LAB GUIDE

VSX Lab2 – Layer3

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

This lab will enable the reader to gain hands-on experience with VSX and basic Layer3 configuration.

Lab Overview

This lab guide explains how to configure a VSX cluster of a pair of AOS-CX switches following the VSX Configuration Best Practices (https://support.hpe.com/hpesc/public/docDisplay?docId=a00094242en_us), for Layer3 IPv4 and IPv6 networks.

Please read also the AOS-CX 10.6 Virtual Switching Extension (VSX) Guide (https://www.arubanetworks.com/techdocs/AOS-CX/10.06/HTML/5200-7727/index.html#book.html).

In this lab, you'll be able to:

- Configure VSX and VSX LAG (MCLAG) for IPv4 and IPv6 networks
- Test L3 connectivity between HostA and HostB that are in different subnets
- Test solution resiliency by isolating one of the VSX node (or power-off).

This lab uses the configuration of VSX Lab1 as startup configuration and it is highly recommended to proceed with VSX-Lab1 before proceeding with this VSX-Lab2.

Note: HostB IP address is different than in VSX-Lab1.

The minimum recommended AOS-CX Switch Simulator version for this lab is 10.06.0110.

This lab uses EVE-NG Pro for Graph of links utilization. This is optional and EVE-NG Community or GNS3 can be used as well without graphs by using show interface command instead.

VSX LAG CAVEAT:

If you need to stop the AOS-CX virtual switches already configured with VSX LAGs and you need to start them again later, then there is currently a limitation in the AOS-CX Switch Simulator that prevents the switches, starting with the VSX LAGs configuration, to forward traffic on the VSX LAGs. The following workaround is required to restore the nodes for appropriate forwarding state:

- Before CX virtual switch shutdown, shutdown all interfaces (1/1/1-1/1/9) and remove interface from VSX LAG (no lag command under the interfaces that are part of a multi-chassis LAG).

- Then AOS-CX virtual switch can be stopped.

- After restarting CX virtual switch, re-enable all interfaces (this will clean-up the INVALID MTU state of interfaces) and re-assign the physical port to the desired VSX LAGs (lag command under interface context).

This will restore the AOS-CX virtual nodes with VSX LAGs in a proper state, ready to forward traffic.

if you face an issue with traffic forwarding on a CX Switch Simulator lab configured with VSX LAGs, the following tip might be very useful to remind:
- on the VSX nodes: remove ports from VSX LAGs, shut all ports, write mem, reboot, no shut all ports and finally re-assign ports to the VSX LAGs.
- on the LACP neighbors of VSX nodes, shut/no shut all ports that are members of LAG connected to the VSX nodes.

Lab Network Layout

Here is the proposed topology to study VSX technology and basic Layer3.

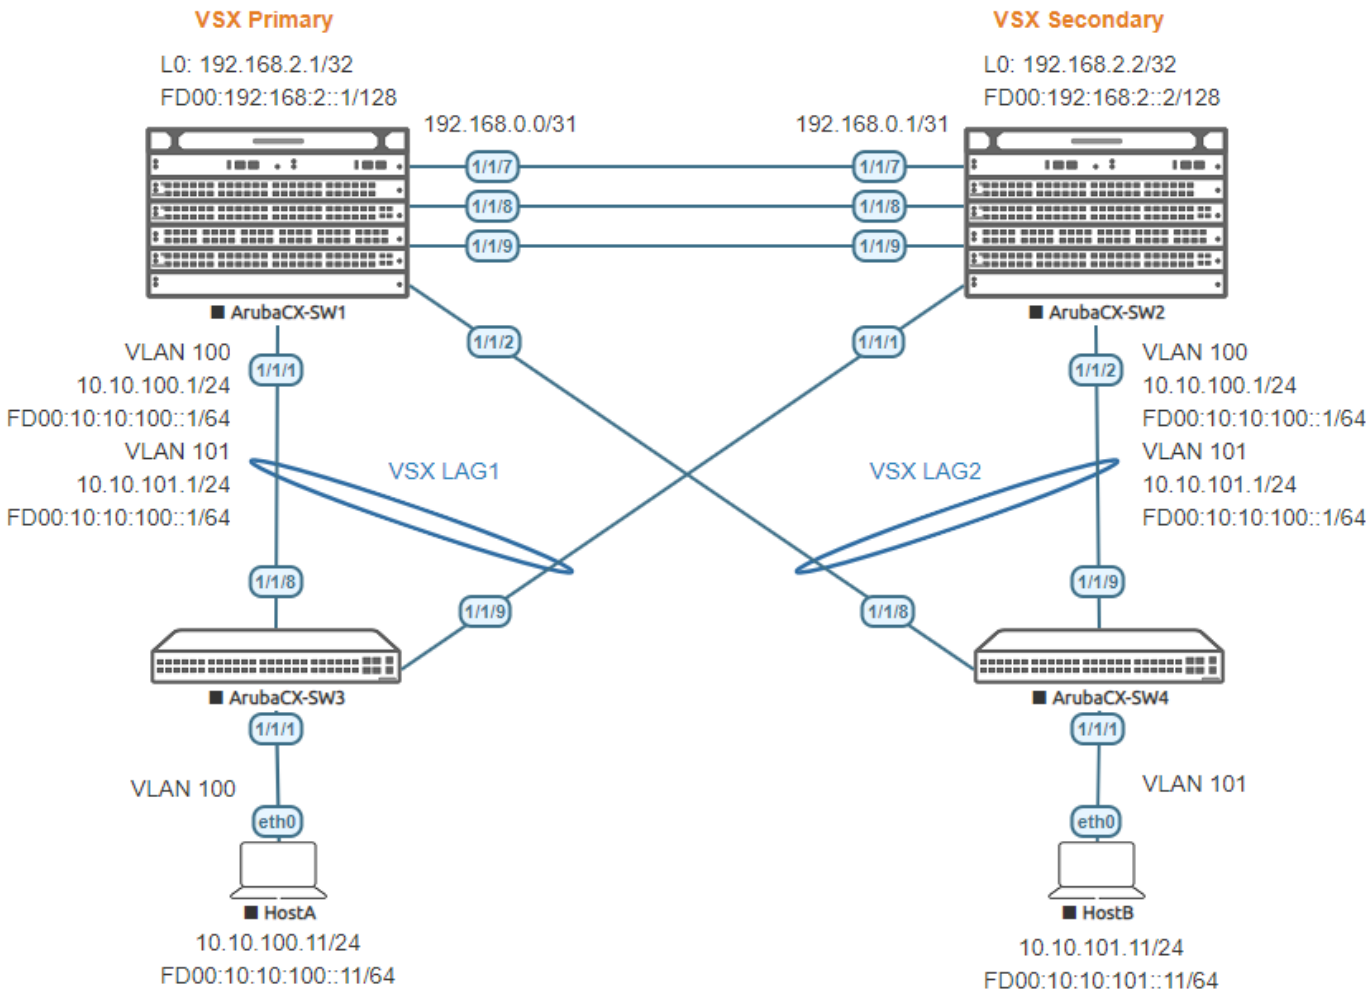

Lab Tasks

Task 1 – Lab setup

- In EVE-NG, i**mport 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 (4 AOS-CX switches and 2 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 from VSX-Lab1 (for initial copy/paste):

SW1

```
hostname SW1
!
vrf KA
!
vlan 1
vlan 100
    vsx-sync
interface mgmt
    no shutdown
    ip dhcp
interface lag 1 multi-chassis
    no shutdown
    description SW3 VSX LAG
    no routing
    vlan trunk native 1
    vlan trunk allowed 100
    lacp mode active
    lacp rate fast
interface lag 2 multi-chassis
    no shutdown
    description SW4 VSX LAG
    no routing
    vlan trunk native 1
    vlan trunk allowed 100
    lacp mode active
    lacp fallback
    lacp rate fast
interface lag 256
    no shutdown
    description ISL
    no routing
    vlan trunk native 1 tag
    vlan trunk allowed all
    lacp mode active
    lacp rate fast
interface 1/1/1
    no shutdown
    mtu 9100
    description to SW3
    lag 1
interface 1/1/2
    no shutdown
    mtu 9100
    description to SW4
    lag 2
interface 1/1/7
    no shutdown
    vrf attach KA
    description keepalive link
    ip address 192.168.0.0/31
interface 1/1/8
    no shutdown
    mtu 9198
    description ISL
    lag 256
interface 1/1/9
    no shutdown
    mtu 9198
    description ISL
    lag 256
vsx
    system-mac 02:01:00:00:01:00
    inter-switch-link lag 256
    role primary
    keepalive peer 192.168.0.1 source 192.168.0.0
vrf KA
    vsx-sync aaa acl-log-timer bfd-global bgp
control-plane-acls copp-policy dhcp-relay dhcp-
server dhcp-snooping dns icmp-tcp lldp loop-
protect-global mac-lockout mclag-interfaces
```

SW2

```
hostname SW2
!
vrf KA
!
vlan 1
vlan 100
    vsx-sync
interface mgmt
    no shutdown
    ip dhcp
interface lag 1 multi-chassis
    no shutdown
    description SW3 VSX LAG
    no routing
    vlan trunk native 1
    vlan trunk allowed 100
    lacp mode active
    lacp rate fast
interface lag 2 multi-chassis
    no shutdown
    description SW4 VSX LAG
    no routing
    vlan trunk native 1
    vlan trunk allowed 100
    lacp mode active
    lacp fallback
    lacp rate fast
interface lag 256
    no shutdown
    description ISL
    no routing
    vlan trunk native 1 tag
    vlan trunk allowed all
    lacp mode active
    lacp rate fast
interface 1/1/1
    no shutdown
    mtu 9100
    description to SW3
    lag 1
interface 1/1/2
    no shutdown
    mtu 9100
    description to SW4
    lag 2
interface 1/1/7
    no shutdown
    vrf attach KA
    description keepalive link
    ip address 192.168.0.1/31
interface 1/1/8
    no shutdown
    mtu 9198
    description ISL
    lag 256
interface 1/1/9
    no shutdown
    mtu 9198
    description ISL
    lag 256
vsx
    system-mac 02:01:00:00:01:00
    inter-switch-link lag 256
    role secondary
    keepalive peer 192.168.0.0 source 192.168.0.1
vrf KA
    vsx-sync aaa acl-log-timer bfd-global bgp
control-plane-acls copp-policy dhcp-relay dhcp-
server dhcp-snooping dns icmp-tcp lldp loop-
protect-global mac-lockout mclag-interfaces
```

```
neighbor ospf qos-global route-map sflow-global
snmp ssh stp-global time vsx-global
```

```
neighbor ospf qos-global route-map sflow-global
snmp ssh stp-global time vsx-global
```

SW3

```
hostname SW3
!
vlan 1,100
interface mgmt
    no shutdown
    ip dhcp
interface lag 1
    no shutdown
    no routing
    vlan trunk native 1
    vlan trunk allowed 100
    lacp mode active
    lacp rate fast
interface 1/1/1
    no shutdown
    no routing
    vlan access 100
interface 1/1/8
    no shutdown
    mtu 9100
    description to SW1
    lag 1
interface 1/1/9
    no shutdown
    mtu 9100
    description to SW2
    lag 1
```

SW4

```
hostname SW4
!
vlan 1,100
interface mgmt
    no shutdown
    ip dhcp
interface lag 1
    no shutdown
    no routing
    vlan trunk native 1
    vlan trunk allowed 100
    lacp mode active
    lacp rate fast
interface 1/1/1
    no shutdown
    no routing
    vlan access 100
interface 1/1/8
    no shutdown
    mtu 9100
    description to SW1
    lag 1
interface 1/1/9
    no shutdown
    mtu 9100
    description to SW2
    lag 1
```

From the baseline configuration of the end of VSX Lab1, the workaround of the VSX LAG limitation has to be applied for the AOS-CX Switch Simulator. The following steps are currently required:

SW1

```
SW1# conf
SW1(config)# interface 1/1/1-1/1/9
SW1(config-if-<1/1/1-1/1/9>)# shut
SW1(config-if-<1/1/1-1/1/9>)# interface 1/1/1
SW1(config-if)# no lag 1
SW1(config-if)# interface 1/1/2
SW1(config-if)# no lag 2
SW1(config-if)# end
SW1# wr mem
Copying configuration: [Success]
SW1# boot system
Checking if the configuration needs to be saved

This will reboot the entire switch and render it
unavailable until the process is complete.
Continue (y/n)? y
The system is going down for reboot.

 Apr 19 13:51:30 hpe-mgmtmd[9952]: RebootLibPh1:
Reboot reason: Reboot requested by user


<after reboot>

SW1# conf
SW1(config)# int 1/1/1-1/1/9
SW1(config-if-<1/1/1-1/1/9>)# no shut
SW1(config-if-<1/1/1-1/1/9>)# interface 1/1/1
SW1(config-if)# lag 1
SW1(config-if)# interface 1/1/2
SW1(config-if)# lag 2
SW1(config-if)# end
```

SW2

```
SW2# conf
SW2(config)# interface 1/1/1-1/1/9
SW2(config-if-<1/1/1-1/1/9>)# shut
SW2(config-if-<1/1/1-1/1/9>)# interface 1/1/1
SW2(config-if)# no lag 1
SW2(config-if)# interface 1/1/2
SW2(config-if)# no lag 2
SW2(config-if)# end
SW2# wr mem
Copying configuration: [Success]
SW2# boot system
Checking if the configuration needs to be saved

This will reboot the entire switch and render it
unavailable until the process is complete.
Continue (y/n)? y
The system is going down for reboot.

 Apr 19 13:52:22 hpe-mgmtmd[9987]: RebootLibPh1:
Reboot reason: Reboot requested by user


<after reboot>

SW2# conf
SW2(config)# int 1/1/1-1/1/9
SW2(config-if-<1/1/1-1/1/9>)# no shut
SW2(config-if-<1/1/1-1/1/9>)# interface 1/1/1
SW2(config-if)# lag 1
SW2(config-if)# interface 1/1/2
SW2(config-if)# lag 2
SW2(config-if)# end
```

```
SW1# wr mem
Copying configuration: [Success]
```

```
SW2# wr mem
Copying configuration: [Success]
```

No need to apply any workaround on SW3 and SW4.

- Verify the connectivity through LLDP neighbor information as follows:

SW1

```
SW1# show lldp neighbor-info

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

Total Neighbor Entries          : 5
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       08:00:09:5b:7e:2d  1/1/8                   to SW1                    120      SW3
1/1/2       08:00:09:ed:b5:6e  1/1/8                   to SW1                    120      SW4
1/1/7       08:00:09:54:97:83  1/1/7                   keepalive link            120      SW2
1/1/8       08:00:09:54:97:83  1/1/8                   ISL                       120      SW2
1/1/9       08:00:09:54:97:83  1/1/9                   ISL                       120      SW2
```

SW2

```
SW2# show lldp neighbor-info

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

Total Neighbor Entries          : 5
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       08:00:09:5b:7e:2d  1/1/9                   to SW2                    120      SW3
1/1/2       08:00:09:ed:b5:6e  1/1/9                   to SW2                    120      SW4
1/1/7       08:00:09:d7:5f:0f  1/1/7                   keepalive link            120      SW1
1/1/8       08:00:09:d7:5f:0f  1/1/8                   ISL                       120      SW1
1/1/9       08:00:09:d7:5f:0f  1/1/9                   ISL                       120      SW1
```

- Check VSX status:

SW1

```
SW1# 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          02:01:00:00:01:00   02:01:00:00:01:00
Platform            X86-64              X86-64
Software Version    Virtual.10.06.0110  Virtual.10.06.0110
Device Role         primary             secondary
```

SW2

```
SW2# 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          02:01:00:00:01:00   02:01:00:00:01:00
Platform            X86-64              X86-64
Software Version    Virtual.10.06.0110  Virtual.10.06.0110
Device Role         secondary           primary
```

- Check that LACP is collecting and distributing (flags should be ALFNCD or ASFNCD).

SW1 / SW2

```
SW1# 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
------------------------------------------------------------------------------
1/1/1      lag1(mc)   1      1     ASFNCD  02:01:00:00:01:00 65534  1    up
1/1/2      lag2(mc)   2      1     ASFNCD  02:01:00:00:01:00 65534  2    up
1/1/8      lag256     9      1     ASFNCD  08:00:09:d7:5f:0f 65534  256  up
1/1/9      lag256     10     1     ASFNCD  08:00:09:d7:5f:0f 65534  256  up


Partner details of all interfaces:
------------------------------------------------------------------------------
Intf       Aggr       Port   Port  State   System-ID         System Aggr
           Name       Id     Pri                              Pri    Key
------------------------------------------------------------------------------
1/1/1      lag1(mc)   9      1     ASFNCD  08:00:09:5b:7e:2d 65534  1
1/1/2      lag2(mc)   9      1     ASFNCD  08:00:09:ed:b5:6e 65534  1
1/1/8      lag256     9      1     ASFNCD  08:00:09:54:97:83 65534  256
1/1/9      lag256     10     1     ASFNCD  08:00:09:54:97:83 65534  256
```

- You may optionally check that HostA (10.10.100.11) can ping HostB being configured with an IP address in same subnet (example: in VSX-Lab1 hostB IP address was 10.10.100.12)

Task 2 – Configure L3 on VSX Cluster

Step #1: Disable IP ICMP redirect

Disable ip icmp redirect to avoid duplicate packet.

SW1	SW2
`SW1(config)# no ip icmp redirect`	synchronized

Note: This setting is synchronized on the VSX secondary only if the associated FeatureGroup has been configured to be vsx-synced in the VSX configuration.

Step #2: add VLAN101

Add VLAN101 on VSX primary and get it vsx-synced on VSX secondary.

SW1(config)#	SW2(config)#

`SW1(config)# vlan 101`
`SW1(config-vlan-101)# vsx-sync`
`SW1(config-vlan-101)# exit`
`SW1(config)# int lag 2 multi-chassis`
`SW1(config-lag-if)# vlan trunk allowed 101`
`SW1(config-lag-if)# show run current`
`interface lag 2 multi-chassis`
	synchronized
	
	`SW2# show run interface lag 2`
	`interface lag 2 multi-chassis`

```
    no shutdown
    description SW4 VSX LAG
    no routing
    vlan trunk native 1
    vlan trunk allowed 100-101
    lacp mode active
    lacp fallback
    lacp rate fast
```

```
    no shutdown
    description SW4 VSX LAG
    no routing
    vlan trunk native 1
    vlan trunk allowed 100-101
    lacp mode active
    lacp fallback
    lacp rate fast
    exit
```

Add VLAN101 on SW4

SW4(config)#

```
SW4(config)# vlan 101
SW4(config-vlan-101)# int lag 1
SW4(config-lag-if)# vlan trunk allowed 101
SW4(config-lag-if)# int 1/1/1
SW4(config-if)# vlan access 101
```

Step #3: SVI (VLAN L3 interface) configuration

Note: In the following Lab sections, dual stack IPv4+IPv6 is proposed. You may configure IPv4 only or IPv6 only if preferred.

The **best practice for SVI active-gateway** is to set the active-gateway Virtual IP and Virtual MAC on the VSX primary and get the value synchronized on the VSX secondary with vsx-sync command.

The **best practice for active-gateway VMAC** is to use the **same VMAC for all IPv4 SVIs**. The scope of this VMAC is purely link-local. If some servers or systems have dual-attachment to two different SVIs, and the system administrator would like to see distinct MAC addresses for the next-hops over these separate interfaces, then 16 VMACs are available. For dual-stack IPv4 and IPv6, 16 VMACs can be used for IPv4 and the same VMACs can be used for IPv6. Although these VMACs are optionally identical for non-VTEP scenario, they must be identical between IPv4 and IPv6 for VTEP case. The recommendation for the Best Practice is consequently to set the same VMAC for IPv4 active-gateway than for IPv6 active-gateway.

The **best practice for IP MTU** is to configure on all SVIs the matching size of the L2 MTU: IP MTU recommended value = 9100. This parameter must be identical and manually set on both VSX nodes.

The **best practice for DHCP relay** is to configure the ip helper-address on the VSX primary and let vsx-sync configuring the same on the VSX secondary. DHCP setting is skipped in this lab and will be addressed in another advanced VSX Lab.

SW1(config)#

```
interface vlan100
    vsx-sync active-gateways
    ip mtu 9100
    ip address 10.10.100.2/24
    active-gateway ip mac 12:01:00:00:01:00
    active-gateway ip 10.10.100.1
    ipv6 address fd00:10:10:100::2/64
    active-gateway ipv6 mac 12:01:00:00:01:00
    active-gateway ipv6 fd00:10:10:100::1
interface vlan101
    vsx-sync active-gateways
    ip mtu 9100
    ip address 10.10.101.2/24
    active-gateway ip mac 12:01:00:00:01:00
    active-gateway ip 10.10.101.1
    ipv6 address fd00:10:10:101::2/64
    active-gateway ipv6 mac 12:01:00:00:01:00
    active-gateway ipv6 fd00:10:10:101::1


SW1# sh run interface vlan 100
interface vlan100
    vsx-sync active-gateways
    ip address 10.10.100.2/24
    active-gateway ip mac 12:01:00:00:01:00
    active-gateway ip 10.10.100.1
```

SW2(config)#

```
interface vlan100

    ip mtu 9100
    ip address 10.10.100.3/24
                synchronized
    ipv6 address fd00:10:10:100::3/64


interface vlan101

   ip mtu 9100
    ip address 10.10.101.3/24
                synchronized
    ipv6 address fd00:10:10:101::3/64




SW2# sh run int vlan 100
interface vlan100
    vsx-sync active-gateways
    ip address 10.10.100.3/24
    active-gateway ip mac 12:01:00:00:01:00
    active-gateway ip 10.10.100.1
```

```
    ip mtu 9100
    ipv6 address fd00:10:10:100::2/64
    active-gateway ipv6 mac 12:01:00:00:01:00
    active-gateway ipv6 fd00:10:10:100::1
    exit
SW1# sh run interface vlan 101
interface vlan101
    vsx-sync active-gateways
    ip address 10.10.101.2/24
    active-gateway ip mac 12:01:00:00:01:00
    active-gateway ip 10.10.101.1
    ip mtu 9100
    ipv6 address fd00:10:10:101::2/64
    active-gateway ipv6 mac 12:01:00:00:01:00
    active-gateway ipv6 fd00:10:10:101::1
    exit


SW1# show ip interface vlan100

Interface vlan100 is up
 Admin state is up
 Hardware: Ethernet, MAC Address:08:00:09:d7:5f:0f
 IP MTU 9100
 IPv4 address 10.10.100.2/24
 active-gateway ip mac 12:01:00:00:01:00
 active-gateway ip 10.10.100.1
 L3 Counters: Rx Disabled, Tx Disabled
 Rx
          ucast: 0 packets, 0 bytes
          mcast: 0 packets, 0 bytes
 Tx
          ucast: 0 packets, 0 bytes
          mcast: 0 packets, 0 bytes

SW1# show ipv6 interface vlan100
Interface vlan100 is up
 Admin state is up
  IPv6 address:
    fd00:10:10:100::2/64 [VALID]
  IPv6 link-local address:
fe80::800:980:64d7:5f0f/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:ffd7:5f0f
   ff02::2
 IPv6 multicast (S,G) entries joined: none
 IPv6 MTU 9100
 IPv6 unicast reverse path forwarding: none
 IPv6 load sharing: none
 active-gateway ipv6 mac 12:01:00:00:01:00
 active-gateway ipv6 fd00:10:10:100::1
 L3 Counters: Rx Disabled, Tx Disabled
 Rx
          ucast: 0 packets, 0 bytes
          mcast: 0 packets, 0 bytes
 Tx
          ucast: 0 packets, 0 bytes
          mcast: 0 packets, 0 bytes
```

```
    ip mtu 9100
    ipv6 address fd00:10:10:100::3/64
    active-gateway ipv6 mac 12:01:00:00:01:00
    active-gateway ipv6 fd00:10:10:100::1
    exit
SW2# sh run int vlan 101
interface vlan101
    vsx-sync active-gateways
    ip address 10.10.101.3/24
    active-gateway ip mac 12:01:00:00:01:00
    active-gateway ip 10.10.101.1
    ip mtu 9100
    ipv6 address fd00:10:10:101::3/64
    active-gateway ipv6 mac 12:01:00:00:01:00
    active-gateway ipv6 fd00:10:10:101::1
    exit


SW2# sh ip int vlan100

Interface vlan100 is up
 Admin state is up
 Hardware: Ethernet, MAC Address: 08:00:09:54:97:83
 IP MTU 9100
 IPv4 address 10.10.100.3/24
 active-gateway ip mac 12:01:00:00:01:00
 active-gateway ip 10.10.100.1
 L3 Counters: Rx Disabled, Tx Disabled
 Rx
          ucast: 0 packets, 0 bytes
          mcast: 0 packets, 0 bytes
 Tx
          ucast: 0 packets, 0 bytes
          mcast: 0 packets, 0 bytes

SW2# sh ipv6 int vlan100
Interface vlan100 is up
 Admin state is up
  IPv6 address:
    fd00:10:10:100::3/64 [VALID]
  IPv6 link-local address:
fe80::800:980:6454:9783/64 [VALID]
 IPv6 virtual address configured: none
 IPv6 multicast routing: disable
 IPv6 Forwarding feature: enabled
 IPv6 multicast groups locally joined:
   ff02::1  ff02::1:ff00:3  ff02::1:ff54:9783
ff02::1:ff00:0
   ff02::2
 IPv6 multicast (S,G) entries joined: none
 IPv6 MTU 9100
 IPv6 unicast reverse path forwarding: none
 IPv6 load sharing: none
 active-gateway ipv6 mac 12:01:00:00:01:00
 active-gateway ipv6 fd00:10:10:100::1
 L3 Counters: Rx Disabled, Tx Disabled
 Rx
          ucast: 0 packets, 0 bytes
          mcast: 0 packets, 0 bytes
 Tx
          ucast: 0 packets, 0 bytes
          mcast: 0 packets, 0 bytes
```

Step #4: OSPF configuration

It is a **best practice to create a dedicated Transit VLAN** between the VSX primary and the VSX secondary to exchange routes information for subnets that are not attached to both VSX nodes (example: loopback addresses of each VSX node). This dedicated Transit VLAN (here VLAN 2) provides better control and will not carry user data traffic in nominal situation or very limited in case of east-west traffic between single-attached endpoints.

There are two strategies to inject endpoint subnets into the routing table: either through OSPF or through BGP.

- OSPF: Most of the Campus deployments use OSPF to exchange route information for end-devices. This is simple and can scale very well with appropriate usage of areas. This is the target of this current Lab guide.
- BGP: Lot of DC deployment use BGP as a routing protocol due to the usage of EVPN based VXLAN. Such a design is coming in the Campus as well. Also, for more complex and granular routing engineering, BGP communities and route-map can offer a level of control that OSPF can not provide. This can be exposed in a future white paper.

There are two options to inject end-user subnets into OSPF DataBase: using OSPF command on the SVI (VLAN L3 interface), or redistributing the connected into OSPF with route-map control. In this lab, it is proposed to use the OSPF command on SVI as offering a simpler configuration like for the area the subnets belongs to. More details on OSPF best practices can be found on IP routing configuration guide.

The **best practice for point-to-point interconnectivity subnet** is to use /31 subnet.

The **best practice for OSPF configuration** is to use vsx-sync ospf synchronization option and have OSPF parameters automatically synced on the VSX secondary. As shown on the configuration step, very few elements have to be configured on the secondary.

The **best practice for OSPF cost** is to have VSX primary <-> VSX secondary cost lower than external parallel path cost (like Core-1 <-> Core-2 cost), as it is frequent that the ISL bandwidth is higher than the available bandwidth through the external devices. In the lab, OSPF cost for Transit VLAN over ISL is set to 50 as an example. OSPF cost is synchronized from the VSX primary to the VSX secondary.

SW1(config)#	SW2(config)#
router ospf 1 router-id 192.168.2.1 max-metric router-lsa on-startup passive-interface default area 0	router ospf 1 router-id 192.168.2.2 synchronized
router ospfv3 1 router-id 192.168.2.1 max-metric router-lsa on-startup passive-interface default area 0	router ospfv3 1 router-id 192.168.2.2 synchronized
!	!
interface loopback 0 ip address 192.168.2.1/32 ipv6 address fd00:192:168:2::1/128 ip ospf 1 area 0.0.0.0 ipv6 ospfv3 1 area 0.0.0.0	interface loopback 0 ip address 192.168.2.2/32 ipv6 address fd00:192:168:2::2/128
!	! synchronized
vlan 1105 vsx-sync description TRANSIT VLAN	
interface vlan1105 ip address 192.168.4.0/31 ipv6 address fd00:192:168:4::/127 ip ospf 1 area 0 no ip ospf passive ip ospf cost 50 ip ospf network point-to-point ip ospf authentication message-digest ip ospf message-digest-key 1 md5 plaintext *yourpass* ipv6 ospfv3 1 area 0 no ipv6 ospfv3 passive ipv6 ospfv3 cost 50 ipv6 ospfv3 network point-to-point ipv6 ospfv3 authentication ipsec spi 256 sha1 plaintext *yourpass*	interface vlan1105 ip address 192.168.4.1/31 ipv6 address fd00:192:168:4::1/127 synchronized
!	!
interface vlan100 ip ospf 1 area 0 ipv6 ospfv3 1 area 0 interface vlan101	synchronized

```
    ip ospf 1 area 0
    ipv6 ospfv3 1 area 0
```

Verify OSPF adjacencies.

SW1 / SW2

```
SW1# show ip ospf neighbors
OSPF Process ID 1 VRF default
==============================

Total Number of Neighbors: 1

Neighbor ID      Priority  State             Nbr Address      Interface
-------------------------------------------------------------------------
192.168.2.2      n/a       FULL              192.168.4.1      vlan1105


SW1# show ipv6 ospfv3 neighbors
OSPFv3 Process ID 1 VRF default
================================

Total Number of Neighbors: 1

Neighbor ID      Priority  State             Interface
-----------------------------------------------------
192.168.2.2      n/a       FULL              vlan1105
  Neighbor address fe80::800:984:5154:9783
```

For the IPv6 OSPF VSX peer, identify the IPv6 Link-Local address related to the Transit VLAN1105:

SW1 / SW2

```
SW1# show ipv6 interface brief vsx-peer
IPv6 Interface Status for VRF "default"
Interface        Link-local Address/IPv6 Address        Interface Status
                                                        link/admin
1/1/3                                                   down/ Admin state is down

1/1/4                                                   down/ Admin state is down

1/1/5                                                   down/ Admin state is down

1/1/6                                                   down/ Admin state is down

1/1/7                                                   up/ Admin state is up

loopback0        fe80::800:9b0:54:9783/64               up/ Admin state is up
                 fd00:192:168:2::2/128

vlan100          fe80::800:980:6454:9783/64             up/ Admin state is up
                 fd00:10:10:100::3/64

vlan101          fe80::800:980:6554:9783/64             up/ Admin state is up
                 fd00:10:10:101::3/64

vlan1105         fe80::800:984:5154:9783/64             up/ Admin state is up
                 fd00:192:168:4::1/127
```

Check the routing table (verify that Loopback of the VSX peer is learnt from OSPF).

SW1 / SW2

```
SW1# show ip route

Displaying ipv4 routes selected for forwarding

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

10.10.100.0/24, vrf default
       via  vlan100,  [0/0],  connected
10.10.100.2/32, vrf default
       via  vlan100,  [0/0],  local
10.10.101.0/24, vrf default
       via  vlan101,  [0/0],  connected
10.10.101.2/32, vrf default
       via  vlan101,  [0/0],  local
```

```
192.168.0.0/31, vrf default
       via  1/1/7,  [0/0],  connected
192.168.0.0/32, vrf default
       via  1/1/7,  [0/0],  local
192.168.2.1/32, vrf default
       via  loopback0,  [0/0],  local
192.168.2.2/32, vrf default
       via  192.168.4.1,  [110/50],  ospf
192.168.4.0/31, vrf default
       via  vlan1105,  [0/0],  connected
192.168.4.0/32, vrf default
       via  vlan1105,  [0/0],  local


SW1# show ipv6 route

Displaying ipv6 routes selected for forwarding

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

fd00:10:10:100::/64, vrf default
       via  vlan100,  [0/0],  connected
fd00:10:10:100::2/128, vrf default
       via  vlan100,  [0/0],  local
fd00:10:10:101::/64, vrf default
       via  vlan101,  [0/0],  connected
fd00:10:10:101::2/128, vrf default
       via  vlan101,  [0/0],  local
fd00:192:168:2::1/128, vrf default
       via  loopback0,  [0/0],  local
fd00:192:168:2::2/128, vrf default
       via  fe80::800:984:5154:9783%vlan1105,  [110/50],  ospf
fd00:192:168:4::/127, vrf default
       via  vlan1105,  [0/0],  connected
fd00:192:168:4::/128, vrf default
       via  vlan1105,  [0/0],  local
```

The main configuration for Layer3 on VSX cluster is completed. More advanced configuration might be exposed in other labs, showing VSX interaction with other devices like core devices or features like BGP.

Task 3 - Resiliency tests

IMPORTANT: The CX Simulator does not sense the state of the interfaces. It means that if the interface of the neighboring switch is shutdown, the local facing interface is not teared down and stays up. In other words, the interface state is not reflected between the CX neighbors. Consequently, when performing resiliency tests in CX Simulator Labs, it is recommended to:

- shutdown both ends of a link in a coordinated manner
- or when the link is part of a LACP LAG, use LACP short timer to let LACP protocol to unselect the interface.

This note is specific to CX Simulator and, in production, default LACP rate (slow) is used for physical CX switches.

In the previous VSX-Lab1, SW1/SW2/SW3/SW4 switches were configured with short **LACP timer** on all LAG interfaces.

As a reference, configuration of SW1/SW2/SW3/SW4 should look like:

SW1	SW2

hostname SW1
!
no ip icmp redirect
vrf KA
!
vlan 1
vlan 100
  vsx-sync
vlan 101
  vsx-sync
vlan 1105
  vsx-sync
  description TRANSIT VLAN
	hostname SW2
	!
	no ip icmp redirect
	vrf KA
	!
	vlan 1
	vlan 100
	  vsx-sync
	vlan 101
	  vsx-sync
	vlan 1105
	  vsx-sync
	  description TRANSIT VLAN

```
interface mgmt
    no shutdown
    ip dhcp
interface lag 1 multi-chassis
    no shutdown
    description SW3 VSX LAG
    no routing
    vlan trunk native 1
    vlan trunk allowed 100
    lacp mode active
    lacp rate fast
interface lag 2 multi-chassis
    no shutdown
    description SW4 VSX LAG
    no routing
    vlan trunk native 1
    vlan trunk allowed 100-101
    lacp mode active
    lacp fallback
    lacp rate fast
interface lag 256
    no shutdown
    description ISL
    no routing
    vlan trunk native 1 tag
    vlan trunk allowed all
    lacp mode active
    lacp rate fast
interface 1/1/1
    no shutdown
    mtu 9100
    description to SW3
    lag 1
interface 1/1/2
    no shutdown
    mtu 9100
    description to SW4
    lag 2
interface 1/1/7
    no shutdown
    vrf attach KA
    description keepalive link
    ip address 192.168.0.0/31
interface 1/1/8
    no shutdown
    mtu 9198
    description ISL
    lag 256
interface 1/1/9
    no shutdown
    mtu 9198
    description ISL
    lag 256
interface loopback 0
    ip address 192.168.2.1/32
    ipv6 address fd00:192:168:2::1/128
    ip ospf 1 area 0.0.0.0
    ipv6 ospfv3 1 area 0.0.0.0
interface vlan 100
    vsx-sync active-gateways
    ip mtu 9100
    ip address 10.10.100.2/24
    active-gateway ip mac 12:01:00:00:01:00
    active-gateway ip 10.10.100.1
    ipv6 address fd00:10:10:100::2/64
    active-gateway ipv6 mac 12:01:00:00:01:00
    active-gateway ipv6 fd00:10:10:100::1
    ip ospf 1 area 0.0.0.0
    ipv6 ospfv3 1 area 0.0.0.0
interface vlan 101
    vsx-sync active-gateways
    ip mtu 9100
    ip address 10.10.101.2/24
```

```
interface mgmt
    no shutdown
    ip dhcp
interface lag 1 multi-chassis
    no shutdown
    description SW3 VSX LAG
    no routing
    vlan trunk native 1
    vlan trunk allowed 100
    lacp mode active
    lacp rate fast
interface lag 2 multi-chassis
    no shutdown
    description SW4 VSX LAG
    no routing
    vlan trunk native 1
    vlan trunk allowed 100-101
    lacp mode active
    lacp fallback
    lacp rate fast
interface lag 256
    no shutdown
    description ISL
    no routing
    vlan trunk native 1 tag
    vlan trunk allowed all
    lacp mode active
    lacp rate fast
interface 1/1/1
    no shutdown
    mtu 9100
    description to SW3
    lag 1
interface 1/1/2
    no shutdown
    mtu 9100
    description to SW4
    lag 2
interface 1/1/7
    no shutdown
    vrf attach KA
    description keepalive link
    ip address 192.168.0.1/31
interface 1/1/8
    no shutdown
    mtu 9198
    description ISL
    lag 256
interface 1/1/9
    no shutdown
    mtu 9198
    description ISL
    lag 256
interface loopback 0
    ip address 192.168.2.2/32
    ipv6 address fd00:192:168:2::2/128
    ip ospf 1 area 0.0.0.0
    ipv6 ospfv3 1 area 0.0.0.0
interface vlan 100
    vsx-sync active-gateways
    ip mtu 9100
    ip address 10.10.100.3/24
    active-gateway ip mac 12:01:00:00:01:00
    active-gateway ip 10.10.100.1
    ipv6 address fd00:10:10:100::3/64
    active-gateway ipv6 mac 12:01:00:00:01:00
    active-gateway ipv6 fd00:10:10:100::1
    ip ospf 1 area 0.0.0.0
    ipv6 ospfv3 1 area 0.0.0.0
interface vlan 101
    vsx-sync active-gateways
    ip mtu 9100
    ip address 10.10.101.3/24
```

```
    active-gateway ip mac 12:01:00:00:01:00
    active-gateway ip 10.10.101.1
    ipv6 address fd00:10:10:101::2/64
    active-gateway ipv6 mac 12:01:00:00:01:00
    active-gateway ipv6 fd00:10:10:101::1
    ip ospf 1 area 0.0.0.0
    ipv6 ospfv3 1 area 0.0.0.0
interface vlan 1105
    ip address 192.168.4.0/31
    ipv6 address fd00:192:168:4::/127
    ip ospf 1 area 0.0.0.0
    no ip ospf passive
    ip ospf cost 50
    ip ospf network point-to-point
    ip ospf authentication message-digest
    ip ospf message-digest-key 1 md5 ciphertext
AQBapUbZyuMyDkoDN0zeQbI8qY0p5vpa77xnpPQEngEkpWjWBQA
AAIouj7OC
    ipv6 ospfv3 1 area 0.0.0.0
    no ipv6 ospfv3 passive
    ipv6 ospfv3 cost 50
    ipv6 ospfv3 network point-to-point
    ipv6 ospfv3 authentication ipsec spi 256 sha1
ciphertext
AQBapUbZyuMyDkoDN0zeQbI8qY0p5vpa77xnpPQEngEkpWjWBQA
AAIouj7OC
vsx
    system-mac 02:01:00:00:01:00
    inter-switch-link lag 256
    role primary
    keepalive peer 192.168.0.1 source 192.168.0.0
vrf KA
    vsx-sync aaa acl-log-timer bfd-global bgp
control-plane-acls copp-policy dhcp-relay dhcp-
server dhcp-snooping dns icmp-tcp lldp loop-
protect-global mac-lockout mclag-interfaces
neighbor ospf qos-global route-map sflow-global
snmp ssh stp-global time vsx-global
!
router ospf 1
    router-id 192.168.2.1
    max-metric router-lsa on-startup
    passive-interface default
    area 0.0.0.0
router ospfv3 1
    router-id 192.168.2.1
    max-metric router-lsa on-startup
    passive-interface default
    area 0.0.0.0
```

```
    active-gateway ip mac 12:01:00:00:01:00
    active-gateway ip 10.10.101.1
    ipv6 address fd00:10:10:101::3/64
    active-gateway ipv6 mac 12:01:00:00:01:00
    active-gateway ipv6 fd00:10:10:101::1
    ip ospf 1 area 0.0.0.0
    ipv6 ospfv3 1 area 0.0.0.0
interface vlan 1105
    ip address 192.168.4.1/31
    ipv6 address fd00:192:168:4::1/127
    ip ospf 1 area 0.0.0.0
    no ip ospf passive
    ip ospf cost 50
    ip ospf network point-to-point
    ip ospf authentication message-digest
    ip ospf message-digest-key 1 md5 ciphertext
AQBapUbZyuMyDkoDN0zeQbI8qY0p5vpa77xnpPQEngEkpWjWBQA
AAIouj7OC
    ipv6 ospfv3 1 area 0.0.0.0
    no ipv6 ospfv3 passive
    ipv6 ospfv3 cost 50
    ipv6 ospfv3 network point-to-point
    ipv6 ospfv3 authentication ipsec spi 256 sha1
ciphertext
AQBapUbZyuMyDkoDN0zeQbI8qY0p5vpa77xnpPQEngEkpWjWBQA
AAIouj7OC
vsx
    system-mac 02:01:00:00:01:00
    inter-switch-link lag 256
    role secondary
    keepalive peer 192.168.0.0 source 192.168.0.1
vrf KA
    vsx-sync aaa acl-log-timer bfd-global bgp
control-plane-acls copp-policy dhcp-relay dhcp-
server dhcp-snooping dns icmp-tcp lldp loop-
protect-global mac-lockout mclag-interfaces
neighbor ospf qos-global route-map sflow-global
snmp ssh stp-global time vsx-global
!
router ospf 1
    router-id 192.168.2.2
    max-metric router-lsa on-startup
    passive-interface default
    area 0.0.0.0
router ospfv3 1
    router-id 192.168.2.2
    max-metric router-lsa on-startup
    passive-interface default
    area 0.0.0.0
```

SW3

```
hostname SW3
!
vlan 1,100
interface mgmt
    no shutdown
    ip dhcp
interface lag 1
    no shutdown
    no routing
    vlan trunk native 1
    vlan trunk allowed 100
    lacp mode active
    lacp rate fast
interface 1/1/1
    no shutdown
    no routing
    vlan access 100
interface 1/1/8
    no shutdown
    mtu 9100
    description to SW1
    lag 1
interface 1/1/9
```

SW4

```
hostname SW4
!
vlan 1,100-101
interface mgmt
    no shutdown
    ip dhcp
interface lag 1
    no shutdown
    no routing
    vlan trunk native 1
    vlan trunk allowed 100-101
    lacp mode active
    lacp rate fast
interface 1/1/1
    no shutdown
    no routing
    vlan access 101
interface 1/1/8
    no shutdown
    mtu 9100
    description to SW1
    lag 1
interface 1/1/9
```

```
    no shutdown
    mtu 9100
    description to SW2
    lag 1
```

```
    no shutdown
    mtu 9100
    description to SW2
    lag 1
```

Test #1: Layer3 connectivity between HostA and HostB

Set-up IPv4/IPv6 addresses on HostA (10.10.100.11/24) and HostB (10.10.101.11/24):

HostA

```
VPCS> ip 10.10.100.11/24 10.10.100.1
Checking for duplicate address...
VPCS : 10.10.100.11 255.255.255.0 gateway
10.10.100.1

VPCS> ip fd00:10:10:100::11/64 fd00:10:10:100::1
PC1 : fd00:10:10:100::11/64

VPCS> show ip

NAME        : VPCS[1]
IP/MASK     : 10.10.100.11/24
GATEWAY     : 10.10.100.1
DNS         :
MAC         : 00:50:79:66:68:07
LPORT       : 20000
RHOST:PORT  : 127.0.0.1:30000
MTU         : 1500

VPCS> show ipv6

NAME              : VPCS[1]
LINK-LOCAL SCOPE  : fe80::250:79ff:fe66:6807/64
GLOBAL SCOPE      : fd00:10:10:100::11/64
DNS               :
ROUTER LINK-LAYER :
MAC               : 00:50:79:66:68:07
LPORT             : 20000
RHOST:PORT        : 127.0.0.1:30000
MTU:              : 1500
```

HostB

```
VPCS> ip 10.10.101.11/24 10.10.101.1
Checking for duplicate address...
VPCS : 10.10.101.11 255.255.255.0 gateway
10.10.101.1

VPCS> ip fd00:10:10:101::11/64 fd00:10:10:101::1
PC1 : fd00:10:10:101::11/64

VPCS> show ip

NAME        : VPCS[1]
IP/MASK     : 10.10.101.11/24
GATEWAY     : 10.10.101.1
DNS         :
MAC         : 00:50:79:66:68:05
LPORT       : 20000
RHOST:PORT  : 127.0.0.1:30000
MTU         : 1500

VPCS> show ipv6

NAME              : VPCS[1]
LINK-LOCAL SCOPE  : fe80::250:79ff:fe66:6805/64
GLOBAL SCOPE      : fd00:10:10:101::11/64
DNS               :
ROUTER LINK-LAYER :
MAC               : 00:50:79:66:68:05
LPORT             : 20000
RHOST:PORT        : 127.0.0.1:30000
MTU:              : 1500
```

Ping HostB from HostA

HostA

```
VPCS> ping 10.10.101.11

84 bytes from 10.10.101.11 icmp_seq=1 ttl=63 time=17.344 ms
84 bytes from 10.10.101.11 icmp_seq=2 ttl=63 time=5.126 ms
84 bytes from 10.10.101.11 icmp_seq=3 ttl=63 time=4.099 ms
84 bytes from 10.10.101.11 icmp_seq=4 ttl=63 time=3.885 ms
84 bytes from 10.10.101.11 icmp_seq=5 ttl=63 time=4.252 ms
```

Please note the higher response time for the first packet which corresponds to the ARP requests performed in both subnets.

Ping SW2 L0 from HostA

HostA

```
VPCS> ping 192.168.2.2

84 bytes from 192.168.2.2 icmp_seq=1 ttl=64 time=3.253 ms
84 bytes from 192.168.2.2 icmp_seq=2 ttl=64 time=3.121 ms
84 bytes from 192.168.2.2 icmp_seq=3 ttl=64 time=3.089 ms
84 bytes from 192.168.2.2 icmp_seq=4 ttl=64 time=3.465 ms
84 bytes from 192.168.2.2 icmp_seq=5 ttl=64 time=3.546 ms
```

As VPCS does not provide ping6 capability, use ping6 from the switch to the Host. Ping6 HostA from SW1.

SW1

```
SW1# ping6 fd00:10:10:100::11 source fd00:192:168:2::1
PING fd00:10:10:100::11(fd00:10:10:100::11) from fd00:192:168:2::1 : 100 data bytes
108 bytes from fd00:10:10:100::11: icmp_seq=1 ttl=64 time=1003 ms
108 bytes from fd00:10:10:100::11: icmp_seq=2 ttl=64 time=1002 ms
```

```
108 bytes from fd00:10:10:100::11: icmp_seq=3 ttl=64 time=1002 ms
108 bytes from fd00:10:10:100::11: icmp_seq=4 ttl=64 time=1017 ms
108 bytes from fd00:10:10:100::11: icmp_seq=5 ttl=64 time=1012 ms

--- fd00:10:10:100::11 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4029ms
rtt min/avg/max/mdev = 1002.256/1007.666/1017.386/6.114 ms
```

Test #2: resiliency tests while disconnecting the VSX primary

Right-click on SW3, select Graph Menu (EVE-NG pro only) and click on 1/1/8.

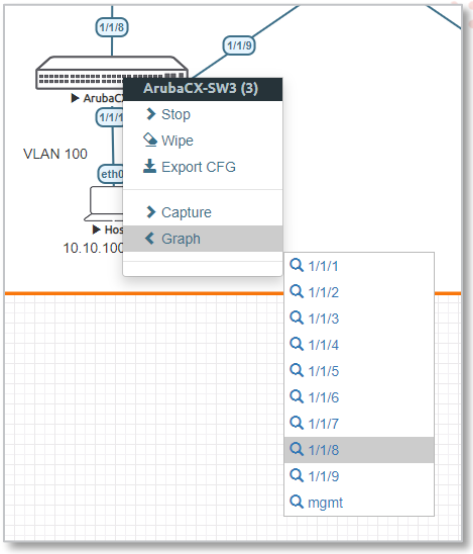

A Graph of the link utilization will appear. Repeat the same action for 1/1/9. You should have the following graphs:

Start continuous ping to HostB from HostA with large datagram size like: ping 10.10.101.11 –l 1400 -t

You should now see links being loaded by this ICMP traffic, here interface 1/1/8 loaded in egress direction, and interface 1/1/9 loaded in ingress direction.

The egress port being loaded on SW3 is determined by the hashing algorithm on SW3, whereas the ingress port being loaded on SW3 is determined by the hashing mechanism performed on egress direction from SW4. The VSX nodes will forward the traffic on the local VSX node in nominal situation, and will not use the ISL if all downstream switches are dual-attached to VSX primary and to VSX secondary.

Shutdown all interfaces on the VSX primary SW1.

SW1

```
SW1# conf
SW1(config)# int 1/1/1-1/1/9
SW1(config-if-<1/1/1-1/1/9>)# shut
```

Check the impact on the ping:

```
1428 bytes from 10.10.101.11 icmp_seq=431 ttl=63 time=4.459 ms
1428 bytes from 10.10.101.11 icmp_seq=432 ttl=63 time=10.504 ms
1428 bytes from 10.10.101.11 icmp_seq=433 ttl=63 time=4.247 ms
1428 bytes from 10.10.101.11 icmp_seq=434 ttl=63 time=4.388 ms
1428 bytes from 10.10.101.11 icmp_seq=435 ttl=63 time=4.225 ms
1428 bytes from 10.10.101.11 icmp_seq=436 ttl=63 time=4.157 ms
1428 bytes from 10.10.101.11 icmp_seq=437 ttl=63 time=4.774 ms
1428 bytes from 10.10.101.11 icmp_seq=438 ttl=63 time=4.712 ms
1428 bytes from 10.10.101.11 icmp_seq=439 ttl=63 time=3.792 ms
1428 bytes from 10.10.101.11 icmp_seq=440 ttl=63 time=4.285 ms
10.10.101.11 icmp_seq=441 timeout
10.10.101.11 icmp_seq=442 timeout
10.10.101.11 icmp_seq=443 timeout
1428 bytes from 10.10.101.11 icmp_seq=444 ttl=63 time=4.312 ms
1428 bytes from 10.10.101.11 icmp_seq=445 ttl=63 time=8.074 ms
1428 bytes from 10.10.101.11 icmp_seq=446 ttl=63 time=4.311 ms
1428 bytes from 10.10.101.11 icmp_seq=447 ttl=63 time=3.782 ms
1428 bytes from 10.10.101.11 icmp_seq=448 ttl=63 time=12.451 ms
1428 bytes from 10.10.101.11 icmp_seq=449 ttl=63 time=4.220 ms
1428 bytes from 10.10.101.11 icmp_seq=450 ttl=63 time=4.974 ms
```

You will see about 3 seconds' outage on the Simulator. This duration corresponds to the LACP-block detection performed on SW1 to use the link to SW2 instead of the link to SW1, and the time for VSX secondary to forward to SW4. In production network, recovery for such power failure would take less than 200 milliseconds as the link state detection will be much faster.

Once the VSX primary is disconnected/isolated, check VSX status on the VSX secondary and verify the status of the VSX LAGs:

SW2

```
SW2# show vsx status
VSX Operational State
---------------------
  ISL channel             : Out-Of-Sync
  ISL mgmt channel        : inter_switch_link_down
  Config Sync Status      : Out-Of-Sync
  NAE                     : peer_unreachable
  HTTPS Server            : peer_unreachable

Attribute           Local               Peer
------------        --------            --------
ISL link            lag256
ISL version         2
System MAC          02:01:00:00:01:00
Platform            X86-64
Software Version    Virtual.10.06.0110
Device Role         secondary


SW2# show lacp int

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
-------------------------------------------------------------------------------
1/1/1      lag1(mc)   1001  1     ASFNCD  02:01:00:00:01:00 65534  1    up
1/1/2      lag2(mc)   1002  1     ASFNCD  02:01:00:00:01:00 65534  2    up
1/1/8      lag256     9     1     ASFOE   08:00:09:54:97:83 65534  256  lacp-block
1/1/9      lag256     10    1     ASFOE   08:00:09:54:97:83 65534  256  lacp-block


Partner details of all interfaces:
-------------------------------------------------------------------------------
Intf       Aggr       Port  Port  State   System-ID          System Aggr
           Name       Id    Pri                              Pri    Key
-------------------------------------------------------------------------------
1/1/1      lag1(mc)   10    1     ASFNCD  08:00:09:5b:7e:2d 65534  1
1/1/2      lag2(mc)   10    1     ASFNCD  08:00:09:ed:b5:6e 65534  1
1/1/8      lag256     0     0     PLFOEX  00:00:00:00:00:00 0      0
1/1/9      lag256     0     0     PLFOEX  00:00:00:00:00:00 0      0
```

On the graphs, you should see traffic transition: interface 1/1/8 (facing SW1) is no longer used, where interface 1/1/9 is sending and receiving ICMP traffic.

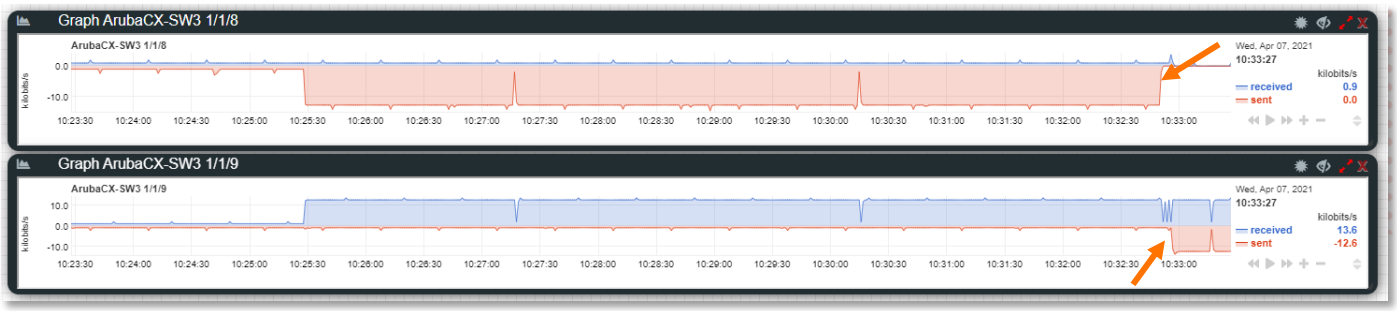

Restore VSX primary:

SW1

```
SW1(config-if-<1/1/1-1/1/9>)# no shut
```

Check on the client the impact:

```
1428 bytes from 10.10.101.11 icmp_seq=628 ttl=63 time=10.638 ms
1428 bytes from 10.10.101.11 icmp_seq=629 ttl=63 time=4.716 ms
1428 bytes from 10.10.101.11 icmp_seq=630 ttl=63 time=4.466 ms
1428 bytes from 10.10.101.11 icmp_seq=631 ttl=63 time=4.867 ms
1428 bytes from 10.10.101.11 icmp_seq=632 ttl=63 time=4.402 ms
1428 bytes from 10.10.101.11 icmp_seq=633 ttl=63 time=4.105 ms
1428 bytes from 10.10.101.11 icmp_seq=634 ttl=63 time=3.592 ms
10.10.101.11 icmp_seq=635 timeout
10.10.101.11 icmp_seq=636 timeout
10.10.101.11 icmp_seq=637 timeout
10.10.101.11 icmp_seq=638 timeout
1428 bytes from 10.10.101.11 icmp_seq=639 ttl=63 time=29.107 ms
1428 bytes from 10.10.101.11 icmp_seq=640 ttl=63 time=3.578 ms
1428 bytes from 10.10.101.11 icmp_seq=641 ttl=63 time=4.506 ms
1428 bytes from 10.10.101.11 icmp_seq=642 ttl=63 time=3.621 ms
1428 bytes from 10.10.101.11 icmp_seq=643 ttl=63 time=4.251 ms
1428 bytes from 10.10.101.11 icmp_seq=644 ttl=63 time=3.809 ms
```

You should see about 3~4 seconds outage. This is the time required for SW1 to restore LACP and re-ARP the destination endpoint hosts.

SW1

```
SW1# 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          02:01:00:00:01:00   02:01:00:00:01:00
Platform            X86-64              X86-64
Software Version    Virtual.10.06.0110  Virtual.10.06.0110
Device Role         primary             secondary
```

Check the VSX secondary status and in particular the linkup-delay timer status:

SW2

```
SW1# 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          02:01:00:00:01:00   02:01:00:00:01:00
Platform            X86-64              X86-64
Software Version    Virtual.10.06.0110  Virtual.10.06.0110
Device Role         primary             secondary


SW2# show vsx status linkup-delay
Configured linkup delay-timer                                : 180 seconds
Initial sync status                                          : Completed
Delay timer status                                           : Running
Linkup Delay time left                                       : 2 minutes 57 seconds
Interfaces that will be brought up after delay timer expires : lag1-lag2
Interfaces that are excluded from delay timer                :
```

The VSX secondary has to join back the VSX cluster, and its VSX LAGs are blocked during the linkup-delay timer. On the graphs you should see interface 1/1/9 not used for about one minute (during linkup delay timer).

Before linkup-delay timer expires, check on the VSX secondary the status of the IP interfaces:

SW2

```
SW2# show ip interface brief
Interface         IP Address            Interface Status
                                          link/admin
1/1/3             No Address              down/down

1/1/4             No Address              down/down

1/1/5             No Address              down/down

1/1/6             No Address              down/down

1/1/7             192.168.0.1/31          up/up
```

```
loopback0         192.168.2.2/32            up/up

vlan100           10.10.100.3/24            down/up

vlan101           10.10.101.3/24            down/up

vlan1105          192.168.4.1/31            up/up


SW2# show ip route

Displaying ipv4 routes selected for forwarding

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

192.168.0.0/31, vrf default
       via  1/1/7,  [0/0],  connected
192.168.0.1/32, vrf default
       via  1/1/7,  [0/0],  local
192.168.2.2/32, vrf default
       via  loopback0,  [0/0],  local
192.168.4.0/31, vrf default
       via  vlan1105,  [0/0],  connected
192.168.4.1/32, vrf default
       via  vlan1105,  [0/0],  local


SW2# show ipv6 interface brief
IPv6 Interface Status for VRF "default"
Interface         Link-local Address/IPv6 Address         Interface Status
                                                          link/admin
1/1/3                                                     down/ Admin state is down

1/1/4                                                     down/ Admin state is down

1/1/5                                                     down/ Admin state is down

1/1/6                                                     down/ Admin state is down

1/1/7                                                     up/ Admin state is up

loopback0         fe80::800:9b0:54:9783/64                up/ Admin state is up
                  fd00:192:168:2::2/128

vlan100           fe80::800:980:6454:9783/64              down/ Admin state is up
                  fd00:10:10:100::3/64

vlan101           fe80::800:980:6554:9783/64              down/ Admin state is up
                  fd00:10:10:101::3/64
```

Note: During linkup-delay timer, the SVI (L3 VLAN interfaces) that are associated to the VLANs that are carried over the VSX LAGs, are shutdown. Inter-VLAN routing is performed by the VSX primary during that linkup-delay timer in this scenario (VSX secondary joining the VSX cluster). This may impact any routing protocol for upstream core if any (not demonstrated in this lab but in another VSX lab), as these connected subnet won't be in the routing table during the linkup-delay timer on the VSX secondary.

After linkup-delay timer expires:

SW2

```
SW2# show ip interface brief
Interface          IP Address             Interface Status
                                            link/admin
1/1/3              No Address               down/down

1/1/4              No Address               down/down

1/1/5              No Address               down/down

1/1/6              No Address               down/down
```

```
1/1/7             192.168.0.1/31            up/up

loopback0         192.168.2.2/32            up/up

vlan100           10.10.100.3/24            up/up

vlan101           10.10.101.3/24            up/up

vlan1105          192.168.4.1/31            up/up


SW2# show ip route

Displaying ipv4 routes selected for forwarding

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

10.10.100.0/24, vrf default
        via  vlan100,  [0/0],  connected
10.10.100.3/32, vrf default
        via  vlan100,  [0/0],  local
10.10.101.0/24, vrf default
        via  vlan101,  [0/0],  connected
10.10.101.3/32, vrf default
        via  vlan101,  [0/0],  local
192.168.0.0/31, vrf default
        via  1/1/7,  [0/0],  connected
192.168.0.1/32, vrf default
        via  1/1/7,  [0/0],  local
192.168.2.1/32, vrf default
        via  192.168.4.0,  [110/50],  ospf
192.168.2.2/32, vrf default
        via  loopback0,  [0/0],  local
192.168.4.0/31, vrf default
        via  vlan1105,  [0/0],  connected
192.168.4.1/32, vrf default
        via  vlan1105,  [0/0],  local


SW2# show ipv6 interface brief
IPv6 Interface Status for VRF "default"
Interface         Link-local Address/IPv6 Address         Interface Status
                                                          link/admin
1/1/3                                                     down/ Admin state is down

1/1/4                                                     down/ Admin state is down

1/1/5                                                     down/ Admin state is down

1/1/6                                                     down/ Admin state is down

1/1/7                                                     up/ Admin state is up

loopback0         fe80::800:9b0:54:9783/64                up/ Admin state is up
                  fd00:192:168:2::2/128

vlan100           fe80::800:980:6454:9783/64              up/ Admin state is up
                  fd00:10:10:100::3/64

vlan101           fe80::800:980:6554:9783/64              up/ Admin state is up
                  fd00:10:10:101::3/64

vlan1105          fe80::800:984:5154:9783/64              up/ Admin state is up
                  fd00:192:168:4::1/127
```

The VSX secondary is back to nominal routing and forwarding as shown on interfaces 1/1/8 and 1/1/9 of SW3:

Test #3: resiliency tests during power-off of the VSX primary

This particular sequence (isolating VSX primary and restoring primary's links without reboot) induces that the VSX secondary has to join back the VSX cluster as this event is considered as a VSX split: indeed, the VSX primary did not reboot.

In order to simulate a power-off, simply STOP SW1 from EVE-NG:

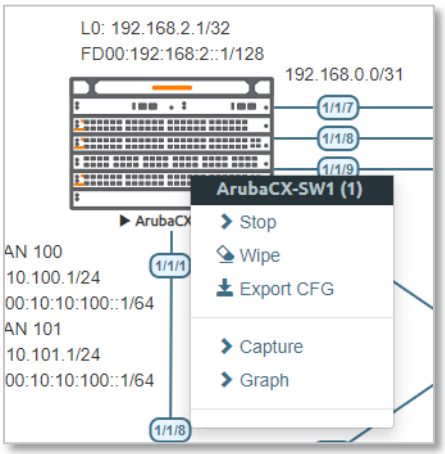

You should see the same transition than before.

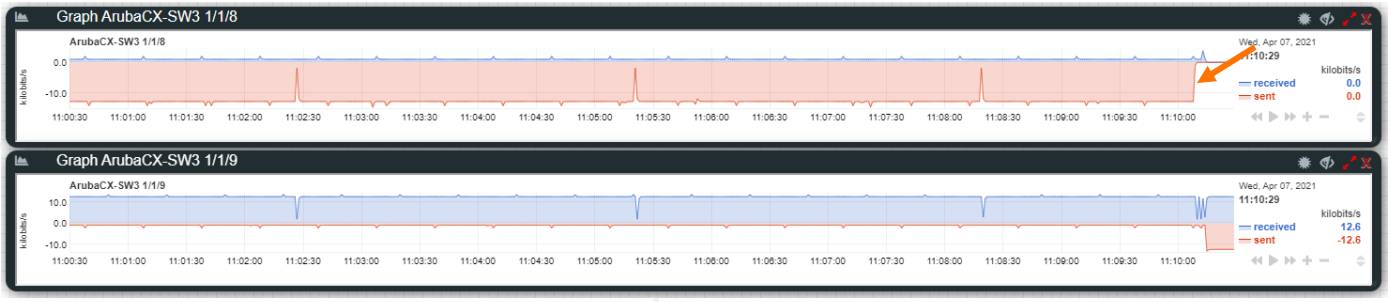

Start again SW1 (by right-click on SW1), and after few seconds start the console.

When SW1 is rebooted, login as admin and check the vsx status:

SW1

```
SW1# 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          02:01:00:00:01:00   02:01:00:00:01:00
Platform            X86-64              X86-64
Software Version    Virtual.10.06.0110  Virtual.10.06.0110
Device Role         primary             secondary


SW1# sh vsx status linkup-delay
Configured linkup delay-timer                              : 180 seconds
Initial sync status                                        : Completed
Delay timer status                                         : Running
Linkup Delay time left                                     : 2 minutes 33 seconds
Interfaces that will be brought up after delay timer expires : lag1-lag2
Interfaces that are excluded from delay timer              :
```

You should now see VSX primary joining the VSX cluster instead of the VSX secondary like in Test#2, as VSX secondary remains UP while VSX primary rebooted.

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
!
no ip icmp redirect
vrf KA
!
vlan 1
vlan 100
    vsx-sync
vlan 101
    vsx-sync
vlan 1105
    vsx-sync
    description TRANSIT VLAN
interface mgmt
    no shutdown
    ip dhcp
interface lag 1 multi-chassis
    no shutdown
    description SW3 VSX LAG
    no routing
    vlan trunk native 1
    vlan trunk allowed 100
    lacp mode active
    lacp rate fast
interface lag 2 multi-chassis
    no shutdown
    description SW4 VSX LAG
    no routing
    vlan trunk native 1
    vlan trunk allowed 100-101
    lacp mode active
    lacp fallback
    lacp rate fast
interface lag 256
    no shutdown
    description ISL
    no routing
    vlan trunk native 1 tag
    vlan trunk allowed all
    lacp mode active
    lacp rate fast
interface 1/1/1
    no shutdown
    mtu 9100
    description to SW3
    lag 1
interface 1/1/2
    no shutdown
    mtu 9100
    description to SW4
    lag 2
interface 1/1/7
    no shutdown
    vrf attach KA
    description keepalive link
    ip address 192.168.0.0/31
interface 1/1/8
    no shutdown
    mtu 9198
    description ISL
    lag 256
interface 1/1/9
    no shutdown
    mtu 9198
    description ISL
    lag 256
interface loopback 0
    ip address 192.168.2.1/32
```

```
    ipv6 address fd00:192:168:2::1/128
    ip ospf 1 area 0.0.0.0
    ipv6 ospfv3 1 area 0.0.0.0
interface vlan 100
    vsx-sync active-gateways
    ip mtu 9100
    ip address 10.10.100.2/24
    active-gateway ip mac 12:01:00:00:01:00
    active-gateway ip 10.10.100.1
    ipv6 address fd00:10:10:100::2/64
    active-gateway ipv6 mac 12:01:00:00:01:00
    active-gateway ipv6 fd00:10:10:100::1
    ip ospf 1 area 0.0.0.0
    ipv6 ospfv3 1 area 0.0.0.0
interface vlan 101
    vsx-sync active-gateways
    ip mtu 9100
    ip address 10.10.101.2/24
    active-gateway ip mac 12:01:00:00:01:00
    active-gateway ip 10.10.101.1
    ipv6 address fd00:10:10:101::2/64
    active-gateway ipv6 mac 12:01:00:00:01:00
    active-gateway ipv6 fd00:10:10:101::1
    ip ospf 1 area 0.0.0.0
    ipv6 ospfv3 1 area 0.0.0.0
interface vlan 1105
    ip address 192.168.4.0/31
    ipv6 address fd00:192:168:4::/127
    ip ospf 1 area 0.0.0.0
    no ip ospf passive
    ip ospf cost 50
    ip ospf network point-to-point
    ip ospf authentication message-digest
    ip ospf message-digest-key 1 md5 ciphertext
AQBapUbZyuMyDkoDN0zeQbI8qY0p5vpa77xnpPQEngEkpWjWBQAAAIouj7OC
    ipv6 ospfv3 1 area 0.0.0.0
    no ipv6 ospfv3 passive
    ipv6 ospfv3 cost 50
    ipv6 ospfv3 network point-to-point
    ipv6 ospfv3 authentication ipsec spi 256 sha1 ciphertext
AQBapUbZyuMyDkoDN0zeQbI8qY0p5vpa77xnpPQEngEkpWjWBQAAAIouj7OC
vsx
    system-mac 02:01:00:00:01:00
    inter-switch-link lag 256
    role primary
    keepalive peer 192.168.0.1 source 192.168.0.0 vrf KA
    vsx-sync aaa acl-log-timer bfd-global bgp control-plane-acls copp-policy dhcp-relay dhcp-server
dhcp-snooping dns icmp-tcp lldp loop-protect-global mac-lockout mclag-interfaces neighbor ospf qos-
global route-map sflow-global snmp ssh stp-global time vsx-global
!
router ospf 1
    router-id 192.168.2.1
    max-metric router-lsa on-startup
    passive-interface default
    area 0.0.0.0
router ospfv3 1
    router-id 192.168.2.1
    max-metric router-lsa on-startup
    passive-interface default
    area 0.0.0.0
```

SW2

```
hostname SW2
!
no ip icmp redirect
vrf KA
!
vlan 1
vlan 100
    vsx-sync
vlan 101
```

```
    vsx-sync
vlan 1105
    vsx-sync
    description TRANSIT VLAN
interface mgmt
    no shutdown
    ip dhcp
interface lag 1 multi-chassis
    no shutdown
    description SW3 VSX LAG
    no routing
    vlan trunk native 1
    vlan trunk allowed 100
    lacp mode active
    lacp rate fast
interface lag 2 multi-chassis
    no shutdown
    description SW4 VSX LAG
    no routing
    vlan trunk native 1
    vlan trunk allowed 100-101
    lacp mode active
    lacp fallback
    lacp rate fast
interface lag 256
    no shutdown
    description ISL
    no routing
    vlan trunk native 1 tag
    vlan trunk allowed all
    lacp mode active
    lacp rate fast
interface 1/1/1
    no shutdown
    mtu 9100
    description to SW3
    lag 1
interface 1/1/2
    no shutdown
    mtu 9100
    description to SW4
    lag 2
interface 1/1/7
    no shutdown
    vrf attach KA
    description keepalive link
    ip address 192.168.0.1/31
interface 1/1/8
    no shutdown
    mtu 9198
    description ISL
    lag 256
interface 1/1/9
    no shutdown
    mtu 9198
    description ISL
    lag 256
interface loopback 0
    ip address 192.168.2.2/32
    ipv6 address fd00:192:168:2::2/128
    ip ospf 1 area 0.0.0.0
    ipv6 ospfv3 1 area 0.0.0.0
interface vlan 100
    vsx-sync active-gateways
    ip mtu 9100
    ip address 10.10.100.3/24
    active-gateway ip mac 12:01:00:00:01:00
    active-gateway ip 10.10.100.1
    ipv6 address fd00:10:10:100::3/64
    active-gateway ipv6 mac 12:01:00:00:01:00
    active-gateway ipv6 fd00:10:10:100::1
    ip ospf 1 area 0.0.0.0
    ipv6 ospfv3 1 area 0.0.0.0
```

```
interface vlan 101
    vsx-sync active-gateways
    ip mtu 9100
    ip address 10.10.101.3/24
    active-gateway ip mac 12:01:00:00:01:00
    active-gateway ip 10.10.101.1
    ipv6 address fd00:10:10:101::3/64
    active-gateway ipv6 mac 12:01:00:00:01:00
    active-gateway ipv6 fd00:10:10:101::1
    ip ospf 1 area 0.0.0.0
    ipv6 ospfv3 1 area 0.0.0.0
interface vlan 1105
    ip address 192.168.4.1/31
    ipv6 address fd00:192:168:4::1/127
    ip ospf 1 area 0.0.0.0
    no ip ospf passive
    ip ospf cost 50
    ip ospf network point-to-point
    ip ospf authentication message-digest
    ip ospf message-digest-key 1 md5 ciphertext
AQBapUbZyuMyDkoDN0zeQbI8qY0p5vpa77xnpPQEngEkpWjWBQAAAIouj7OC
    ipv6 ospfv3 1 area 0.0.0.0
    no ipv6 ospfv3 passive
    ipv6 ospfv3 cost 50
    ipv6 ospfv3 network point-to-point
    ipv6 ospfv3 authentication ipsec spi 256 sha1 ciphertext
AQBapUbZyuMyDkoDN0zeQbI8qY0p5vpa77xnpPQEngEkpWjWBQAAAIouj7OC
vsx
    system-mac 02:01:00:00:01:00
    inter-switch-link lag 256
    role secondary
    keepalive peer 192.168.0.0 source 192.168.0.1 vrf KA
    vsx-sync aaa acl-log-timer bfd-global bgp control-plane-acls copp-policy dhcp-relay dhcp-server
dhcp-snooping dns icmp-tcp lldp loop-protect-global mac-lockout mclag-interfaces neighbor ospf qos-
global route-map sflow-global snmp ssh stp-global time vsx-global
!
router ospf 1
    router-id 192.168.2.2
    max-metric router-lsa on-startup
    passive-interface default
    area 0.0.0.0
router ospfv3 1
    router-id 192.168.2.2
    max-metric router-lsa on-startup
    passive-interface default
    area 0.0.0.0
```

SW3

```
hostname SW3
!
vlan 1,100
interface mgmt
    no shutdown
    ip dhcp
interface lag 1
    no shutdown
    no routing
    vlan trunk native 1
    vlan trunk allowed 100
    lacp mode active
    lacp rate fast
interface 1/1/1
    no shutdown
    no routing
    vlan access 100
interface 1/1/8
    no shutdown
    mtu 9100
    description to SW1
    lag 1
interface 1/1/9
```

```
    no shutdown
    mtu 9100
    description to SW2
    lag 1
```

SW4

```
hostname SW4
!
vlan 1,100-101
interface mgmt
    no shutdown
    ip dhcp
interface lag 1
    no shutdown
    no routing
    vlan trunk native 1
    vlan trunk allowed 100-101
    lacp mode active
    lacp rate fast
interface 1/1/1
    no shutdown
    no routing
    vlan access 101
interface 1/1/8
    no shutdown
    mtu 9100
    description to SW1
    lag 1
interface 1/1/9
    no shutdown
    mtu 9100
    description to SW2
    lag 1
```

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