LAB GUIDE

# Deploying basic BGP

a Hewlett Packard Enterprise company

**!!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

At the end of this workshop, you will be able to implement the basic configuration to enable both iBGP and EBGP. The main goal in to ensure Hosts within different Layer 3 autonomous systems (AS) can reach each other.

For stability, a best practice is to peer using loopbacks for iBGP. To achieve this, you will first configure an IGP protocol (OSPF) to enable loopback reachability for iBGP peering. You will also need to ensure the host networks are advertised into the BGP fabric. You will then create an EBGP peering relationship into another AS and ensure the hosts can communicate across ASs.

Finally, once the hosts can communicate across the ASs you will add a layer of security by creating a policy which blocks traffic from a certain host. Once enabled you will see that the hosts cannot communicate, however the switches can still reach the hosts across the ASs.

## Lab Overview

Border Gateway Protocol (BGP) is a standardized exterior gateway protocol designed to exchange routing and reachability information among autonomous systems (AS) both on the Internet and even within and across enterprise organizations. BGP is classified as a path-vector routing protocol, and it makes routing decisions based on paths, network policies, or rule-sets configured by a network administrator.

BGP may be used for routing within an autonomous system. In this application it is referred to as Interior Border Gateway Protocol, Internal BGP, or iBGP. In contrast, when routing between different autonomous systems the protocol may be referred to as Exterior Border Gateway Protocol, External BGP, or eBGP.

The characteristics of BGP are:

- The current version of BGP is BGP version 4, based on RFC4271.
- BGP is the path-vector protocol that provides routing information for autonomous systems on the Internet via its AS-Path attribute.
- BGP is a Layer 4 protocol that sits on top of TCP. It is much simpler than OSPF, because it doesn't have to worry about the things TCP will handle.
- Peers that have been manually configured to exchange routing information will form a TCP connection and begin speaking BGP. There is no discovery in BGP.
- Medium-sized businesses usually get into BGP for the purpose of true multi-homing for their entire network.
- An important aspect of BGP is that the AS-Path itself is an anti-loop mechanism. Routers will not import any routes that contain themselves in the AS-Path.

## Lab Network Layout

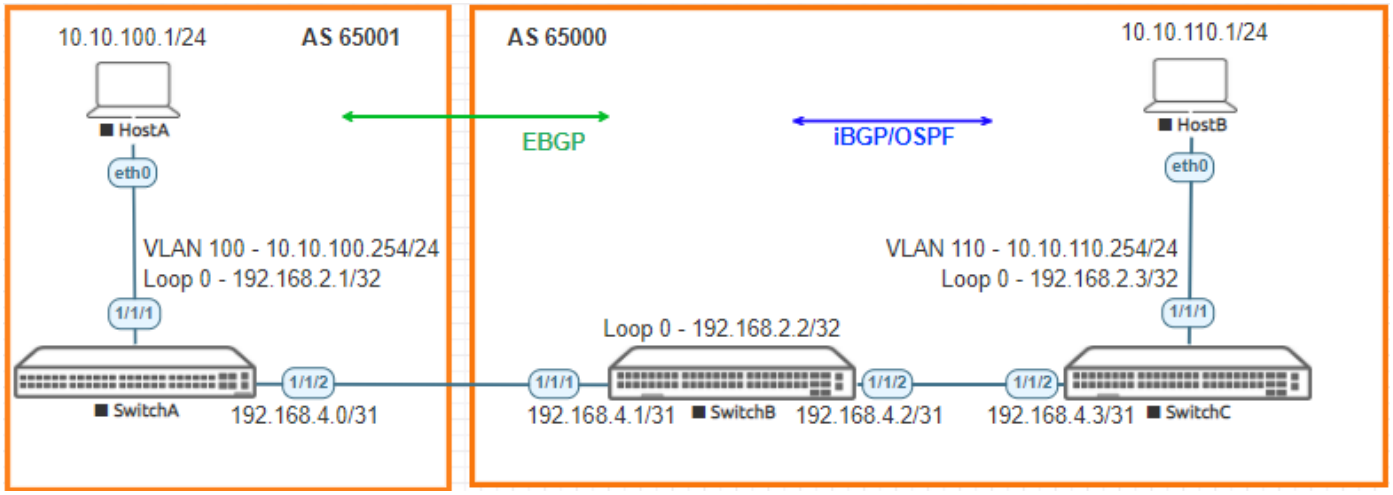


*Figure 1. Lab topology and addresses*

Note: With the OVs in the lab, after adding a route-map you will need to clear BGP from user mode to make the map take effect.

## Lab Tasks

### Task 1 - Lab setup

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 1/1/1-1/1/2
  no shutdown
```

- Validate LLDP neighbors appear as expected

```
show lldp neighbor
```

**SwitchB**

```
SwitchB# show lldp neighbor

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

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

LOCAL-PORT  CHASSIS-ID         PORT-ID          PORT-DESC                TTL      SYS-NAME
-------------------------------------------------------------------------------------------
1/1/1       08:00:09:ee:11:82  1/1/1            1/1/1                    120      SwitchA
1/1/2       08:00:09:16:7b:7e  1/1/2            1/1/2                    120      SwitchC
```

## Task 2 - Configure Host_A and Host_B

- Apply the proper IP address and gateway to both Host_A and Host_B

**HostA**

```
ip 10.10.100.1/24 10.10.100.254
```

**HostB**

```
ip 10.10.110.1/24 10.10.110.254
```

- Verify with show ip

```
show ip
```

**HostA**

```
VPCS> sho ip

NAME        : VPCS[1]
IP/MASK     : 10.10.100.1/24
GATEWAY     : 10.10.100.254
DNS         :
MAC         : 00:50:79:66:68:07
LPORT       : 20000
RHOST:PORT  : 127.0.0.1:30000
MTU         : 1500
```

## Task 3 - Configure interfaces and verify direct connectivity

- Configure switch interfaces and ensure direct connectivity works
- Apply proper IPv4 Addresses to all interfaces, including loopback
- On Switch A and C:
    - Create Host facing VLAN/Interface
    - Apply proper VLAN to host facing access interface
- Ensure direct connectivity works between each link

**SwitchA**

```
vlan 100
    description HostA VLAN
interface 1/1/1
    no shutdown
    no routing
    vlan access 100
```

```
interface 1/1/2
    no shutdown
    ip address 192.168.4.0/31
interface loopback 0
    ip address 192.168.2.1/32
interface vlan 100
    description To Client segment
    ip address 10.10.100.254/24
```

**SwitchB**

```
interface 1/1/1
    no shutdown
    ip address 192.168.4.1/31
interface 1/1/2
    no shutdown
    ip address 192.168.4.2/31
interface loopback 0
    ip address 192.168.2.2/32
```

**SwitchC**

```
vlan 110
    description HostB VLAN
interface 1/1/1
    no shutdown
    no routing
    vlan access 110
interface 1/1/2
    no shutdown
    ip address 192.168.4.3/31
interface loopback 0
    ip address 192.168.2.3/32
interface vlan 110
    description To Client segment
    ip address 10.10.110.254/24
```

**SwitchA**

```
SwitchA(config)# ping 192.168.4.1
PING 192.168.4.1 (192.168.4.1) 100(128) bytes of data.
108 bytes from 192.168.4.1: icmp_seq=1 ttl=64 time=1.78 ms
108 bytes from 192.168.4.1: icmp_seq=2 ttl=64 time=1.18 ms
108 bytes from 192.168.4.1: icmp_seq=3 ttl=64 time=1.70 ms
108 bytes from 192.168.4.1: icmp_seq=4 ttl=64 time=1.71 ms
108 bytes from 192.168.4.1: icmp_seq=5 ttl=64 time=1.89 ms


--- 192.168.4.1 ping statistics ---

5 packets transmitted, 5 received, 0% packet loss, time 4005ms

rtt min/avg/max/mdev = 1.187/1.656/1.893/0.244 ms

SwitchA(config)# ping 10.10.100.1
PING 10.10.100.1 (10.10.100.1) 100(128) bytes of data.
108 bytes from 10.10.100.1: icmp_seq=1 ttl=64 time=11.4 ms
108 bytes from 10.10.100.1: icmp_seq=2 ttl=64 time=1.84 ms
108 bytes from 10.10.100.1: icmp_seq=3 ttl=64 time=0.930 ms
108 bytes from 10.10.100.1: icmp_seq=4 ttl=64 time=0.924 ms
108 bytes from 10.10.100.1: icmp_seq=5 ttl=64 time=0.916 ms

--- 10.10.100.1 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4004ms
rtt min/avg/max/mdev = 0.916/3.212/11.448/4.133 ms
```

## Task 4 - Configure OSPF on Switch_B and C

- Create OSPF process 1 and area 0
- Create an OSPF IPv4 Router-ID (same as Loopback)
- Ensure interfaces are added to OSPF table
- Use OSPF point to point
- Ensure Loopback 0 and switch to switch interfaces on Switch B and C are advertised into OSPF
- Ensure VLAN 7 interface on Switch C is advertised into OSPF
- Verify OSPF peering is up between Switch B and C

**SwitchB**

```
router ospf 1
    router-id 192.168.2.2
    area 0.0.0.0
interface 1/1/2
    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
```

**SwitchC**

```
router ospf 1
    router-id 192.168.2.3
    area 0.0.0.0
interface 1/1/2
    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
interface vlan 7
    ip ospf 1 area 0.0.0.0
    ip ospf network point-to-point
```

**SwitchB**

```
SwitchB(config)# sho ip os neighbors
OSPF Process ID 1 VRF default
==============================

Total Number of Neighbors: 1

Neighbor ID      Priority  State            Nbr Address       Interface
-------------------------------------------------------------------------
192.168.2.3      n/a       FULL             192.168.4.3        1/1/2
```

## Task 5 - Configure IBGP on Switch_B and Switch_C

- To prevent route flapping caused by port state changes, this example uses loopback interfaces to establish IBGP connections.
- Loopback interfaces are virtual interfaces. Therefore, use the update-source loopback command to specify the loopback interface as the source interface for establishing BGP connections.
- Create a BGP router-ID
- Ensure HostB segment is added to the BGP table
- Add the neighbor default originate command to send a default route to SwitchC
- Verify BGP peering is up between SwitchB and C

**SwitchB**
```
router bgp 65000
    bgp router-id 192.168.2.2
    neighbor 192.168.2.3 remote-as 65000
    neighbor 192.168.2.3 update-source loopback 0
    address-family ipv4 unicast
        neighbor 192.168.2.3 activate
        neighbor 192.168.2.3 default-originate
    exit-address-family
```

**SwitchC**
```
router bgp 65000
    bgp router-id 192.168.2.3
    neighbor 192.168.2.2 remote-as 65000
    neighbor 192.168.2.2 update-source loopback 0
    address-family ipv4 unicast
        neighbor 192.168.2.2 activate
        network 10.10.110.0/24
    exit-address-family
```

**SwitchC**
```
SwitchC(config-bgp)# show bgp ipv4 unicast summary
VRF : default
BGP Summary
-----------
 Local AS               : 65000        BGP Router Identifier  : 192.168.2.3
 Peers                  : 1            Log Neighbor Changes   : No
 Cfg. Hold Time         : 180          Cfg. Keep Alive        : 60

 Neighbor        Remote-AS MsgRcvd MsgSent   Up/Down Time State        AdminStatus
 192.168.2.2     65000       5       5       00h:01m:16s  Established   Up
```

## Task 6 - Configure EBGP

- The EBGP peers, SwitchA and SwitchB are located in different ASs and usually belong to different carriers. Typically, their loopback interfaces are not reachable to each other, so directly connected interfaces are used for establishing EBGP sessions.
- Create a BGP router-ID for Switch A
- To enable SwitchC (and HostB) access to network 8.1.1.0/24 which is connected directly to SwitchA, inject network 8.1.1.0/24 in to the BGP routing table of SwitchA.
- To enable SwitchA (and HostA) access to network 7.1.1.0/24 which is connected directly to SwitchC, inject network 7.1.1.0/24 in to the BGP routing table of SwitchC.
- Verify BGP peering is up between Switch A and B
- Verify that HostA can reach HostB

**SwitchA**
```
router bgp 65001
    bgp router-id 192.168.2.1
    neighbor 192.168.4.1 remote-as 65000
    address-family ipv4 unicast
        neighbor 192.168.4.1 activate
        network 10.10.100.0/24
    exit-address-family
```

**SwitchB**
```
router bgp 65000
    neighbor 192.168.4.0 remote-as 65001
    address-family ipv4 unicast
        neighbor 192.168.4.0 activate
```

```
    exit-address-family
```

**SwitchA**

```
SwitchA(config)# show bgp ipv4 unicast summary
VRF : default
BGP Summary
-----------
 Local AS               : 65001        BGP Router Identifier  : 192.168.2.1
 Peers                  : 1            Log Neighbor Changes   : No
 Cfg. Hold Time         : 180          Cfg. Keep Alive        : 60

 Neighbor        Remote-AS MsgRcvd MsgSent   Up/Down Time State       AdminStatus
 192.168.4.1     65000       4       5       00h:00m:34s  Established   Up
```

**HostA**

```
VPCS> ping 10.10.110.1
84 bytes from 10.10.110.1 icmp_seq=1 ttl=61 time=10.922 ms
84 bytes from 10.10.110.1 icmp_seq=2 ttl=61 time=2.932 ms
84 bytes from 10.10.110.1 icmp_seq=3 ttl=61 time=2.617 ms
84 bytes from 10.10.110.1 icmp_seq=4 ttl=61 time=3.659 ms
84 bytes from 10.10.110.1 icmp_seq=5 ttl=61 time=4.695 ms
```

**HostB**

```
VPCS> ping 10.10.100.1

84 bytes from 10.10.100.1 icmp_seq=1 ttl=61 time=2.829 ms
84 bytes from 10.10.100.1 icmp_seq=2 ttl=61 time=3.203 ms
84 bytes from 10.10.100.1 icmp_seq=3 ttl=61 time=2.736 ms
84 bytes from 10.10.100.1 icmp_seq=4 ttl=61 time=2.876 ms
84 bytes from 10.10.100.1 icmp_seq=5 ttl=61 time=3.177 ms
```

## Task 7 - Finish by adding a layer of security which blocks unwanted communication.

- In this example, create a route-map which blocks HostB from communication into AS65001.
- After the route-map has been applied to the IPv4 Address Family, move to user mode and clear the BGP session to its neighbor 100.1.1.1. This will reset BGP and ensure the route-map is now working.
- Once completed, HostA and HostB should not be able to communicate, however, SwitchB will still be able to reach HostA

**SwitchA**

```
ip prefix-list net110 seq 10 permit 10.10.110.0/24

route-map ebgp-in deny seq 10
     match ip address prefix-list net110
route-map ebgp-in permit seq 20

router bgp 65001
    address-family ipv4 unicast
        neighbor 192.168.4.1 route-map ebgp-in in
    exit-address-family

clear bgp 192.168.4.1
```

**HostB**

```
VPCS> ping 10.10.110.1

10.10.110.1 icmp_seq=1 timeout
10.10.110.1 icmp_seq=2 timeout
10.10.110.1 icmp_seq=3 timeout
```

```
10.10.110.1 icmp_seq=4 timeout
10.10.110.1 icmp_seq=5 timeout
```

**SwitchB**

```
SwitchB(config-bgp)# ping 10.10.100.1
PING 10.10.100.1 (10.10.100.1) 100(128) bytes of data.
108 bytes from 10.10.100.1: icmp_seq=1 ttl=63 time=5.10 ms
108 bytes from 10.10.100.1: icmp_seq=2 ttl=63 time=1.97 ms
108 bytes from 10.10.100.1: icmp_seq=3 ttl=63 time=2.14 ms
108 bytes from 10.10.100.1: icmp_seq=4 ttl=63 time=2.15 ms
108 bytes from 10.10.100.1: icmp_seq=5 ttl=63 time=1.66 ms
```

## Appendix – Complete Configurations

**SwitchA**

```
SwitchA(config)# sh run
Current configuration:
!
!Version ArubaOS-CX Virtual.10.05.0001
!export-password: default
hostname SwitchA
user admin group administrators password ciphertext
AQBapU8kDQq8poY298UqEEFemRYpiDj0peEcT7TkWZrMRMsUYgAAAGJnbVBC6OyltRotNLsG8SnNVekmT1PwXqEVfEHoslK
MS9l4Bja5TljApzbhdVgOFwJiuZNr29MHQgDA/NwRWKzF8IZT9vPjdNXj01d3glD3OH89tYmtQciDnj2BFiFQHSUl
led locator on
!
!
!
!
ssh server vrf mgmt
vlan 1
vlan 100
    description HostA VLAN
interface mgmt
    no shutdown
    ip dhcp
interface 1/1/1
    no shutdown
    no routing
    vlan access 100
interface 1/1/2
    no shutdown
    ip address 192.168.4.0/31
interface loopback 0
    ip address 192.168.2.1/32
interface vlan 100
    description Client Segment
    ip address 10.10.100.254/24
ip prefix-list DENY seq 10 deny 10.10.110.1/32
!
!
!
!
route-map Deny deny seq 10
     match ip address prefix-list DENY
!
router bgp 65001
    bgp router-id 192.168.2.1
    neighbor 192.168.4.1 remote-as 65000
    address-family ipv4 unicast
        neighbor 192.168.4.1 activate
        neighbor 192.168.4.1 route-map Deny in
```

```
        network 10.10.100.0/24
    exit-address-family
!
https-server vrf mgmt
```

**SwitchB**

```
SwitchB(config)# sh run
Current configuration:
!
!Version ArubaOS-CX Virtual.10.05.0001
!export-password: default
hostname SwitchB
user admin group administrators password ciphertext
AQBapSpaFeGWn6tgLNlV/xyP97U+9GFVEjhilYmwPGoB1A6NYgAAAE7T4+3xUZVnlbwkaHhpEPZ2jI9joCfuLjyNwYjr1TU
8DwVHfVv+bsm9MpOzcoVa2Qy7fnxbX32zNHQtXNHZ2DJbcGDWP4nCxYWnKaqNqMO98lw5fWut36RBPHb9dMbVl03Y
led locator on
!
!
!
!
ssh server vrf mgmt
vlan 1
interface mgmt
    no shutdown
    ip dhcp
interface 1/1/1
    no shutdown
    ip address 192.168.4.1/31
interface 1/1/2
    no shutdown
    ip address 192.168.4.2/31
    ip ospf 1 area 0.0.0.0
    ip ospf network point-to-point
interface loopback 0
    ip address 192.168.2.2/32
    ip ospf 1 area 0.0.0.0
!
!
!
!
!
router ospf 1
    router-id 192.168.2.2
    area 0.0.0.0
router bgp 65000
    bgp router-id 192.168.2.2
    neighbor 192.168.2.3 remote-as 65000
    neighbor 192.168.2.3 update-source loopback 0
    neighbor 192.168.4.0 remote-as 65001
    address-family ipv4 unicast
        neighbor 192.168.2.3 activate
        neighbor 192.168.2.3 default-originate
        neighbor 192.168.4.0 activate
    exit-address-family
!
https-server vrf mgmt
```

**SwitchC**

```
SwitchC(config)# sh run
Current configuration:
!
!Version ArubaOS-CX Virtual.10.05.0001
```

```
!export-password: default
hostname SwitchC
user admin group administrators password ciphertext
AQBapaC16ngPle3gHPIf4u7jyKh3JWyfTqYsT8ufm2UXvCllYgAAACbpcU9C/HnXyLp+LSqVZB5fSC/guvr4y9HFrwzRTq0
EME69nFy0fsa+su7eObi1+ie6kpehwhMS4rjwOS2upofs4JlIwXcFS9ZyIO4KmJzOYv91xydt9HMLBZvJE3NOyxWL
led locator on
!
!
!
!
ssh server vrf mgmt
vlan 1
vlan 110
    description HostB VLAN
interface mgmt
    no shutdown
    ip dhcp
interface 1/1/1
    no shutdown
    no routing
    vlan access 110
interface 1/1/2
    no shutdown
    ip address 192.168.4.3/31
    ip ospf 1 area 0.0.0.0
    ip ospf network point-to-point
interface loopback 0
    ip address 192.168.2.3/32
    ip ospf 1 area 0.0.0.0
interface vlan 110
    description To Client Segment
    ip address 10.10.110.254/24
!
!
!
!
!
router ospf 1
    router-id 192.168.2.3
    area 0.0.0.0
router bgp 65000
    bgp router-id 192.168.2.3
    neighbor 192.168.2.2 remote-as 65000
    neighbor 192.168.2.2 update-source loopback 0
    address-family ipv4 unicast
        neighbor 192.168.2.2 activate
        network 10.10.110.0/24
    exit-address-family
!
https-server vrf mgmt
```

**HostA**

```
VPCS> sh ip

NAME        : VPCS[1]
IP/MASK     : 10.10.100.1/24
GATEWAY     : 10.10.100.254
DNS         :
MAC         : 00:50:79:66:68:07
LPORT       : 20000
RHOST:PORT  : 127.0.0.1:30000
MTU         : 1500
```

**HostB**

```
VPCS> sh ip
```

```
NAME        : VPCS[1]
IP/MASK     : 10.10.110.1/24
GATEWAY     : 10.10.110.254
DNS         :
MAC         : 00:50:79:66:68:05
LPORT       : 20000
RHOST:PORT  : 127.0.0.1:30000
MTU         : 1500
```

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