LAB GUIDE

a Hewlett Packard Enterprise company

Deploying OSPFv2 Areas

!!IMPORTANT!!

THIS GUIDE ASSUMES THAT THE AOS-CX OVA HAS BEEN INSTALLED AND WORKS IN GNS3 OR EVE-NG. PLEASE REFER TO GNS3/EVE-NG INITIAL SETUP LABS IF REQUIRED.

AT THIS TIME, EVE-NG DOES NOT SUPPORT EXPORTING/IMPORTING AOS-CX STARTUP-CONFIG. THE LAB USER SHOULD COPY/PASTE THE AOS-CX NODE CONFIGURATION FROM THE LAB GUIDE AS DESCRIBED IN THE LAB GUIDE IF REQUIRED.

TABLE OF CONTENTS

Lab Objective

The OSPF (Open Shortest Path Protocol) is one of the most popular routing protocols for IP Networks. It uses a link state routing (LSR) algorithm which is performed by every switch router mode in the network. OSPF leverages areas and it is these area concepts that form the basis of the LAB which introduces the 'Backbone area', 'Regular Areas' including Stub areas and not so stubby areas.

This lab should be considered as a basic OSPFv2 lab as an introduction to the configuration and operation of OSPF on Aruba CX switches.

At the end of this workshop you will be able to understand and configure ospf areas, understand basic ospf metric calculations of routes, simple route redistributions and the use of stub areas and NSSAs (Not so stubby areas).

Lab Overview

The lab comprises of two Autonomous systems presented as AS1 and AS2. AS1 comprises of two areas , Area 0 & 1 with AS2 redistributing into AS1 and vice-versa.

AS – Autonomous Systems

The two AS systems in this lab are discreet/separate routing systems each running its own LSR (Link State Routing) algorithm for each router node to build a topology map of all available data paths in the network. The data is saved on each router in database which is also referred to as a Link-State Database (LSDB).

Routing information Is not shared between OSPF Autonomous Systems unless explicitly configured with route redistribution for each AS. This activity is covered in the lab between Switch C and Switch E where switch C is configured as an ASBR and redistributes route between AS1 and AS2. .(An ASBR is an Autonomous System Boundary Router)

Area0 backbone –

OSPF area 0 or backbone area is typically designed as a high-speed transit area for router traffic and is at the core of an OSPF network. All other areas are connected to it and inter area traffic must traverse the backbone area. (If a single area only is deployed there Is no requirement to have an area 0)

The lab has two Area backbones or Area 0 networks, One for AS1 & 1 for AS2

OSPF areas (Not Area 0-)

OSPF areas that are not the backbone are numbered other than 0 and are often referred to as ‘Regular Areas’ if they are not configured as a ‘Stub Area’ or ‘Not So Stubby Area ‘ (NSSA)..

In this lab, AS1 has Area 0 and Area 1 connecting to it via Switch B which performs the function of an ABR (Area Border Router).

The initial build of the lab for AS1 involves Area 0 and Area 1 as a regular area. Switch B and Switch C in Area 1 are re-configured from a regular OSPF area to a ‘Stub’ area and then as a NSSA in subsequent lab tasks.

- A ‘Stub’ area is an area where there are no routers or areas beyond it and it does not advertise external routes (external link advertisements LSA Type 5).
- A NSSA accepts external routes (in the form of external link advertisements LSA Type 7) and it is useful sometimes to import external routes from one AS to another whilst still keeping some benefits of a stub area.

Lab Network Layout

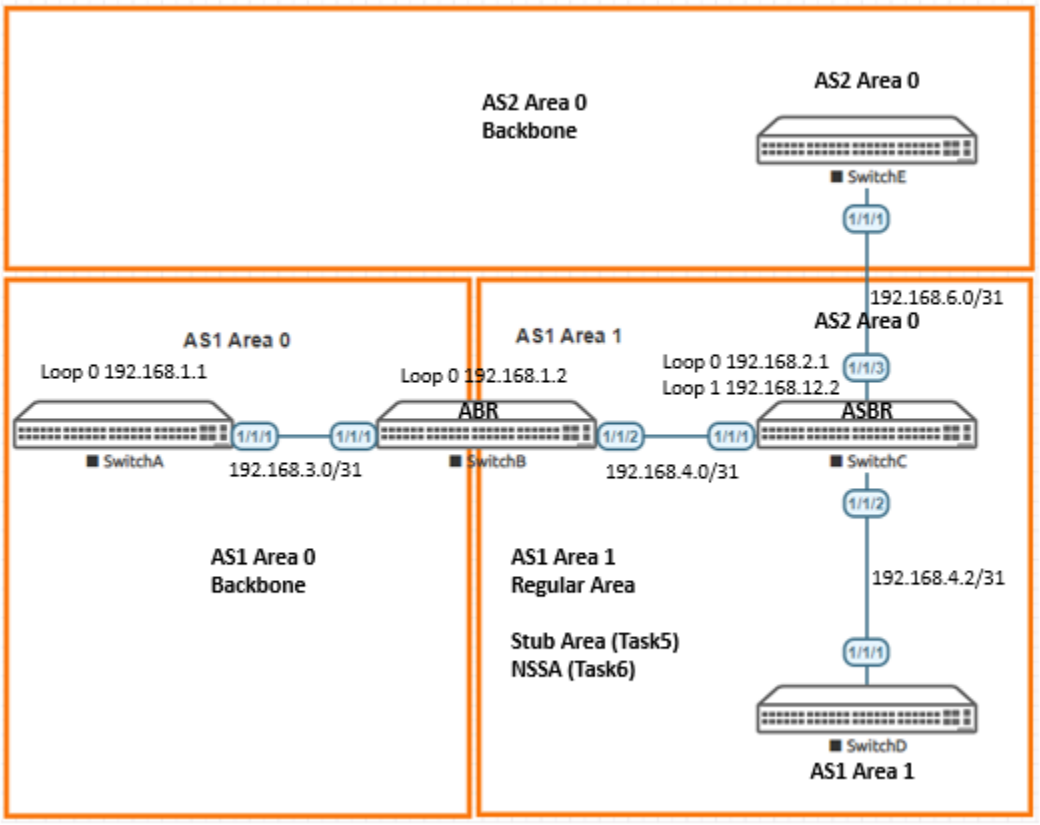

Figure 1 OSPF Area and IP addressing

Lab Tasks

Task 1 Lab Set-up

For this lab refer to Figure 1 for topology and IP address details.

- Start all the devices, including host and client
- Open each switch console and log in with user “admin” and no password
- Change all hostnames as shown in the topology:

```
hostname …
```

- On all devices, bring up required ports:

```
int 1/1/1-1/1/3
  no shutdown
```

- Validate LLDP neighbors appear as expected on each switch

```
show lldp neighbor
```

Task 2–Configure loopback 0 interfaces on Switch A-E

Configure loopback addressing on loopback 0 on each switch

Loopback0 ip addressing

Switch A ip address 192.168.1.1
Switch B ip address 192.168.1.2
Switch C ip address 192.168.2.1
Switch D ip address 192.168.2.2
Switch E ip address 192.168.12.1

Example Switch B

```
SwitchB# conf t

SwitchB(config)# interface loopback 0
SwitchB(config-loopback-if)# ip address 192.168.1.2/32
```

End of Task2

Task 3 - Configure OSPF Area 0 and Area 1 for Switches A, B, C D

The following tasks will be completed in task3 to configure OSPF on switches A,B, C.& D

On each switch A, B,C, D

- Configure a OSPF routing process with appropriate areas and assign a router-id which will be ‘loopback0’
- Configure appropriate switch interfaces with OSPF enabled and ensure connectivity is established
- Ensure neighbor adjacencies are formed between each switch rtr
- Review inter-area and intra-area routes in the ospf routing table
- Review the OSPF Cost of specific routes (Switch A)

Task 3.1 Configure OSPF routing

- Configure OSPF routing on Switch A, B, C & D and assign a router-id with loopback 0
- Configure IP ospf interfaces

SwitchA area 0

```
router ospf 1
    router-id 192.168.1.1
    area 0.0.0.0

interface 1/1/1
    ip address 192.168.3.0/31
    ip ospf 1 area 0.0.0.0
    ip ospf network point-to-point


interface loopback 0
     ip ospf 1 area 0.0.0.0
```

SwitchB area 0 Area 1 – ABR router

```
router ospf 1
    router-id 192.168.1.2
    area 0.0.0.0
    area 0.0.0.1
interface 1/1/1
    ip address 192.168.3.1/31
    ip ospf 1 area 0.0.0.0
    ip ospf network point-to-point

interface 1/1/2
     ip address 192.168.4.0/31
     ip ospf 1 area 0.0.0.1
     ip ospf network point-to-point


interface loopback 0
         ip ospf 1 area 0.0.0.0
```

SwitchC

```
router ospf 1
     router-id 192.168.2.1
```

```
    area 0.0.0.1

interface 1/1/1
    ip address 192.168.4.1/31
    ip ospf 1 area 0.0.0.1
    ip ospf network point-to-point


interface 1/1/2
    ip address 192.168.4.2/31
    ip ospf 1 area 0.0.0.1
    ip ospf network point-to-point


interface loopback 0
    ip ospf 1 area 0.0.0.1
```

SwitchD

```
router ospf 1
    router-id 192.168.2.2
    area 0.0.0.1

interface 1/1/1
    ip address 192.168.4.3/31
    ip ospf 1 area 0.0.0.1
    ip ospf network point-to-point


interface loopback 0
    ip ospf 1 area 0.0.0.1
```

Task 3.2 Validate connectivity -Check ospf neighbor adjacencies are formed

On all switches confirm ospf neighbor adjacencies are formed

On all switches

```
sh ip ospf neighbors
```

Example output

Switch A OSPF neighbor(s)

```
SwitchA# sh ip ospf neighbors
```

```
OSPF Process ID 1 VRF default
==============================

Total Number of Neighbors: 1

Neighbor ID      Priority  State            Nbr Address       Interface
-------------------------------------------------------------------------
192.168.1.2      n/a       FULL             192.168.3.1       1/1/1
```

Switch B OSPF neighbors

```
. SwitchB# sh ip ospf neighbors
OSPF Process ID 1 VRF default
==============================

Total Number of Neighbors: 2

Neighbor ID      Priority  State            Nbr Address       Interface
-------------------------------------------------------------------------
192.168.1.1      n/a       FULL             192.168.3.0        1/1/1

192.168.2.1      n/a       FULL             192.168.4.1        1/1/2
```

Repeat for Switch C & D.

SwitchC will have neighbor adjacencies with Switch B &D

SwitchE will be configured in subsequent tasks.

Task 3.3 Review Routing tables on switches

Review ospf routing table output on sample switches in area 0 and Area 1 and note the intra-area and inter area routes presented.

On selected switches use the command:-

```
sh ip ospf route
```

Sample output

Switch A

```
SwitchA# sh ip ospf route
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 : 6

192.168.1.2/32     (i) area: 0.0.0.0
     via 192.168.3.1 interface 1/1/1, cost 100 distance 110
192.168.2.1/32     (I)
     via 192.168.3.1 interface 1/1/1, cost 200 distance 110
192.168.2.2/32     (I)
     via 192.168.3.1 interface 1/1/1, cost 300 distance 110
192.168.3.0/31     (i) area: 0.0.0.0
     directly attached to interface 1/1/1, cost 100 distance 110
192.168.4.0/31     (I)
     via 192.168.3.1 interface 1/1/1, cost 200 distance 110
192.168.4.2/31     (I)
     via 192.168.3.1 interface 1/1/1, cost 300 distance 110
```

Each switch has a loopback 0 ip address configured with an appropriate ospf area configuration. The loopback address are

injected into the ospf routing table, advertised and presented as a 'reachable' subnet on each switch receiving the ospf updates.

Note the intra-area and inter area-routes from Switch A

Intra-area routes refer to updates (routing) that are passed between ospf routers within the same area and do not need to traverse the backbone (Area 0).

Inter-area routes refer to updates that are passed between areas and required to traverse Area 0

External routes refer to updates passed from another routing protocol into the OSPF domain using an Autonomous System Border Router. An example of external routes will be configured in subsequent steps

Switch B

Switch B is an Area Border Router with area 0 & 1 configured. .

Output extracted from 'sh ip ospf route'

```
Total Number of Routes : 6

192.168.1.1/32     (i) area: 0.0.0.0
     via 192.168.3.0 interface 1/1/1, cost 100 distance 110
192.168.2.1/32     (i) area: 0.0.0.1
     via 192.168.4.1 interface 1/1/2, cost 100 distance 110
192.168.2.2/32     (i) area: 0.0.0.1
     via 192.168.4.1 interface 1/1/2, cost 200 distance 110
192.168.3.0/31     (i) area: 0.0.0.0
     directly attached to interface 1/1/1, cost 100 distance 110
192.168.4.0/31     (i) area: 0.0.0.1
     directly attached to interface 1/1/2, cost 100 distance 110
192.168.4.2/31     (i) area: 0.0.0.1
     via 192.168.4.1 interface 1/1/2, cost 200 distance 110
```

As Switch B has interfaces in area 0 & area 1 configured, all routes are learnt as intra-area routes.

Task 3.4 Review OSPF Path and link Costs

On **Switch A** routing output, note the OSPF costs between Switch A and Switch D. Use the loopback 0 address of 192.168.2.2 on switch D in the switch A routing table as a metric reference.

```
192.168.2.2/32      (I)
     via 192.168.3.1 interface 1/1/1, cost 300 distance 110
```

Route to 192.168.2.2/32 will be presented as a cost of '300' from the output in Switch A's route table. OSPF

OSPF routing metric OSPF uses following formula to calculate the cost

Cost = Reference bandwidth / Interface bandwidth in bps.

Reference bandwidth was defined as arbitrary value in OSPF documentation (RFC 2338). Vendors need to use their own reference bandwidth. Aruba uses the 100 Mbs value as a reference bandwidth (100000000 bps).

Switch A Interface Speed for interface 1/1/1

Run '`sh interface brief`' or '`sh interface 1/1/1`' from the CLI to find the default interface speed.

Example shown with '`sh interface brief`' command.

```
SwitchA# sh interface brief
--------------------------------------------------------------------------------------------------------
-----
Port      Native  Mode   Type          Enabled Status  Reason               Speed   Description
          VLAN                                                              (Mb/s)
--------------------------------------------------------------------------------------------------------
-----
1/1/1     --      routed --            yes     up                           1000    --
```

Using the bandwidth formula we have 100,000/1000 (Reference bandwidth in mbps/interface bandwidth in mbps) = a link cost of 100

As we have standard default settings and common link costs across our lab network, we can ascertain that the route 192.168.2.2/32 has traversed x 3 links to reach Switch A from Switch D.

In a 'live' network, interface speeds will vary and may not be consistent which will impact the overall bandwidth cost of any given route.

Cost calculation using a reference speed of 100 Mbps

Interface Speed	Link Cost
25 Gbit/s	1
10 Gbit/s	1
5 Gbit/s	2
1 Gbit/s	10
1000 Mbit/s	100

Note: For VLAN interfaces, the default interface speed is taken as 1 Gbit/s

Confirm the ip ospf default link cost on switch A interface 1/1/1

```
SwitchA# sh ip ospf interface 1/1/1
Interface 1/1/1 is up, line protocol is up
-----------------------------------------

IP address 192.168.3.0/31, Process ID 1 VRF default, area 0.0.0.0
    State Point-to-point, Status up, Network type Point-to-point
    Link Speed: 1000 Mbps
    Cost Configured NA, Calculated 100
    Transit delay 1 sec, Router priority n/a
    No designated router on this network
    No backup designated router on this network
    Timer Intervals: Hello 10, Dead 40, Retransmit 5
    No authentication
    Number of Link LSAs: 0, checksum sum 0
    BFD is disabled
```

The default reference speed can be changed in the respective ospf process configuration using the reference-bandwidth command.

```
SwitchA# SwitchA(config)# router ospf 1
SwitchA(config-ospf-1)# reference-bandwidth ?
  <1-4000000>  Set reference bandwidth in Mbps. (Default: 100000Mbps)
```

The default interface costs can be changed for each interface (or interface VLAN) by using the ip ospf cost command:-

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

```
SwitchA(config-if)# ip ospf cost ?
  <1-65535>  Set interface cost
```

The `no ip ospf cost` command resets the cost value back to the default

End of Task 3

Task 4 Create different OSPF Autonomous Systems (AS) and redistribute routes

Importing routes and redistributing into OSPF is supported by creating an ASBR, an Autonomous System Boundary Router.

In this task you will create:-

- a separate OSPF routing process: on Switch C – (process 2) in area 0
- A routing ospf process in Switch D in area 0
- Route redistribute ospf routes (from ospf process 2) into ospf process 1 on switch C
- Route redistribute ospf routes (from ospf process 1) into ospf process 2 on switch C
- Review ospf redistributed route metrics

-

Task 4.1 Configure ospf routing between Switch C & Switch D

Switch C

From the configuration context, create an additional loopback address for the router-id for ospf process 2

```
interface loopback 1
    ip address 192.168.12.2/32
```

Create an additional router ospf process

```
router ospf 2
    router-id 192.168.12.2
    area 0.0.0.0
```

add interface loopback 1 in to ospf process 2

```
interface loopback 1
ip ospf 2 area 0.0.0.0
```

Configure OSPF on interface 1/1/3 to Switch E

```
interface 1/1/3
    ip address 192.168.6.0/31
```

```
    ip ospf 2 area 0.0.0.0
    ip ospf network point-to-point
```

Switch E

From the configuration context, create the ospf routing process

```
router ospf 1
    router-id 192.168.12.1
    area 0.0.0.0
```

add interface loopback 0 in to ospf process 1

```
ip ospf 1 area 0.0.0.0
```

Configure interface 1/1/1

```
interface 1/1/1
    ip address 192.168.6.1/31
    ip ospf 1 area 0.0.0.0
    ip ospf network point-to-point
```

Task 4.1 Validate ospf neighbors and routes

Validate neighbor adjacency has been formed between Switch C and Switch D

```
show ip ospf neighbors -
```

Sample output Switch E

```
SwitchE# sh ip ospf neighbors
OSPF Process ID 1 VRF default
==============================

Total Number of Neighbors: 1

Neighbor ID      Priority  State             Nbr Address       Interface
-------------------------------------------------------------------------
192.168.12.2     n/a       FULL              192.168.6.0        1/1/1
```

Review ospf routing table

On Switches B, C,& D

```
sh ip ospf routes
```

Switch C sample output

Note that Switch C now has output for 2 ospf process IDs

```
SwitchC# sh ip ospf route
Codes: i - Intra-area route, I - Inter-area route
       E1 - External type-1, E2 - External type-2

OSPF Process ID 1 VRF default, Routing Table
-------------------------------------------

Total Number of Routes : 5

192.168.1.1/32     (I)
     via 192.168.4.0 interface 1/1/1, cost 200 distance 110
192.168.1.2/32     (I)
     via 192.168.4.0 interface 1/1/1, cost 100 distance 110
192.168.3.0/31     (I)
     via 192.168.4.0 interface 1/1/1, cost 200 distance 110
192.168.4.0/31     (i) area: 0.0.0.1
     directly attached to interface 1/1/1, cost 100 distance 110
192.168.4.2/31     (i) area: 0.0.0.1
     directly attached to interface 1/1/2, cost 100 distance 110

OSPF Process ID 2 VRF default, Routing Table
-------------------------------------------

Total Number of Routes : 2

192.168.6.0/31     (i) area: 0.0.0.0
     directly attached to interface 1/1/3, cost 100 distance 110
192.168.12.1/32    (i) area: 0.0.0.0
```

```
via 192.168.6.1 interface 1/1/3, cost 100 distance 110
```

- On switch B , the ospf route table will not include routes learnt from Switch C ospf process ID 2 as these routes are learnt within a different Autonomous System.

- On Switch E ,the ospf route table will not include routes from ospf process id 1 as they are again routes learnt within a different Autonomous System.

Task 4.2 Create an ASBR with route redistribute commands

To include routes from different AS (Autonomous Systems) so they propagate within our routed lab network we need to redistribute routes on Switch C and by doing so we make Switch C an ASBR: an Autonomous System Boundary Router.

This is a 2-step process:-

1. Redistribute routes from ospf process 2 into ospf process 1
2. Redistribute routes from ospf process 1 into ospf process 2

On Switch C

- First, we route redistribute ospf routes (from ospf process 2) into ospf process 1 on switch C

Within the '`router ospf 1`' context add the following commands#

```
redistribute ospf 2
```

- ospf learned routes from ospf process 2 will be redistributed into ospf process 1

Seconds step , we repeat the process for ospf process 2, we route redistribute ospf routes (from ospf process 1) into ospf process 2 on switch C

Within the '`router ospf 2`' context add the following commands#

```
redistribute ospf 1
```

Task 4.3 Validate neighbors and route redistribution on Switch B & Switch E

On switch C & E, run the 'sh ip ospf neighbors' command.

Sample switch C

Sample Output Switch C – note the process id split on neighbors

```
SwitchC# sh ip ospf neighbors
OSPF Process ID 1 VRF default
==============================

Total Number of Neighbors: 2

Neighbor ID      Priority  State            Nbr Address       Interface
-------------------------------------------------------------------------
192.168.1.2      n/a       FULL             192.168.4.0       1/1/1

192.168.2.2      n/a       FULL             192.168.4.3       1/1/2

OSPF Process ID 2 VRF default
==============================

Total Number of Neighbors: 1

Neighbor ID      Priority  State            Nbr Address       Interface
-------------------------------------------------------------------------
192.168.12.1     n/a       FULL             192.168.6.1       1/1/3
```

On switch B and D run the ‘`sh ip ospf route`’ command and note the output

Switch B output

```
SwitchB# sh ip ospf route
Codes: i - Intra-area route, I - Inter-area route
       E1 - External type-1, E2 - External type-2

OSPF Process ID 1 VRF default, Routing Table
--------------------------------------------
```

```
Total Number of Routes : 8

192.168.1.1/32     (i) area: 0.0.0.0
     via 192.168.3.0 interface 1/1/1, cost 100 distance 110
192.168.2.1/32     (i) area: 0.0.0.1
     via 192.168.4.1 interface 1/1/2, cost 100 distance 110
192.168.2.2/32     (i) area: 0.0.0.1
     via 192.168.4.1 interface 1/1/2, cost 200 distance 110
192.168.3.0/31     (i) area: 0.0.0.0
     directly attached to interface 1/1/1, cost 100 distance 110
192.168.4.0/31     (i) area: 0.0.0.1
     directly attached to interface 1/1/2, cost 100 distance 110
192.168.4.2/31     (i) area: 0.0.0.1
     via 192.168.4.1 interface 1/1/2, cost 200 distance 110
192.168.6.0/31     (E2)
     via 192.168.4.1 interface 1/1/2, cost 100 distance 110
192.168.12.1/32    (E2)
     via 192.168.4.1 interface 1/1/2, cost 100 distance 110
```

Switch E output

```
SwitchE(config)# sh ip ospf route
Codes: i - Intra-area route, I - Inter-area route
       E1 - External type-1, E2 - External type-2

OSPF Process ID 1 VRF default, Routing Table
-------------------------------------------

Total Number of Routes : 8

192.168.1.1/32     (E2)
     via 192.168.6.0 interface 1/1/1, cost 200 distance 110
192.168.1.2/32     (E2)
     via 192.168.6.0 interface 1/1/1, cost 100 distance 110
192.168.2.2/32     (E2)
     via 192.168.6.0 interface 1/1/1, cost 100 distance 110
```

```
192.168.3.0/31     (E2)
     via 192.168.6.0 interface 1/1/1, cost 200 distance 110
192.168.4.0/31     (E2)
     via 192.168.6.0 interface 1/1/1, cost 100 distance 110
192.168.4.2/31     (E2)
     via 192.168.6.0 interface 1/1/1, cost 100 distance 110
192.168.6.0/31     (i) area: 0.0.0.0
     directly attached to interface 1/1/1, cost 100 distance 110
192.168.12.2/32    (i) area: 0.0.0.0
     via 192.168.6.0 interface 1/1/1, cost 100 distance 110
```

The redistributed routes (from another AS) are tagged as a Type 5 LSA routes and are identified as an external router with the E prefix.

E1 routes is the cost of the external metric and the additional internal cost within OSPF to reach that network.

- E1 route(s) includes the internal cost to the ASBR which is added to the external cost of the route

The cost of E2 routes is always the external metric value of the route and the internal cost to/from the ASBR is ignored.

- E2 route(s) do not include the internal cost of the . They will always have the same external cost.

Routes 192.168.1.1/32 & 192.168.3.0/31 via Switch A have traversed 1x ABR(Switch) and 1 x ASBR (SwitchC) which collectively provides the accumulated metric of ‘200’ on receipt at Switch E.

End of Task 4

Task 5 Stub Area

This task will create a stub area between Switch B and Switch C for area 1. Switch B still operates as an ABR but the neighbor relationship in area 1 is changed to ‘Stub’ for switch C.

Switch D and Switch E are not required for this task

‘shutdown’ interface 1/1/2 & 1/1/3 on Switch C

.

.**Switch B**

On switch B , the area 0.0.0.1 needs to be amended to include ‘stub’

From within ‘router ospf 1’ config context

```
area 0.0.0.1 stub
```

.**Switch C**

On switch C , the area 0.0.0.1 needs to be amended to include ‘stub’

From within ‘router ospf 1’ config context

```
area 0.0.0.1 stub
```

Check neighbor adjacency has formed with Switch B

```
SwitchC# sh ip ospf neighbors
OSPF Process ID 1 VRF default
==============================

Total Number of Neighbors: 1

Neighbor ID      Priority  State            Nbr Address       Interface
-------------------------------------------------------------------------
192.168.1.2      n/a       FULL             192.168.4.0        1/1/1
```

Display switch C ospf routing table

```
sh ip ospf route
```

You should note a significant change in the ospf route table on switch C. Switch B, as the ABR , now injects a default route to

it's neighbor Switch B , as it is configured as a stub area .

Sample output

Total Number of Routes : 5

```
0.0.0.0/0           (I)
     via 192.168.4.0 interface 1/1/1, cost 101 distance 110
192.168.1.1/32      (I)
     via 192.168.4.0 interface 1/1/1, cost 200 distance 110
192.168.1.2/32      (I)
     via 192.168.4.0 interface 1/1/1, cost 100 distance 110
192.168.3.0/31      (I)
     via 192.168.4.0 interface 1/1/1, cost 200 distance 110
192.168.4.0/31      (i) area: 0.0.0.1
     directly attached to interface 1/1/1, cost 100 distance 110
```

As there is a default route advertised from Switch B as the ABR for area 1 'Stub'

As Switch C has a single ingress and egress point, the route table can be reduce further by eliminating 'Inter Area routes'.

On Switch B within the 'router ospf 1' context

Enter

SwitchB(config)# router ospf 1

SwitchB(config-ospf-1)# **area 0.0.0.1 stub no-summ**

Display Switch C ospf routing table

```
sh ip ospf route

0.0.0.0/0           (I)
     via 192.168.4.0 interface 1/1/1, cost 101 distance 110
192.168.4.0/31      (i) area: 0.0.0.1
     directly attached to interface 1/1/1, cost 100 distance 110
OSPF Process ID 2 VRF default, Routing Table
----------------------------------------------
Total Number of Routes : 0
```

Inter area routes are no longer present in the route table. The ‘no-summary’ disables the summary of LSAs on each router that is connect to the ABR in that area. (In this case the summary routes are the host routes)

Stub areas

.

- Typically have a single ingress egress point to connecting to the ABR
- External networks redistributed from other protocols into ospf are not allowed to be advertised into a stub area. The ABR, in this case switch B, stops LSA types 4 & 5.
- Routing is based on the stub router receiving a default route from the ABR (0.0.0.0)
- All OSPF routers inside a stub area must be configured as a stub router .
- Routers (stub areas) are required to connect to an ABR

For hub and spoke connectivity in large OSPF networks, the Stub area is used extensively as they reduce the amount the of LSAs advertised and processed and thereby reduce the overall size of the routing table and assist in keeping the overall routing protocol convergence times down..

..

End of Task 5

Task 6 NSSA – Not So Stubby Area

A NSSA, Not So Stubby Area, is very similar to a standard stub area but has one major difference. It is less restrictive than a stub area which cannot import external routes. NSSA can import external routes into OSPF from either another OSPF process or another routing protocol.

In this task, area 1 between Switch B and Switch C is configured as a NSSA area and the routes learnt for ospf process 2 (for area 0 between Switch C and Switch E) are redistributed into NSSA area 1..

On Switch B remove the stub area configuration and add the NSSA configuration

From the router ospf 1 config context

```
no area 0.0.0.1 stub
area 0.0.0.1 nssa no-summary
```

on Switch C - ‘No shut’ interface 1/1/3

```
interface 1/1/3
    no shutdown
```

On Switch C remove the stub area configuration and add the NSSA configuration

From the router ospf 1 config context

```
no area 0.0.0.1 stub
area 0.0.0.1 nssa
```

Check that Switch B & C have an ospf neighbor adjacency

```
sh ip ospf neighbor
```

On Switch C , the redistribute commands into process ospf 1 and process opsf2 should still be present.

```
router ospf 1
    router-id 192.168.2.1
    redistribute ospf 2
    area 0.0.0.1 nssa
router ospf 2
    router-id 192.168.12.2
    redistribute ospf 1
    area 0.0.0.0
```

On Switch C , display the ip osp route table

```
SwitchC# sh ip ospf route
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 : 2

0.0.0.0/0          (I)
     via 192.168.4.0 interface 1/1/1, cost 101 distance 110
192.168.4.0/31     (i) area: 0.0.0.1
     directly attached to interface 1/1/1, cost 100 distance 110

OSPF Process ID 2 VRF default, Routing Table
```

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

Total Number of Routes : 2

192.168.6.0/31     (i) area: 0.0.0.0
     directly attached to interface 1/1/3, cost 100 distance 110
192.168.12.1/32    (i) area: 0.0.0.0
     via 192.168.6.1 interface 1/1/3, cost 100 distance 110
```

The default advertised route from the Switch B ABR is the same when the switches were configured for 'Stub'. Switch C is re-advertising routes between ospf processes 1 & 2.

On Switch B , display the ip osp route table

```
SwitchB# sh ip ospf route
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 : 6

192.168.1.1/32     (i) area: 0.0.0.0
     via 192.168.3.0 interface 1/1/1, cost 100 distance 110
192.168.2.1/32     (i) area: 0.0.0.1
     via 192.168.4.1 interface 1/1/2, cost 100 distance 110
192.168.3.0/31     (i) area: 0.0.0.0
     directly attached to interface 1/1/1, cost 100 distance 110
192.168.4.0/31     (i) area: 0.0.0.1
     directly attached to interface 1/1/2, cost 100 distance 110
192.168.6.0/31     (E2)
     via 192.168.4.1 interface 1/1/2, cost 100 distance 110
192.168.12.1/32    (E2)
     via 192.168.4.1 interface 1/1/2, cost 100 distance 110
```

External redistributed routes from Switch C (ospf process 2) are tagged as E2 routes and populated in the route table .

- **End of lab task5 and lab tasks**

Appendix – Complete Configurations

Switch A

```
interface 1/1/1
    no shutdown
    ip address 192.168.3.0/31
    ip ospf 1 area 0.0.0.0
    ip ospf network point-to-point
interface 1/1/2
    no shutdown
interface 1/1/3
    no shutdown
interface loopback 0
    ip address 192.168.1.1/32
    ip ospf 1 area 0.0.0.0
!
!router ospf 1
    router-id 192.168.1.1
    area 0.0.0.0
```

Switch B

```
interface 1/1/1
    no shutdown
    ip address 192.168.3.1/31
    ip ospf 1 area 0.0.0.0
    ip ospf network point-to-point
interface 1/1/2
    no shutdown
    ip address 192.168.4.0/31
    ip ospf 1 area 0.0.0.1
    ip ospf network point-to-point
interface 1/1/3
    no shutdown
interface loopback 0
    ip address 192.168.1.2/32
```

```
    ip ospf 1 area 0.0.0.0
!
router ospf 1
    router-id 192.168.1.2
    area 0.0.0.0
  area 0.0.0.1                    Tasks 2-4 full ospf area 1
  area 0.0.0.1 stub               Task 5 'Stub area' with or without 'no-summary'
  area 0.0.0.1 nssa no-summary    Task 6 'NSSA' with no-summary
```

Switch C

```
interface 1/1/1interface loopback
    ip address 192.168.4.1/31
    ip ospf 1 area 0.0.0.1
    ip ospf network point-to-point
interface 1/1/2
    ip address 192.168.4.2/31
    ip ospf 1 area 0.0.0.1
    ip ospf network point-to-point
interface 1/1/3
    no shutdown
    ip address 192.168.6.0/31
    ip ospf 2 area 0.0.0.0
    ip ospf network point-to-point
interface loopback 0
    ip address 192.168.2.1/32
    ip ospf 1 area 0.0.0.1
interface loopback 1
    ip address 192.168.12.2/32
    ip ospf 2 area 0.0.0.0
!
router ospf 1
    router-id 192.168.2.1
    redistribute ospf 2
    area 0.0.0.1                    Tasks 2-4 full ospf area 1
    area 0.0.0.1 stub               Task 5 'Stub area' with or without 'no-summary'
    area 0.0.0.1 nssa no-summary    Task 6 'NSSA' with no-summary
```

```
router ospf 2
    router-id 192.168.12.2
    redistribute ospf 1
    area 0.0.0.0
```

Switch D

```
interface 1/1/1
    no shutdown
    ip address 192.168.4.3/31
    ip ospf 1 area 0.0.0.1
    ip ospf network point-to-point
interface 1/1/2
    no shutdown
interface 1/1/3
    no shutdown
interface loopback 0
    ip address 192.168.2.2/32
    ip ospf 1 area 0.0.0.1
!
router ospf 1
    router-id 192.168.2.2
    area 0.0.0.1
```

Switch E

```
interface 1/1/1
    no shutdown
    ip address 192.168.6.1/31
    ip ospf 1 area 0.0.0.0
    ip ospf network point-to-point
interface 1/1/2
    no shutdown
interface 1/1/3
    no shutdown
interface loopback 0
    ip address 192.168.12.1/32
```

```
    ip ospf 1 area 0.0.0.0
!
router ospf 1
    router-id 192.168.12.1
    area 0.0.0.0
```

END OF DOCUMENT

a Hewlett Packard Enterprise company

www.arubanetworks.com

3333 Scott Blvd. Santa Clara, CA 95054
1.844.472.2782 | T: 1.408.227.4500 | FAX: 1.408.227.4550 | info@arubanetworks.com