Client Status Details:

Client 00:50:56:b1:9b:b7, asts-error4
=====
Session Details
-----
Port      : 1/1/4
Session Time : 3915s
IPv4 Address :
IPv6 Address :

Authentication Details
-----
Status      : dot1x Authenticated
Auth Precedence : dot1x - Authenticated, mac-auth - Not attempted

Authorization Details
-----
Role      : asts_aruba_dur_error4_adv-3129-6
Status    : Failed, Failed to assign VLAN

Role Information:
```

Name : asts_aruba_dur_error4_adv-3129-6
Type : clearpass
Status: Completed

```
-----
Reauthentication Period      :
Authentication Mode          :
Session Timeout              :
Client Inactivity Timeout    :
Description                   :
Gateway Zone                  :
UBT Gateway Role              :
Access VLAN                   :
Native VLAN                   :
Allowed Trunk VLANs           :
Access VLAN Name              : employee
Native VLAN Name              :
Allowed Trunk VLAN Names      :
MTU                           : 9198
QOS Trust Mode                :
STP Administrative Edge Port  :
PoE Priority                   :
Captive Portal Profile        :
Policy                        : employee_asts_aruba_dur_error4_adv-3129-6
```

Access Policy Details:

Policy Name : employee_asts_aruba_dur_error4_adv-3129-6
Policy Type : Downloaded
Policy Status : Applied

SEQUENCE	CLASS	TYPE	ACTION
10	any_asts_aruba_dur_error4...	ipv4	permit

Class Details:

```
class ip any_asts_aruba_dur_error4_adv-3129-6
  10 match any any any
```

>> Other method to see this via debugging
debug portaccess all
clear debug buffer
port-access reauthenticate interface 1/1/4

Sw2-lab14-tt08(config)# show debug buffer severity err

```
-----
show debug buffer
```

```
2020-11-13:12:27:08.832731|port-  
accessd|LOG_ERR|MSTR|1|PORTACCESS|PORTACCESS_SERVICES|logID=14030 VLAN with name employee is  
not found.
```

Solution

```
Sw2-lab14-tt08(config)# vlan 11  
Sw2-lab14-tt08(config-vlan-11)# name employee  
Sw2-lab14-tt08(config-vlan-11)# exit  
Sw2-lab14-tt08(config)#  
Sw2-lab14-tt08(config)# show vlan
```

VLAN	Name	Status	Reason	Type	Interfaces
1	DEFAULT_VLAN_1	up	ok	default	1/1/1- 1/1/3,1/1/5-1/1/24, 1/1/28,lag255
11	employee	up	ok	static	1/1/27,lag255
12	VLAN12	up	ok	static	1/1/4,lag255
15	VLAN15	up	ok	static	1/1/27
19	VLAN19	down	no_member_port	static	
46	VLAN46	up	ok	static	1/1/27
47	VLAN47	up	ok	static	1/1/27
61	VLAN61	up	ok	static	1/1/27
62	VLAN62	up	ok	static	1/1/27
212	VLAN212	up	ok	static	lag255
4000	VLAN4000	down	no_member_port	static	

```
Sw2-lab14-tt08(config)# show port-access clients
```

Port Access Clients

Port	MAC Address	Onboarded Method	Status	Role
1/1/4	00:50:56:b1:9b:b7	dot1x	Success	asts_aruba_dur_error4_a...

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

```
Sw2-lab14-tt08(config)# show port-access clients detail
```

Port Access Client Status Details:

Client 00:50:56:b1:9b:b7, asts-error5

=====

Session Details

```
Port      : 1/1/4
Session Time : 4326s
IPv4 Address :
IPv6 Address :
```

Authentication Details

```
Status      : dot1x Authenticated
Auth Precedence : dot1x - Authenticated, mac-auth - Not attempted
```

Authorization Details

```
Role      :
Status : Not Ready
```

```
Sw2-lab14-tt08(config)# show event -r -n 10 -s err
```

Event logs from current boot

```
2020-11-13T14:35:07.261056-05:00 Sw2-lab14-tt08 port-accessd[3369]:
Event|9301|LOG_ERR|MSTR|1|Failed to apply ClearPass role - asts_aruba_dur_error5_adv-3138-1 -
Error at Line 9 - mtu 9
2020-11-13T14:35:07.260269-05:00 Sw2-lab14-tt08 port-accessd[3369]:
Event|9301|LOG_ERR|MSTR|1|Failed to apply ClearPass role - asts_aruba_dur_error5_adv-3138-1 -
Invalid input: 9
2020-11-13T14:34:56.194677-05:00 Sw2-lab14-tt08 hpe-config[20143]: Event|6802|LOG_ERR|AMM|-
|Error while copying configs. Error: failed to copy asts-lab13-done to running-config. Error:
Transaction Failed due to an error
```

```
Sw2-lab14-tt08(config)# start-shell
```

```
Sw2-lab14-tt08:~$ sudo cat /tmp/*.xml
```

```
<?xml version="1.0" ?>
```

```
<CPPMAPIResponse xmlns="http://www.arubanetworks.com/cppm/apidefs/arubaroledefs/1.0">
```



```
<CPPMAPIHeader exportTime="2020-11-13 15:51:00 +000" version="1.0"/>
<StatusCode>
    Success
</StatusCode>
<CPPMDownloadableRoles>
    <CPPMDownloadableRole roleName="asts_aruba_dur_error3_adv-3128-3"
version="3">
<![CDATA[asts_aruba_dur_error3_adv-3128-3
class ip any
10 match any any any
exit
port-access policy employee
10 class ip any
exit
port-access role employee
associate policy employee
mtu 9198
vlan access pc
exit]]>    </CPPMDownloadableRole>
    </CPPMDownloadableRoles>
</CPPMAPIResponse>
<?xml version="1.0" ?>
<CPPMAPIResponse xmlns="http://www.arubanetworks.com/cppm/apidefs/arubaroledefs/1.0">
    <CPPMAPIHeader exportTime="2020-11-13 19:34:49 +000" version="1.0"/>
    <StatusCode>
        Success
    </StatusCode>
    <CPPMDownloadableRoles>
        <CPPMDownloadableRole roleName="asts_aruba_dur_error5_adv-3138-1"
version="1">
<![CDATA[asts_aruba_dur_error5_adv-3138-1
class ip any
10 match any any any
exit
port-access policy employee
10 class ip any
exit
port-access role employee
associate policy employee
mtu 9298
vlan access name pc
exit]]>    </CPPMDownloadableRole>
    </CPPMDownloadableRoles>
</CPPMAPIResponse>
```

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

```
Sw2-lab14-tt08(config)# show port-access clients
```

Port Access Clients

Port	MAC Address	Onboarded Method	Status	Role
1/1/4	00:50:56:b1:9b:b7		Fail	

```
Sw2-lab14-tt08(config)# show port-access clients detail
```

Port Access Client Status Details:

Client 00:50:56:b1:9b:b7, asts-error6
=====

Session Details

```
-----
Port          : 1/1/4
Session Time  : 5039s
IPv4 Address  :
IPv6 Address  :
```

Authentication Details

```
-----
Status          : dot1x Authenticated
Auth Precedence : dot1x - Authenticated, mac-auth - Not attempted
```

Authorization Details

```
-----
Role           :
Status         : Invalid
```

>> Event log does not show new events

```
Sw2-lab14-tt08(config)# show event -r -n 10
```

Event logs from current boot

```
>> Debug log
debug portaccess all
debug pacpolicy all
clear debug buffer
do port-access reauthenticate interface 1/1/4
show debug buffer sev info
Sw2-lab14-tt08(config)# show debug buffer sev info
-----
show debug buffer
-----
2020-11-13:14:50:29.761263|port-
accessd|LOG_INFO|MSTR|1|PORTACCESS|PORTACCESS_DOT1X_RADIUS|Dropped 3 log messages in last 152
seconds (most recently, 152 seconds ago) due to excessive rate
2020-11-13:14:50:29.761279|port-
accessd|LOG_INFO|MSTR|1|PORTACCESS|PORTACCESS_DOT1X_RADIUS|Received authentication request
Access-Request(New) from 802.1x to send to group cppm
2020-11-13:14:50:29.784001|port-
accessd|LOG_INFO|MSTR|1|PORTACCESS|PORTACCESS_DOT1X_RADIUS|Validation succeeded for the
packet received from 10.252.1.23:1812 (0) for RADIUS request context 0xffff85e12bc0 client
request id 120 radius id 119 .
2020-11-13:14:50:29.784036|port-
accessd|LOG_INFO|MSTR|1|PORTACCESS|PORTACCESS_DOT1X_RADIUS|Received Access-Challenge response
for the request context 0xffff85e12bc0 client request id 120 radius id 119 sent by 802.1x
2020-11-13:14:50:29.786735|port-
accessd|LOG_INFO|MSTR|1|PORTACCESS|PORTACCESS_DOT1X_RADIUS|Received authentication request
Access-Request(Challenge) from 802.1x to send to group cppm
2020-11-13:14:50:29.793273|port-
accessd|LOG_INFO|MSTR|1|PORTACCESS|PORTACCESS_DOT1X_RADIUS|Validation succeeded for the
packet received from 10.252.1.23:1812 (0) for RADIUS request context 0xffff85e12bc0 client
request id 121 radius id 120 .
2020-11-13:14:50:29.793301|port-
accessd|LOG_INFO|MSTR|1|PORTACCESS|PORTACCESS_DOT1X_RADIUS|Received Access-Challenge response
for the request context 0xffff85e12bc0 client request id 121 radius id 120 sent by 802.1x
2020-11-13:14:50:29.796017|port-
accessd|LOG_INFO|MSTR|1|PORTACCESS|PORTACCESS_DOT1X_RADIUS|Received authentication request
Access-Request(Challenge) from 802.1x to send to group cppm
2020-11-13:14:50:29.824767|port-
accessd|LOG_INFO|MSTR|1|PORTACCESS|PORTACCESS_DOT1X_RADIUS|Dropped 6 log messages in last 0
seconds (most recently, 0 seconds ago) due to excessive rate
2020-11-13:14:50:29.824783|port-
accessd|LOG_INFO|MSTR|1|PORTACCESS|PORTACCESS_DOT1X_RADIUS|Received authentication request
Access-Request(Challenge) from 802.1x to send to group cppm
2020-11-13:14:50:29.882614|port-
accessd|LOG_INFO|MSTR|1|PORTACCESS|PORTACCESS_DOT1X_RADIUS|Dropped 15 log messages in last 0
seconds (most recently, 0 seconds ago) due to excessive rate
2020-11-13:14:50:29.882629|port-
accessd|LOG_INFO|MSTR|1|PORTACCESS|PORTACCESS_DOT1X_RADIUS|Received authentication request
Access-Request(Challenge) from 802.1x to send to group cppm
```

```
2020-11-13:14:50:29.891573|port-
accessd|LOG_INFO|MSTR|1|PORTACCESS|PORTACCESS_SERVICES|logID=23243 Secure client state
changed for client 00:50:56:b1:9b:b7 in port 1/1/4 to Authenticated but failed to apply role.
2020-11-13:14:50:29.891937|port-
accessd|LOG_INFO|MSTR|1|PORTACCESS|PORTACCESS_SERVICES|logID=23243 PortClientProfile current
profile type is invalid 0
```

>> Something is wrong with the role. Debug and show don't reveal the details, now let's look at the RADIUS trace

```
# use TCPDUMP to see the RADIUS port traffic in the session
diagnostics
diag utilities tcpdump host-port 1812
```

>> on the PC4, restart the 'Wired AutoConfig' Service to trigger authentication

```
Notice Vendor-Specific Attribute: Vendor 14823 is Aruba
For Aruba, Vendor Attribute 1 is the Aruba-User-Role
>> RADIUS server returns Aruba-User-Role 'phone' for this user.
```

```
12 14:52:49.744481 ethertype IPv4, IP (tos 0x0, ttl 60, id 36833, offset 0, flags [DF],
proto UDP (17), length 272)
  10.252.1.23.radius > Sw2-lab14-tt08.39069: RADIUS, length: 244
    Access-Accept (2), id: 0x88, Authenticator: 094fe8f3320511155a0e0b342fbc87aa
      Vendor-Specific Attribute (26), length: 13, Value: Vendor: Unknown (14823)
        Vendor Attribute: 1, Length: 5, Value: phone
      Vendor-Specific Attribute (26), length: 58, Value: Vendor: Microsoft (311)
        Vendor Attribute: 17, Length: 50, Value: .C..d...FC;.P..!u`.wH.y.!..
1\..H<..)O.E..8.....
      Vendor-Specific Attribute (26), length: 58, Value: Vendor: Microsoft (311)
        Vendor Attribute: 16, Length: 50, Value:
.}JU.OLS....t1.T.`B, @.Ud.q...\.....M!~.&K....4.
      EAP-Message Attribute (79), length: 6, Value: .k
      Message-Authenticator Attribute (80), length: 18, Value: .b...?...K..\w.l
      User-Name Attribute (1), length: 13, Value: asts-error6
      Class Attribute (25), length: 58, Value: ?03.|.B..n..PD5w..

13 14:52:49.744481 IP (tos 0x0, ttl 60, id 36833, offset 0, flags [DF], proto UDP (17),
length 272)
  10.252.1.23.radius > Sw2-lab14-tt08.39069: RADIUS, length: 244
    Access-Accept (2), id: 0x88, Authenticator: 094fe8f3320511155a0e0b342fbc87aa
      Vendor-Specific Attribute (26), length: 13, Value: Vendor: Unknown (14823)
        Vendor Attribute: 1, Length: 5, Value: phone
      Vendor-Specific Attribute (26), length: 58, Value: Vendor: Microsoft (311)
        Vendor Attribute: 17, Length: 50, Value: .C..d...FC;.P..!u`.wH.y.!..
1\..H<..)O.E..8.....
      Vendor-Specific Attribute (26), length: 58, Value: Vendor: Microsoft (311)
```

```
Vendor Attribute: 16, Length: 50, Value:
.}JU.OLS....t1.T.`.B,.@.Ud.q...\.....M!~.&K....4.
EAP-Message Attribute (79), length: 6, Value: .k
Message-Authenticator Attribute (80), length: 18, Value: .b...?...K..\w.l
User-Name Attribute (1), length: 13, Value: asts-error6
Class Attribute (25), length: 58, Value: ?03.|.B..n..PD5w..
```

```
Sw2-lab14-tt08(config)# show port-access role local
>> no local role 'phone'
```

Solution:

Define local role 'phone' and assign it to VLAN 12

```
Sw2-lab14-tt08(config)# port-access role phone
Sw2-lab14-tt08(config-pa-role)#
Sw2-lab14-tt08(config-pa-role)# vlan access 12
Sw2-lab14-tt08(config-pa-role)# exit
```

```
Sw2-lab14-tt08(config)# do port-access reauthenticate interface 1/1/4
Sw2-lab14-tt08(config)# show port-access clients
```

Port Access Clients

Port	MAC Address	Onboarded Method	Status	Role
1/1/4	00:50:56:b1:9b:b7	dot1x	Success	phone

Trouble Ticket 08.07

Hints

```
show port-access clients
show port-access clients detail
```

Solution

```
Sw2-lab14-tt08(config)# show port-access clients
```

Port Access Clients

Port	MAC Address	Onboarded Method	Status	Role
1/1/4	00:50:56:b1:9b:b7		Fail	RADIUS_737958316

```
Sw2-lab14-tt08(config)# show port-access clients detail
```

Port Access Client Status Details:

Client 00:50:56:b1:9b:b7, asts-error7

=====

Session Details

Port : 1/1/4
Session Time : 5930s
IPv4 Address :
IPv6 Address :

Authentication Details

Status : dot1x Authenticated
Auth Precedence : dot1x - Authenticated, mac-auth - Not attempted

Authorization Details

Role : RADIUS_737958316
Status : Failed, Failed to assign VLAN

Role Information:

Name : RADIUS_737958316

Type : radius

Reauthentication Period :
Authentication Mode :
Session Timeout :
Client Inactivity Timeout :

```
Description :
Gateway Zone :
UBT Gateway Role :
Access VLAN : 7
Native VLAN :
Allowed Trunk VLANs :
Access VLAN Name :
Native VLAN Name :
Allowed Trunk VLAN Names :
MTU :
QOS Trust Mode :
STP Administrative Edge Port :
PoE Priority :
Captive Portal Profile :
Policy :
```

Sw2-lab14-tt08(config)#

```
#####
# solution
Define VLAN 7
Sw2-lab14-tt08(config)# vlan 7
Sw2-lab14-tt08(config-vlan-7)# exit
```

```
Sw2-lab14-tt08(config)# do port-access reauthenticate interface 1/1/4
Sw2-lab14-tt08(config)# show port-access clients
```

Port Access Clients

Port	MAC Address	Onboarded Method	Status	Role
1/1/4	00:50:56:b1:9b:b7	dot1x	Success	RADIUS_737958316

Sw2-lab14-tt08(config)# show port-access clients detail

Port Access Client Status Details:

Client 00:50:56:b1:9b:b7, asts-error7

=====

Session Details

```
Port : 1/1/4
Session Time : 6421s
IPv4 Address :
IPv6 Address :
```

Authentication Details

```
Status : dot1x Authenticated
Auth Precedence : dot1x - Authenticated, mac-auth - Not attempted
```

Authorization Details

```
-----
Role   : RADIUS_737958316
Status : Applied
```

Role Information:

```
Name  : RADIUS_737958316
Type  : radius
```

```
-----
Reauthentication Period      :
Authentication Mode          :
Session Timeout              :
Client Inactivity Timeout    :
Description                  :
Gateway Zone                 :
UBT Gateway Role             :
Access VLAN                  : 7
Native VLAN                  :
Allowed Trunk VLANs          :
Access VLAN Name             :
Native VLAN Name             :
Allowed Trunk VLAN Names     :
MTU                           :
QOS Trust Mode               :
STP Administrative Edge Port :
PoE Priority                  :
Captive Portal Profile       :
Policy                       :
```

```
# optional RADIUS trace:
diagnostics
diag utilities tcpdump host-port 1812
```

```
13 15:11:32.038171 IP (tos 0x0, ttl 60, id 51448, offset 0, flags [DF], proto UDP (17),
length 275)
  10.252.1.23.radius > Sw2-lab14-tt08.39069: RADIUS, length: 247
    Access-Accept (2), id: 0xb5, Authenticator: c5846b3ceda265720e13ce2420edb975
      Tunnel-Type Attribute (64), length: 6, Value: Tag[1] VLAN
      Tunnel-Medium-Type Attribute (65), length: 6, Value: Tag[1] 802
      Tunnel-Private-Group-ID Attribute (81), length: 4, Value: Tag[1] 7
      Vendor-Specific Attribute (26), length: 58, Value: Vendor: Microsoft (311)
        Vendor Attribute: 17, Length: 50, Value:
.3..s...S...0.....0{...dB...b.:.....8...{....7...`
      Vendor-Specific Attribute (26), length: 58, Value: Vendor: Microsoft (311)
        Vendor Attribute: 16, Length: 50, Value: ...f....Q
<.hQ.....z...p..J~.....P.'.....VF....
      EAP-Message Attribute (79), length: 6, Value: .n
      Message-Authenticator Attribute (80), length: 18, Value: k...>.....o.o
```


User-Name Attribute (1), length: 13, Value: asts-error7
Class Attribute (25), length: 58, Value: ?03.|.B..n..PD5w..

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

```
# Customer has noticed reachability issues

# PC3
C:\WINDOWS\system32>ping 10.x.212.201

Pinging 10.x.212.201 with 32 bytes of data:
Reply from 10.x.41.4: TTL expired in transit.
Reply from 10.x.41.4: TTL expired in transit.
Reply from 10.x.41.4: TTL expired in transit.

# Review the route entry on Sw1
Sw1-lab14-tt09(config)# show ip route 10.x.212.201

Displaying ipv4 routes selected for forwarding

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

10.x.212.128/25, vrf default, tag 0
    via 10.x.46.5, [110/25], ospf

>> This reveals that there is an OSPF route that points to the branch

# On Sw2(branch), review the route entry
Sw2-lab14-tt09(config)# show ip route 10.x.212.201 vrf branch

Displaying ipv4 routes selected for forwarding

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

10.x.212.128/25, vrf branch, tag 0
    via 10.x.46.4, [1/0], static

>> customer had a static route left in the branch router configuration. Since this router was
redistributing static routes into OSPF, the Sw1 core received it as an OSPF route.
```

Internal routes have a higher priority than external, redistributed routes, but in this case the static route was using a /25 mask, that is more specific than the internal OSPF route /24.

```
# Remove the rogue static route from the branch
Sw2-lab14-tt09(config)# no ip route 10.x.212.128/25 10.x.46.4 vrf branch
```

```
# test again on PC3
C:\WINDOWS\system32>ping 10.x.212.201
```

```
Pinging 10.x.212.201 with 32 bytes of data:
Reply from 10.x.41.4: Destination host unreachable.
Reply from 10.x.41.4: Destination host unreachable.
```

```
>> Reply from Sw1 that the destination is unreachable.
```

```
# review routes on Sw1
Sw1-lab14-tt09(config)# show ip route 10.x.212.201
```

```
Displaying ipv4 routes selected for forwarding
```

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

```
10.x.212.0/24, vrf default, tag 0
    via 10.x.242.2, [1/0], static
```

```
# Attempt to reach PC4
Sw1-lab14-tt09(config)# ping 10.x.212.201
PING 10.x.212.201 (10.x.212.201) 100(128) bytes of data.
From 10.x.242.4 icmp_seq=1 Destination Host Unreachable
From 10.x.242.4 icmp_seq=2 Destination Host Unreachable
From 10.x.242.4 icmp_seq=3 Destination Host Unreachable
From 10.x.242.4 icmp_seq=4 Destination Host Unreachable
```

```
>> customer has a static route left in the Sw1 router configuration that points to an old
router IP address. Remove the static route so the route can be learned via OSPF.
```

```
Sw1-lab14-tt09(config)# show run | i route
ip route 10.x.212.0/24 10.x.242.2
```

```
Sw1-lab14-tt09(config)# no ip route 10.x.212.0/24 10.x.242.2
```

```
# Ping to the PC4 IP should now be successful from Sw1. It should also work from PC3.
```

```
Sw1-lab14-tt09(config)# ping 10.x.212.201
PING 10.x.212.201 (10.x.212.201) 100(128) bytes of data.
```

```
108 bytes from 10.x.212.201: icmp_seq=1 ttl=127 time=0.720 ms
108 bytes from 10.x.212.201: icmp_seq=2 ttl=127 time=0.669 ms
108 bytes from 10.x.212.201: icmp_seq=3 ttl=127 time=0.734 ms
108 bytes from 10.x.212.201: icmp_seq=4 ttl=127 time=0.750 ms
108 bytes from 10.x.212.201: icmp_seq=5 ttl=127 time=0.592 ms
```

```
--- 10.x.212.201 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 3998ms
rtt min/avg/max/mdev = 0.592/0.693/0.750/0.057 ms
Sw1-lab14-tt09(config)#
```

Multicast traffic

>> While ping works, the multicast traffic does not seem to arrive on PC3.

Review on the PIM router connected to PC3 (SVI11) if join was received (this validates the IGMP as well)

```
Agg1-lab14-tt09(config-if-vlan)# show ip pim pending
```

Join Pending

VRF : default

Group 239.1.1.100

(*,G) Pending

Incoming Interface: vlan11

Group 239.255.255.250

(*,G) Pending

Incoming Interface: vlan11

>> Agg1 has received the join request > this is good

check on Sw1

```
Sw1-lab14-tt09(config)# show ip pim pending
```

Join Pending

VRF : default

Group 239.1.1.100

(*,G) Pending

Incoming Interface: vlan41

Group 239.255.255.250

(*,G) Pending

Incoming Interface: vlan41

Group 239.1.1.100

(*,G) Pending

Incoming Interface: vlan42

Group 239.255.255.250

(*,G) Pending

Incoming Interface: vlan42

Group 239.255.255.250

```
(*,G) Pending
    Incoming Interface:  vlan242

>> Sw1 has received join request on SVI41, connected to Agg1 > this is good

#### Check on Building2 Agg routers

Agg1-lab14-tt09(config-if-vlan)#  show ip pim pending vrf b2

Join Pending

VRF : b2
  Group 239.255.255.250
    (*,G) Pending
      Incoming Interface:  vlan241
  Group 239.255.255.250
    (*,G) Pending
      Incoming Interface:  vlan212

Agg1-lab14-tt09(config-if-vlan)#  show ip pim pending vrf b2 vsx

Join Pending

VRF : b2
  Group 239.255.255.250
    (*,G) Pending
      Incoming Interface:  vlan212

>>> No join pending for the 239.1.1.100 ! >> Issue between Sw1 and Building2 routers

#### On Sw1, check the next hop router for the PC4 IP address
Sw1-lab14-tt09(config)# show ip route 10.x.212.201

Displaying ipv4 routes selected for forwarding

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

10.x.212.0/24, vrf default, tag 0
    via 10.x.242.3, [110/110], ospf

#### Compare this with the PIM neighbor IP address list
Sw1-lab14-tt09(config)# show ip pim neighbor brief

-----
VRF : default
-----
Total number of neighbors : 3
-----
```

Interface Time (HH:MM:SS)	Neighbor Secondary Address (IPV4)	Uptime Address (HH:MM:SS) (IPV4)	Expires (HH:MM:SS)	DR Priority	Hold
-----	-----	-----	-----	-----	-----
vlan241 00:01:45	10.x.241.2 Null	00:37:45	00:01:33	1	
vlan41 00:01:45	10.x.41.2 Null	00:37:45	00:01:32	1	
vlan42 00:01:45	10.x.42.3 Null	00:37:46	00:01:33	1	

Sw1-lab14-tt09(config)#

>> the next hop for the route is using SVI212, but there is no PIM neighbor on this interface

Compare OSPF interfaces with PIM interfaces
Sw1-lab14-tt09(config)# **show ip ospf interface brief**
OSPF Process ID 1 VRF default
=====

Total Number of Interfaces: 7

Interface	Area	IP Address/Mask	Cost	State	Status
-----	-----	-----	-----	-----	-----
loopback0	0.0.0.0	10.x.0.4/32	0	Loopback	up
vlan241	0.0.0.0	10.x.241.4/24	100	Point-to-point	up
vlan242	0.0.0.0	10.x.242.4/24	10	Point-to-point	up
vlan41	0.0.0.0	10.x.41.4/24	100	Point-to-point	up
vlan42	0.0.0.0	10.x.42.4/24	100	Point-to-point	up
vlan46	0.0.0.0	10.x.46.4/24	10	Point-to-point	up
vlan47	0.0.0.0	10.x.47.4/24	100	Point-to-point	up

Sw1-lab14-tt09(config)# **show ip pim interface brief**

VRF : default			Total number of interfaces : 4			
Interface	IP Address	DR Address	Neighbor count	Mode	VSX Role	
-----	-----	-----	-----	-----	-----	-----
loopback0	10.x.0.4/32	N/A	N/A	sparse	N/A	
vlan241	10.x.241.4/24	10.x.241.4	1	sparse	N/A	
vlan41	10.x.41.4/24	10.x.41.4	1	sparse	N/A	
vlan42	10.x.42.4/24	10.x.42.4	1	sparse	N/A	

>>> SVI242 is missing as a PIM interface

compare SVI241 and SVI242 configuration

Sw1-lab14-tt09(config)# show run int vlan241

```
interface vlan241
  ip address 10.x.241.4/24
  ip ospf 1 area 0.0.0.0
  ip ospf network point-to-point
  ip pim-sparse enable
exit
```

Sw1-lab14-tt09(config)# show run int vlan242

```
interface vlan242
  ip address 10.x.242.4/24
  ip ospf 1 area 0.0.0.0
  ip ospf cost 10
  ip ospf network point-to-point
exit
```

>> pim SM not enabled on SVI242.

>> Since the SVI 242 has a lower OSPF cost, this route was the preferred one in the routing table.

Adjust the configuration

```
interface vlan 242
  ip pim-sparse enable
```

>> Verify multicast traffic is received now >> no multicast received yet

>> Repeat the review of the pim pending. This was ok on Agg1/Agg2/Sw1. Now verify on Agg1-Building2

Agg1-lab14-tt09(config)# show ip pim pending vrf b2

Join Pending

VRF : b2

Group 239.1.1.100

(*,G) Pending

Incoming Interface: vlan241

Group 239.255.255.250

(*,G) Pending

Incoming Interface: vlan241

Group 239.255.255.250

(*,G) Pending

Incoming Interface: vlan212

>> This confirms that the request was received

review the multicast routing table

Agg1-lab14-tt09(config)# show ip mroute brief vrf b2

VRF : b2				Total number of entries : 3
Group Address	Source Address	Neighbor	Interface	Uptime
State				(HH:MM:SS)
239.1.1.100	10.x.212.201	10.x.212.3	vlan212	00:01:58
route				

Check the details of the multicast routing entry

Agg1-lab14-tt09(config)# show ip mroute 239.1.1.100 vrf b2

IP Multicast Route Entries

VRF : b2

Total number of entries : 1

```

Group Address          : 239.1.1.100
Source Address         : 10.x.212.201
Neighbor               : 10.x.212.3
Incoming interface     : vlan212
Multicast Routing Protocol : PIM-SM
Unicast Routing Protocol  : connected
Metric                 : 0
Metric Pref             : 0
Uptime (HH:MM:SS)      : 00:04:39
Downstream Interface
Interface              State      VSX Role
-----
vlan241                forwarding DR
  
```

>> Agg1(Building2) has processed the join and is routing traffic upstream on SVI241.

>> Next step is to check the Sw1 status again.

Sw1-lab14-tt09(config-if-vlan)# show ip pim pending

Join Pending

VRF : default


```

Group 239.1.1.100
  (*,G) Pending
    Incoming Interface:  vlan41
Group 239.255.255.250
  (*,G) Pending
    Incoming Interface:  vlan41
Group 239.1.1.100
  (*,G) Pending
    Incoming Interface:  vlan42
Group 239.255.255.250
  (*,G) Pending
    Incoming Interface:  vlan42
Group 239.255.255.250
  (*,G) Pending
    Incoming Interface:  vlan242
Sw1-lab14-tt09(config-if-vlan)# show ip mroute brief
Sw1-lab14-tt09(config-if-vlan)#

>> This reveals that Sw1 is still in the joining state, but has not received any dataplane
traffic.

# Review the SVI241 configuration
Sw1-lab14-tt09(config-if-vlan)# show run int vlan241
interface vlan241
  apply access-list ip core routed-in
  ip address 10.x.241.4/24
  ip ospf 1 area 0.0.0.0
  ip ospf network point-to-point
  ip pim-sparse enable
  exit

# Review the ACL
Sw1-lab14-tt09(config-if-vlan)# show access-list interface vlan241 routed-in commands
access-list ip core
  10 permit ospf any any
  20 permit pim any any
  30 permit any any 10.0.0.0/255.0.0.0
interface vlan241
  apply access-list ip core routed-in

>> The customer had correctly configured the ACL to support PIM and OSPF, but they forgot to
add the actual multicast traffic

# Adjust the ACL on Sw1
access-list ip core
  40 permit any any 224.0.0.0/4

>> PC3 should now receive multicast traffic

```

This concludes the trouble tickets activity.

Aruba Training-Confidential

Aruba Training-Confidential

3333 Scott Blvd, Santa Clara, CA 95054 TEL:
408.227.4500 | FAX: 408.227.4550
www.ARUBANETWORKS.com

EDU-ASTS-Solutions-v20.41