

# VRF Lab2 – Dynamic IVRL

**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.

<https://www.eve-ng.net/index.php/documentation/howtos/howto-add-aruba-cx-switch/>

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.....	2
Lab Overview.....	2
Lab Network Layout.....	3
Lab Tasks.....	3
Task 1 – Lab setup.....	3
Task 2 – Configure Layer3 for VRF-lite.....	5
Step #1: Configure VRFs .....	5
Step #2: Configure Host VLANs and Transit VLANs .....	5
Step #3: Configure SVI (Switch Virtual Interface = L3 VLAN interface).....	6
Step #4: Configure ROP (Routed Only Port) L3 interface.....	6
Step #5: Verify VRF attachment .....	7
Step #6: Routing .....	7
Task 3 – Test VRFs isolation.....	10
Task 4 – Configure dynamic route leaking.....	13
Step #1: Configure MP-BGP .....	13
Step #2: Configure VRF RDs and VRF RTs.....	14
Task 5 – Check routing tables and test inter-VRF traffic.....	15
Appendix – Reference Configurations.....	18

## Lab Objective

This lab will enable the reader to gain hands-on experience with VRF and Dynamic inter VRF route leaking (IVRL) with MP-BGP.

## Lab Overview

This lab guide explains how to configure VRFs (Virtual Routing and Forwarding) on AOS-CX switch with dynamic route leaking.

Please read the VRF section of the [AOS-CX 10.6 IP Routing Guide](https://www.arubanetworks.com/techdocs/AOS-CX/10.06/HTML/5200-7702/index.html#GUID-F2CC1540-2EFD-41FF-B3A8-9C38E9133488.html) (<https://www.arubanetworks.com/techdocs/AOS-CX/10.06/HTML/5200-7702/index.html#GUID-F2CC1540-2EFD-41FF-B3A8-9C38E9133488.html>).

During this lab, you'll be able to:

- Configure VRF and attach L3 interfaces to VRF
- Connect network nodes in a VRF-lite model
- Test traffic isolation between hosts in different VRFs
- Configure MP-BGP
- Configure dynamic inter-VRF route leaking to allow communication between hosts and server.

The minimum required AOS-CX Switch Simulator version for this lab is 10.5. It is recommended to use release 10.6 or later.

This lab uses EVE-NG but GNS3 can be used as well.

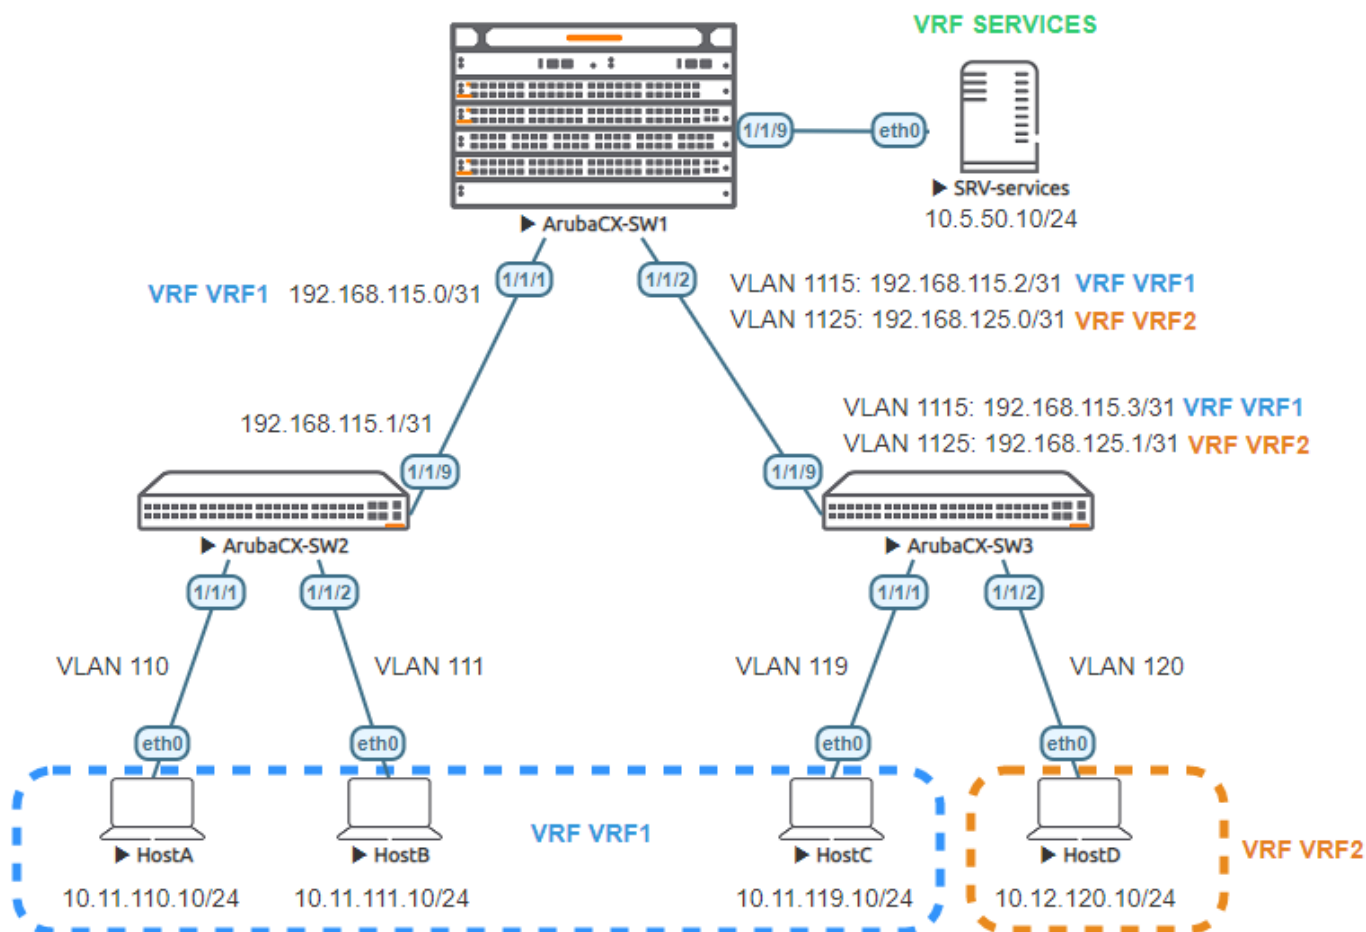
---

This lab uses the same configuration of VRF Lab1 as VRF configuration and it is highly recommended to proceed with VRF-Lab1 (static route leaking) before proceeding with this VRF-Lab2.

---

## Lab Network Layout

Here is the proposed topology (same as VRF Lab1):



## Lab Tasks

This lab uses the same configuration Task#1 and Task#2 of VRF Lab1.

### Task 1 – Lab setup

- In EVE-NG, import the .zip lab file containing the “unl” file.  
All the connections between nodes are already set-up. Appropriate numbers of CPUs (2), RAM (4096 MB) and interfaces are already allocated.
- Check the connectivity as proposed above
- Start all the devices (3 AOS-CX switches and 5 hosts)
- Open each switch console and log in with user “admin”.  
The switches will ask to enter a new password. This new password can be an empty password for simplicity in this lab.
- Apply (copy/paste) the baseline configuration as proposed below

Baseline configuration proposal (for initial copy/paste):

SW1	SW2
<pre>hostname SW1 ! vlan 1 interface mgmt   no shutdown   ip dhcp interface 1/1/1   no shutdown   description to SW2 interface 1/1/2   no shutdown   description to SW3 interface 1/1/9   no shutdown   description to SRV-services</pre>	<pre>hostname SW2 ! vlan 1 interface mgmt   no shutdown   ip dhcp interface 1/1/1   no shutdown   description to HostA interface 1/1/2   no shutdown   description to HostB interface 1/1/9   no shutdown   description to SW1</pre>
SW3	
<pre>hostname SW3 ! vlan 1 interface mgmt   no shutdown   ip dhcp interface 1/1/1   no shutdown   description to HostC interface 1/1/2   no shutdown   description to HostD interface 1/1/9   no shutdown   description to SW1</pre>	

- Verify the connectivity through LLDP neighbor information as follows:

SW1
<pre>SW1# show lldp neighbor-info  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       08:00:09:06:d8:b9  1/1/9           to SW1             120    SW2 1/1/2       08:00:09:8e:d0:6f  1/1/9           to SW1             120    SW3</pre>
SW2
<pre>SW2# show lldp neighbor-info  LLDP Neighbor Information =====  Total Neighbor Entries      : 1 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 ----- 1/1/9       08:00:09:d7:5f:0f  1/1/1           to SW2             120    SW1</pre>
SW3
<pre>SW3# show lldp neighbor-info  LLDP Neighbor Information =====</pre>

```
Total Neighbor Entries      : 1
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
1/1/9	08:00:09:d7:5f:0f	1/1/2	to SW3	120	SW1

## Task 2 – Configure Layer3 for VRF-lite

There are 2 ways to transport VRF in a VRF-lite architecture:

- through ROP (Routed Only Port): one VRF per interface in case of a single VRF or one VRF per sub-interface in case of multiple VRFs (not yet supported on AOS-CX Simulator)
- through Transit VLANs, each Transit VLAN being associated to one VRF for multiple VRFs case.

Both methods are used in this lab for educational purpose. SW2 will use ROP with one VRF only. SW3 will use Transit VLANs.

### Step #1: Configure VRFs

SW1 will host 3 VRFs:

- VRF1, for VRF-lite interconnectivity to SW1
- VRF2, for VRF-lite interconnectivity to SW2
- SERVICES, for hosting SRV-services server in the SERVICES VRF.

SW2 will use only default VRF. Indeed, default VRF in access SW2 is mapped to VRF1 on SW1 interconnection. This is done for simplification. An alternative would have been to configure VRF1 as well on SW2 and attach all L3 interfaces in VRF1. As there is no other VRFs hosted in SW2, it is simpler to just use default VRF and bind it to VRF1 through the VRF attachment on SW1 interconnection.

SW3 will host 2 VRFs:

- VRF1, for VRF-lite interconnectivity to SW1, and for hosting VRF1 endpoint: HostC.
- VRF2, for VRF-lite interconnectivity to SW1, and for hosting VRF2 endpoint: HostD

SW1(config)#	SW3(config)#
vrf VRF1 vrf VRF2 vrf SERVICES	vrf VRF1 vrf VRF2

**Note:** RD (route-distinguisher) in the VRF context is configure later in the task#4.

### Step #2: Configure Host VLANs and Transit VLANs

VLANs are used for endpoint Hosts, and for Transit VLANs.

Transit VLAN 1115 is used for VRF1 and Transit VLAN 1125 is used for VRF2.

VLAN 110, 111, 119 are endpoints VLANs for VRF1, VLANs 110 and 111 used on SW2, VLAN 119 used on SW3.

VLAN 120 is the endpoint VLAN for VRF2 on SW3.

SW1(config)#	SW2(config)#
vlan 1115,1125 ! interface 1/1/2 no shutdown description to SW3	vlan 110-111 ! interface 1/1/1 no shutdown description to HostA

<pre>no routing vlan trunk native 1 vlan trunk allowed 1115,1125</pre>	<pre>no routing vlan access 110 interface 1/1/2 no shutdown description to HostB no routing vlan access 111</pre>
<pre>SW3(config)# vlan 119-120,1115,1125 ! interface 1/1/1 no shutdown description to HostC no routing vlan access 119 interface 1/1/2 no shutdown description to HostD no routing vlan access 120 interface 1/1/9 no shutdown description to SW1 no routing vlan trunk native 1 vlan trunk allowed 1115,1125</pre>	

### Step #3: Configure SVI (Switch Virtual Interface = L3 VLAN interface)

VRF binding is configured in this step. **Reminder:** it was chosen to not configure VRF in SW2 for simplicity and educational purpose.

<pre>SW1(config)# interface vlan 1115 vrf attach VRF1 ip address 192.168.115.2/31 interface vlan 1125 vrf attach VRF2 ip address 192.168.125.0/31</pre>	<pre>SW2(config)# interface vlan 110 ip address 10.11.110.1/24 interface vlan 111 ip address 10.11.111.1/24</pre>
<pre>SW3(config)# interface vlan 119 vrf attach VRF1 ip address 10.11.119.1/24 interface vlan 120 vrf attach VRF2 ip address 10.12.120.1/24 interface vlan 1115 vrf attach VRF1 ip address 192.168.115.3/31 interface vlan 1125 vrf attach VRF2 ip address 192.168.125.1/31</pre>	

### Step #4: Configure ROP (Routed Only Port) L3 interface

On SW1, ROP to SW2 is attached to VRF1, whereas it is attached to default VRF on SW2.

On SW1, a ROP is used for Lab simplicity to connect the server SRV-services.

<pre>SW1(config)# interface 1/1/1 no shutdown vrf attach VRF1 description to SW2 ip address 192.168.115.0/31 interface 1/1/9 no shutdown vrf attach SERVICES</pre>	<pre>SW2(config)# interface 1/1/9 no shutdown description to SW1 ip address 192.168.115.1/31</pre>
--	--

```
description to SRV-services
ip address 10.5.50.1/24
```

### Step #5: Verify VRF attachment

SW1(config)#	SW2(config)#
<pre>SW1# show vrf VRF Configuration: ----- VRF Name   : default Interfaces           Status ----- 1/1/3         down 1/1/4         down 1/1/5         down 1/1/6         down 1/1/7         down 1/1/8         down  VRF Name   : SERVICES Interfaces           Status ----- 1/1/9         up  VRF Name   : VRF1 Interfaces           Status ----- 1/1/1         up vlan1115      up  VRF Name   : VRF2 Interfaces           Status ----- vlan1125      up</pre>	<pre>SW2# show vrf VRF Configuration: ----- VRF Name   : default Interfaces           Status ----- 1/1/3         down 1/1/4         down 1/1/5         down 1/1/6         down 1/1/7         down 1/1/8         down 1/1/9         up vlan110       up vlan111       up</pre>
SW3(config)#	
<pre>SW3# show vrf VRF Configuration: ----- VRF Name   : default Interfaces           Status ----- 1/1/3         down 1/1/4         down 1/1/5         down 1/1/6         down 1/1/7         down 1/1/8         down  VRF Name   : VRF1 Interfaces           Status ----- vlan119       up vlan1115      up  VRF Name   : VRF2 Interfaces           Status ----- vlan120       up vlan1125      up</pre>	

### Step #6: Routing

In this lab, static routing is used for simplicity on each network node within the given VRFs, whereas MP-BGP is used only on SW1 as the underlying protocol to learn NLRI (Network Layer Reachability Information) that enables dynamic route leaking with route-targets between VRFs. This lab uses a basic set-up in order to focus on the dynamic route leaking concept. Of course,

dynamic protocols such as OSPF or BGP could have been set-up between SW1 and SW2, and between SW1 and SW3 instead of static routes.

On SW1, we need to create a route to reach 10.11.110.0/24 and 10.11.111.0/24. This is summarized with 10.11.96.0/20 with Next-Hop being SW2 IP address. Similarly a route entry is created for 10.12.0.0/16 pointing to SW3 IP address as Next-Hop.

On SW2, a default route is enough. On SW3, a default route per VRF is used as well.

SW1(config)#	SW2(config)#
<pre>ip route 10.11.96.0/20 192.168.115.1 vrf VRF1 ip route 10.11.119.0/24 192.168.115.3 vrf VRF1 ip route 10.12.0.0/16 192.168.125.1 vrf VRF2</pre>	<pre>ip route 0.0.0.0/0 192.168.115.0</pre>
SW3(config)#	
<pre>ip route 0.0.0.0/0 192.168.115.2 vrf VRF1 ip route 0.0.0.0/0 192.168.125.0 vrf VRF2</pre>	

Verify the routing table on each node. Here on SW1:

```
SW1
SW1# show ip route

No ipv4 routes configured
```

There is no route in default VRF in SW1 as expected.

```
SW1
SW1# show ip route vrf VRF1

Displaying ipv4 routes selected for forwarding

'[x/y]' denotes [distance/metric]

10.11.96.0/20, vrf VRF1
    via 192.168.115.1, [1/0], static
10.11.119.0/24, vrf VRF1
    via 192.168.115.3, [1/0], static
192.168.115.0/31, vrf VRF1
    via 1/1/1, [0/0], connected
192.168.115.0/32, vrf VRF1
    via 1/1/1, [0/0], local
192.168.115.2/31, vrf VRF1
    via vlan1115, [0/0], connected
192.168.115.2/32, vrf VRF1
    via vlan1115, [0/0], local
```

For VRF1, there are local /32 entry, connected /31 entry and static routes to SW2 and SW3.

```
SW1
SW1# show ip route vrf VRF2

Displaying ipv4 routes selected for forwarding

'[x/y]' denotes [distance/metric]

10.12.0.0/16, vrf VRF2
    via 192.168.125.1, [1/0], static
192.168.125.0/31, vrf VRF2
    via vlan1125, [0/0], connected
192.168.125.0/32, vrf VRF2
    via vlan1125, [0/0], local
```

Similarly for VRF2. And finally for VRF SERVICES:

```
SW1
SW1# show ip route vrf SERVICES
```



```

Displaying ipv4 routes selected for forwarding

'[x/y]' denotes [distance/metric]

10.5.50.0/24, vrf SERVICES
    via 1/1/9, [0/0], connected
10.5.50.1/32, vrf SERVICES
    via 1/1/9, [0/0], local
  
```

On SW2:

```

SW2
SW2# show ip route

Displaying ipv4 routes selected for forwarding

'[x/y]' denotes [distance/metric]

0.0.0.0/0, vrf default
    via 192.168.115.0, [1/0], static
10.11.110.0/24, vrf default
    via vlan110, [0/0], connected
10.11.110.1/32, vrf default
    via vlan110, [0/0], local
10.11.111.0/24, vrf default
    via vlan111, [0/0], connected
10.11.111.1/32, vrf default
    via vlan111, [0/0], local
192.168.115.0/31, vrf default
    via 1/1/9, [0/0], connected
192.168.115.1/32, vrf default
    via 1/1/9, [0/0], local
  
```

On SW3:

```

SW2
SW3# show ip route

No ipv4 routes configured

SW3# show ip route vrf VRF1

Displaying ipv4 routes selected for forwarding

'[x/y]' denotes [distance/metric]

0.0.0.0/0, vrf VRF1
    via 192.168.115.2, [1/0], static
10.11.119.0/24, vrf VRF1
    via vlan119, [0/0], connected
10.11.119.1/32, vrf VRF1
    via vlan119, [0/0], local
192.168.115.2/31, vrf VRF1
    via vlan1115, [0/0], connected
192.168.115.3/32, vrf VRF1
    via vlan1115, [0/0], local

SW3# show ip route vrf VRF2

Displaying ipv4 routes selected for forwarding

'[x/y]' denotes [distance/metric]

0.0.0.0/0, vrf VRF2
    via 192.168.125.0, [1/0], static
10.12.120.0/24, vrf VRF2
    via vlan120, [0/0], connected
10.12.120.1/32, vrf VRF2
    via vlan120, [0/0], local
192.168.125.0/31, vrf VRF2
    via vlan1125, [0/0], connected
192.168.125.1/32, vrf VRF2
  
```

```
via vlan1125, [0/0], local
```

The main configuration on SW1, SW2 and SW3 is ready to start performing connectivity tests.

### Task 3 – Test VRFs isolation

As a reference, the configuration of SW1/SW2/SW3 should look like:

SW1	SW2
<pre>hostname SW1 ! vrf SERVICES vrf VRF1 vrf VRF2 ! vlan 1,1115,1125 interface mgmt   no shutdown   ip dhcp interface 1/1/1   no shutdown   vrf attach VRF1   description to SW2   ip address 192.168.115.0/31 interface 1/1/2   no shutdown   description to SW3   no routing   vlan trunk native 1   vlan trunk allowed 1115,1125 interface 1/1/9   no shutdown   vrf attach SERVICES   description to SRV-services   ip address 10.5.50.1/24 interface vlan 1115   vrf attach VRF1   ip address 192.168.115.2/31 interface vlan 1125   vrf attach VRF2   ip address 192.168.125.0/31 ip route 10.11.96.0/20 192.168.115.1 vrf VRF1 ip route 10.11.119.0/24 192.168.115.3 vrf VRF1 ip route 10.12.0.0/16 192.168.125.1 vrf VRF2 !</pre>	<pre>hostname SW2 ! vlan 1,110-111 interface mgmt   no shutdown   ip dhcp interface 1/1/1   no shutdown   description to HostA   no routing   vlan access 110 interface 1/1/2   no shutdown   description to HostB   no routing   vlan access 111 interface 1/1/9   no shutdown   description to SW1   ip address 192.168.115.1/31 interface vlan 110   ip address 10.11.110.1/24 interface vlan 111   ip address 10.11.111.1/24 ip route 0.0.0.0/0 192.168.115.0 !</pre>
<pre>hostname SW3 ! vrf VRF1 vrf VRF2 vlan 1,119-120,1115,1125 interface mgmt   no shutdown   ip dhcp interface 1/1/1   no shutdown   description to HostC   no routing   vlan access 119 interface 1/1/2   no shutdown   description to HostD   no routing   vlan access 120 interface 1/1/9   no shutdown   description to SW1   no routing   vlan trunk native 1   vlan trunk allowed 1115,1125 interface vlan 119</pre>	

```
vrf attach VRF1
ip address 10.11.119.1/24
interface vlan 120
vrf attach VRF2
ip address 10.12.120.1/24
interface vlan 1115
vrf attach VRF1
ip address 192.168.115.3/31
interface vlan 1125
vrf attach VRF2
ip address 192.168.125.1/31
ip route 0.0.0.0/0 192.168.115.2 vrf VRF1
ip route 0.0.0.0/0 192.168.125.0 vrf VRF2
```

### Set-up IP address on HostA and HostB:

HostA	HostB
<pre>VPCS&gt; ip 10.11.110.10/24 10.11.110.1 Checking for duplicate address. VPCS : 10.11.110.10 255.255.255.0 gateway 10.11.110.1  VPCS&gt; show ip  NAME       : VPCS[1] IP/MASK    : 10.11.110.10/24 GATEWAY    : 10.11.110.1 DNS        : MAC        : 00:50:79:66:68:07 LPORT     : 20000 RHOST:PORT : 127.0.0.1:30000 MTU        : 1500</pre>	<pre>VPCS&gt; ip 10.11.111.10/24 10.11.111.1 Checking for duplicate address VPCS : 10.11.111.10 255.255.255.0 gateway 10.11.111.1  VPCS&gt; show ip  NAME       : VPCS[1] IP/MASK    : 10.11.111.10/24 GATEWAY    : 10.11.111.1 DNS        : MAC        : 00:50:79:66:68:06 LPORT     : 20000 RHOST:PORT : 127.0.0.1:30000 MTU        : 1500</pre>
HostC	HostD
<pre>VPCS&gt; ip 10.11.119.10/24 10.11.119.1 Checking for duplicate address... VPCS : 10.11.119.10 255.255.255.0 gateway 10.11.119.1  VPCS&gt; show ip  NAME       : VPCS[1] IP/MASK    : 10.11.119.10/24 GATEWAY    : 10.11.119.1 DNS        : MAC        : 00:50:79:66:68:05 LPORT     : 20000 RHOST:PORT : 127.0.0.1:30000 MTU        : 1500</pre>	<pre>VPCS&gt; ip 10.12.120.10/24 10.12.120.1 Checking for duplicate address... VPCS : 10.12.120.10 255.255.255.0 gateway 10.12.120.1  VPCS&gt; show ip  NAME       : VPCS[1] IP/MASK    : 10.12.120.10/24 GATEWAY    : 10.12.120.1 DNS        : MAC        : 00:50:79:66:68:08 LPORT     : 20000 RHOST:PORT : 127.0.0.1:30000 MTU        : 1500</pre>

### SRV-services

```
VPCS> ip 10.5.50.10/24 10.5.50.1
Checking for duplicate address...
VPCS : 10.5.50.10 255.255.255.0 gateway 10.5.50.1

VPCS> show ip

NAME       : VPCS[1]
IP/MASK    : 10.5.50.10/24
GATEWAY    : 10.5.50.1
DNS        :
MAC        : 00:50:79:66:68:04
LPORT     : 20000
RHOST:PORT : 127.0.0.1:30000
MTU        : 1500
```

### Ping inside the same VRF:

#### Ping HostB from HostA (VRF1)

```
HostA
VPCS> ping 10.11.111.10
```

```
84 bytes from 10.11.111.10 icmp_seq=1 ttl=63 time=2.815 ms
84 bytes from 10.11.111.10 icmp_seq=2 ttl=63 time=6.434 ms
84 bytes from 10.11.111.10 icmp_seq=3 ttl=63 time=1.307 ms
84 bytes from 10.11.111.10 icmp_seq=4 ttl=63 time=1.224 ms
84 bytes from 10.11.111.10 icmp_seq=5 ttl=63 time=5.006 ms
```

#### Ping HostC from HostA (VRF1)

##### HostA

```
VPCS> ping 10.11.119.10

84 bytes from 10.11.119.10 icmp_seq=1 ttl=61 time=10.754 ms
84 bytes from 10.11.119.10 icmp_seq=2 ttl=61 time=9.072 ms
84 bytes from 10.11.119.10 icmp_seq=3 ttl=61 time=4.065 ms
84 bytes from 10.11.119.10 icmp_seq=4 ttl=61 time=3.620 ms
84 bytes from 10.11.119.10 icmp_seq=5 ttl=61 time=3.573 ms
```

#### Ping SW1 VRF2 IP address from HostD (VRF2)

##### HostD

```
VPCS> ping 192.168.125.0

84 bytes from 192.168.125.0 icmp_seq=1 ttl=63 time=2.741 ms
84 bytes from 192.168.125.0 icmp_seq=2 ttl=63 time=7.833 ms
84 bytes from 192.168.125.0 icmp_seq=3 ttl=63 time=2.987 ms
84 bytes from 192.168.125.0 icmp_seq=4 ttl=63 time=2.900 ms
84 bytes from 192.168.125.0 icmp_seq=5 ttl=63 time=2.792 ms
```

#### Ping between VRFs:

The purpose of VRFs is to isolate routing domains. As a consequence, without any inter-VRF route leaking, hosts in VRF1 should not communicate with hosts in other VRFs.

#### Ping HostD (VRF2) from HostA(VRF1):

##### HostA

```
VPCS> ping 10.12.120.10

*192.168.115.0 icmp_seq=1 ttl=63 time=3.025 ms (ICMP type:3, code:0, Destination network unreachable)
*192.168.115.0 icmp_seq=2 ttl=63 time=2.367 ms (ICMP type:3, code:0, Destination network unreachable)
*192.168.115.0 icmp_seq=3 ttl=63 time=2.305 ms (ICMP type:3, code:0, Destination network unreachable)
*192.168.115.0 icmp_seq=4 ttl=63 time=2.328 ms (ICMP type:3, code:0, Destination network unreachable)
10.12.120.10 icmp seq=5 timeout
```

#### Ping SRV-services(SERVICES VRF) from HostA(VRF1):

##### HostA

```
VPCS> ping 10.5.50.10

*192.168.115.0 icmp_seq=1 ttl=63 time=2.514 ms (ICMP type:3, code:0, Destination network unreachable)
*192.168.115.0 icmp_seq=2 ttl=63 time=7.301 ms (ICMP type:3, code:0, Destination network unreachable)
*192.168.115.0 icmp_seq=3 ttl=63 time=2.651 ms (ICMP type:3, code:0, Destination network unreachable)
*192.168.115.0 icmp_seq=4 ttl=63 time=2.048 ms (ICMP type:3, code:0, Destination network unreachable)
10.5.50.10 icmp seq=5 timeout
```

#### Ping SRV-services(SERVICES VRF) from HostD(VRF2):

##### HostA

```
VPCS> ping 10.5.50.10

10.5.50.10 icmp_seq=1 timeout
10.5.50.10 icmp_seq=2 timeout
10.5.50.10 icmp_seq=3 timeout
10.5.50.10 icmp_seq=4 timeout
10.5.50.10 icmp_seq=5 timeout
```

Between VRF the network is unreachable or timeout, as expected.

The next section explain how to make communication between VRF1 and SERVICES, and between VRF2 and SERVICES, while maintaining isolation between VRF1 and VRF2.

## Task 4 – Configure dynamic route leaking

Here are the route-leaking lab objectives:

- Hosts in VRF1 need to access server in SERVICES VRF.
- Hosts in VRF2 need to access server in SERVICES VRF.
- Hosts in VRF1 should not be able to communicate with hosts in VRF2.

The network node used in this lab to perform inter-VRF route leaking is **SW1**.

In order for each virtual routing domain to know how to reach SRV-services, routes information must be learnt in each VRF. In VRF Lab1 it was achieved with static routes. In this lab, it is achieved with MP-BGP on SW1 node performing inter-VRF communication.

### Step #1: Configure MP-BGP

In this lab, no BGP peering is created, consequently the AS number does not matter. In real deployment, AS number should be selected appropriately based on the existing AS domain if already set-up. It is a best-practice to define the router-id, as this router-id is used as Route-Distinguisher for ease of troubleshooting (not driven by technical reason, purely for operational simplicity). Router-ID is usually defined as the IP address of Loopback 0 interface.

MP-BGP IPv4 unicast Address-Family is configured for each VRF. Connected and Static redistributions are used to inject prefixes in BGP for each VRF.

```
SW1(config)#
router bgp 65001
  bgp router-id 192.168.2.1
!
  vrf SERVICES
    address-family ipv4 unicast
      redistribute connected
      redistribute static
    exit-address-family
!
  vrf VRF1
    address-family ipv4 unicast
      redistribute connected
      redistribute static
    exit-address-family
!
  vrf VRF2
    address-family ipv4 unicast
      redistribute connected
      redistribute static
    exit-address-family
```

Check BGP RIB for each VRF. Each VRF should have redistributed static routes and connected in BGP.

```
SW1
SW1# show bgp vrf VRF1 ipv4 unicast
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed, a additional-paths
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : VRF1
Local Router-ID 192.168.115.2

   Network          Nexthop          Metric    LocPrf    Weight Path
Route Distinguisher: 192.168.2.1:1
*> 10.11.96.0/20     192.168.115.1   0          100       0       ?
*> 10.11.119.0/24   192.168.115.3   0          100       0       ?
*> 192.168.115.0/31 0.0.0.0         0          100       0       ?
*> 192.168.115.2/31 0.0.0.0         0          100       0       ?
Total number of entries 4
```

```
SW1# show bgp vrf VRF2 ipv4 unicast
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed, a additional-paths
Origin codes: i - IGP, e - EGP, ? - incomplete
```

```
VRF : VRF2
Local Router-ID 192.168.125.0
```

Network	Nexthop	Metric	LocPrf	Weight	Path
Route Distinguisher: 192.168.2.1:2					
*> 10.12.0.0/16	192.168.125.1	0	100	0	?
*> 192.168.125.0/31	0.0.0.0	0	100	0	?

Total number of entries 2

```
SW1# show bgp vrf SERVICES ipv4 unicast
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed, a additional-paths
Origin codes: i - IGP, e - EGP, ? - incomplete
```

```
VRF : SERVICES
Local Router-ID 10.5.50.1
```

Network	Nexthop	Metric	LocPrf	Weight	Path
Route Distinguisher: 192.168.2.1:5					
*> 10.5.50.0/24	0.0.0.0	0	100	0	?

Total number of entries 1

## Step #2: Configure VRF RDs and VRF RTs

In order to perform route leaking, Route Distinguisher (RD) is configured per VRF. This unique number, prepended to the routes within the VRF, ensures the support for route identification across different VRFs.

Routes can then be selectively imported and exported across VRFs using Route Target (RT) that are filters, defined in each VRF.

In order for SERVICES routing domain to know how to reach hosts in VRF1 and VRF2, the route-target in SERVICES VRF must import routes that are exported by VRF1 and VRF2.

In order for VRF1 and VRF2 routing domains to know how to reach SERVICES hosts, the route-target in VRF1/VRF2 must import routes that are exported by SERVICES VRF.

```
SW1(config)#
vrf SERVICES
  rd 192.168.2.1:5
  address-family ipv4 unicast
    route-target export 65001:5
    route-target import 65001:1
    route-target import 65001:2
  exit-address-family
vrf VRF1
  rd 192.168.2.1:1
  address-family ipv4 unicast
    route-target export 65001:1
    route-target import 65001:5
  exit-address-family
vrf VRF2
  rd 192.168.2.1:2
  address-family ipv4 unicast
    route-target export 65001:2
    route-target import 65001:5
  exit-address-family
```

## Task 5 – Check routing tables and test inter-VRF traffic

Check the updated BGP table per VRF and compare with the previous one from Step#1 in Task 4. Pay attention to new entries highlighted in blue.

```
SW1
W1# show bgp vrf VRF1 ipv4 unicast
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed, a additional-paths
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : VRF1
Local Router-ID 192.168.115.2

      Network          Nexthop          Metric      LocPrf      Weight Path
Route Distinguisher: 192.168.2.1:1
*> 10.5.50.0/24        0.0.0.0          0            100         0          i
*> 10.11.96.0/20       192.168.115.1    0            100         0          ?
*> 10.11.119.0/24     192.168.115.3    0            100         0          ?
*> 192.168.115.0/31   0.0.0.0          0            100         0          ?
*> 192.168.115.2/31   0.0.0.0          0            100         0          ?
Total number of entries 5

SW1# show bgp vrf VRF2 ipv4 unicast
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed, a additional-paths
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : VRF2
Local Router-ID 192.168.125.0

      Network          Nexthop          Metric      LocPrf      Weight Path
Route Distinguisher: 192.168.2.1:2
*> 10.5.50.0/24        0.0.0.0          0            100         0          i
*> 10.12.0.0/16       192.168.125.1    0            100         0          ?
*> 192.168.125.0/31   0.0.0.0          0            100         0          ?
Total number of entries 3

SW1# show bgp vrf SERVICES ipv4 unicast
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed, a additional-paths
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : SERVICES
Local Router-ID 10.5.50.1

      Network          Nexthop          Metric      LocPrf      Weight Path
Route Distinguisher: 192.168.2.1:5
*> 10.5.50.0/24        0.0.0.0          0            100         0          ?
*> 10.11.96.0/20       192.168.115.1    0            100         0          i
*> 10.11.119.0/24     192.168.115.3    0            100         0          i
*> 10.12.0.0/16       192.168.125.1    0            100         0          i
*> 192.168.115.0/31   0.0.0.0          0            100         0          i
*> 192.168.115.2/31   0.0.0.0          0            100         0          i
*> 192.168.125.0/31   0.0.0.0          0            100         0          i
Total number of entries 7
```

Check the corresponding routing table per VRF (you may check differences with VRF Lab1: bgp versus static routing)

```
SW1
SW1# show ip route vrf VRF1

Displaying ipv4 routes selected for forwarding

'[x/y]' denotes [distance/metric]

10.5.50.0/24, vrf VRF1
```

```

    via 1/1/9[vrf SERVICES], [200/0], bgp
10.11.96.0/20, vrf VRF1
    via 192.168.115.1, [1/0], static
10.11.119.0/24, vrf VRF1
    via 192.168.115.3, [1/0], static
192.168.115.0/31, vrf VRF1
    via 1/1/1, [0/0], connected
192.168.115.0/32, vrf VRF1
    via 1/1/1, [0/0], local
192.168.115.2/31, vrf VRF1
    via vlan1115, [0/0], connected
192.168.115.2/32, vrf VRF1
    via vlan1115, [0/0], local

```

You can see a route entry coming from the egress SERVICES VRF.

```

SW1
SW1# show ip route vrf VRF2

Displaying ipv4 routes selected for forwarding

'[x/y]' denotes [distance/metric]

10.5.50.0/24, vrf VRF2
    via 1/1/9[vrf SERVICES], [200/0], bgp
10.12.0.0/16, vrf VRF2
    via 192.168.125.1, [1/0], static
192.168.125.0/31, vrf VRF2
    via vlan1125, [0/0], connected
192.168.125.0/32, vrf VRF2
    via vlan1125, [0/0], local

```

Similarly for VRF2, a route entry is present from the egress SERVICES VRF.

```

SW1
SW1# show ip route vrf SERVICES

Displaying ipv4 routes selected for forwarding

'[x/y]' denotes [distance/metric]

10.5.50.0/24, vrf SERVICES
    via 1/1/9, [0/0], connected
10.5.50.1/32, vrf SERVICES
    via 1/1/9, [0/0], local
10.11.96.0/20, vrf SERVICES
    via 192.168.115.1[vrf VRF1], [200/0], bgp
10.11.119.0/24, vrf SERVICES
    via 192.168.115.3[vrf VRF1], [200/0], bgp
10.12.0.0/16, vrf SERVICES
    via 192.168.125.1[vrf VRF2], [200/0], bgp
192.168.115.0/31, vrf SERVICES
    via 1/1/1[vrf VRF1], [200/0], bgp
192.168.115.2/31, vrf SERVICES
    via vlan1115[vrf VRF1], [200/0], bgp
192.168.125.0/31, vrf SERVICES
    via vlan1125[vrf VRF2], [200/0], bgp

```

Finally, SERVICES routing table includes routes for egress VRFs VRF1 and VRF2, including connected subnet which are important to resolve the reachability of the next-hop within the VRF. (You may try removing connected redistribution in BGP, traffic will not work).

Test again the connectivity between Hosts and then between hosts and server:

Ping HostD (VRF2) from HostA(VRF1):

```

HostA
VPCS> ping 10.12.120.10

*192.168.115.0 icmp_seq=1 ttl=63 time=3.064 ms (ICMP type:3, code:0, Destination net work unreachable)
*192.168.115.0 icmp_seq=2 ttl=63 time=6.026 ms (ICMP type:3, code:0, Destination net work unreachable)
*192.168.115.0 icmp_seq=3 ttl=63 time=2.927 ms (ICMP type:3, code:0, Destination net work unreachable)
*192.168.115.0 icmp_seq=4 ttl=63 time=2.455 ms (ICMP type:3, code:0, Destination net work unreachable)

```



```
10.12.120.10 icmp_seq=5 timeout
```

This is still not possible as expected and desired.

Ping SRV-services(SERVICES VRF) from HostA(VRF1):

```
HostA
VPCS> ping 10.5.50.10

84 bytes from 10.5.50.10 icmp_seq=1 ttl=61 time=11.072 ms
84 bytes from 10.5.50.10 icmp_seq=2 ttl=61 time=3.646 ms
84 bytes from 10.5.50.10 icmp_seq=3 ttl=61 time=3.019 ms
84 bytes from 10.5.50.10 icmp_seq=4 ttl=61 time=2.774 ms
84 bytes from 10.5.50.10 icmp_seq=5 ttl=61 time=2.805 ms
```

The communication is now possible between Hosts in VRF1 and SRV-services in SERVICES VRF.

Similarly for HostD in VRF2:

Ping SRV-services(SERVICES VRF) from HostD(VRF2):

```
HostA
VPCS> ping 10.5.50.10

84 bytes from 10.5.50.10 icmp_seq=1 ttl=61 time=14.803 ms
84 bytes from 10.5.50.10 icmp_seq=2 ttl=61 time=3.532 ms
84 bytes from 10.5.50.10 icmp_seq=3 ttl=61 time=3.393 ms
84 bytes from 10.5.50.10 icmp_seq=4 ttl=61 time=3.542 ms
84 bytes from 10.5.50.10 icmp_seq=5 ttl=61 time=3.558 ms
```

This is the end of this lab.

## Appendix – Reference Configurations

If you face issues during your lab, you can verify your configuration with the configuration extract listed in this section.

```
SW1
hostname SW1
!
vrf SERVICES
  rd 192.168.2.1:5
  address-family ipv4 unicast
    route-target export 65001:5
    route-target import 65001:1
    route-target import 65001:2
  exit-address-family
vrf VRF1
  rd 192.168.2.1:1
  address-family ipv4 unicast
    route-target export 65001:1
    route-target import 65001:5
  exit-address-family
vrf VRF2
  rd 192.168.2.1:2
  address-family ipv4 unicast
    route-target export 65001:2
    route-target import 65001:5
  exit-address-family
!
vlan 1,1115,1125
interface mgmt
  no shutdown
  ip dhcp
interface 1/1/1
  no shutdown
  vrf attach VRF1
  description to SW2
  ip address 192.168.115.0/31
interface 1/1/2
  no shutdown
  description to SW3
  no routing
  vlan trunk native 1
  vlan trunk allowed 1115,1125
interface 1/1/9
  no shutdown
  vrf attach SERVICES
  description to SRV-services
  ip address 10.5.50.1/24
interface vlan 1115
  vrf attach VRF1
  ip address 192.168.115.2/31
interface vlan 1125
  vrf attach VRF2
  ip address 192.168.125.0/31
!
ip route 10.11.96.0/20 192.168.115.1 vrf VRF1
ip route 10.11.119.0/24 192.168.115.3 vrf VRF1
ip route 10.12.0.0/16 192.168.125.1 vrf VRF2
!
router bgp 65001
  bgp router-id 192.168.2.1
!
  vrf SERVICES
    address-family ipv4 unicast
      redistribute connected
      redistribute static
    exit-address-family
!
  vrf VRF1
    address-family ipv4 unicast
      redistribute connected
      redistribute static
    exit-address-family
!
```

```
vrf VRF2
  address-family ipv4 unicast
    redistribute connected
    redistribute static
  exit-address-family
```

### SW2

```
hostname SW2
!
vlan 1,110-111
interface mgmt
  no shutdown
  ip dhcp
interface 1/1/1
  no shutdown
  description to HostA
  no routing
  vlan access 110
interface 1/1/2
  no shutdown
  description to HostB
  no routing
  vlan access 111
interface 1/1/9
  no shutdown
  description to SW1
  ip address 192.168.115.1/31
interface vlan 110
  ip address 10.11.110.1/24
interface vlan 111
  ip address 10.11.111.1/24
ip route 0.0.0.0/0 192.168.115.0
```

### SW3

```
hostname SW3
!
vrf VRF1
vrf VRF2
!
vlan 1,119-120,1115,1125
interface mgmt
  no shutdown
  ip dhcp
interface 1/1/1
  no shutdown
  description to HostC
  no routing
  vlan access 119
interface 1/1/2
  no shutdown
  description to HostD
  no routing
  vlan access 120
interface 1/1/9
  no shutdown
  description to SW1
  no routing
  vlan trunk native 1
  vlan trunk allowed 1115,1125
interface vlan 119
  vrf attach VRF1
  ip address 10.11.119.1/24
interface vlan 120
  vrf attach VRF2
  ip address 10.12.120.1/24
interface vlan 1115
  vrf attach VRF1
  ip address 192.168.115.3/31
interface vlan 1125
  vrf attach VRF2
```

```
ip address 192.168.125.1/31
!  
ip route 0.0.0.0/0 192.168.115.2 vrf VRF1  
ip route 0.0.0.0/0 192.168.125.0 vrf VRF2
```

