LAB GUIDE



. . . . . . . . . . . . . . . .

# VSX Lab1 - Layer2

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	
Lab Overview	
Lab Network Layout	
Lab Tasks	
Task 1 – Lab setup	
Task 2 – Configure VSX	
Prerequisite: same firmware release	5
Step #1: create LAG for ISI	5
Step #2://SX keepalive preparation	6
Step #3: VSX Cluster creation	7
Stop #4: VSX koopaliyo	Q
Step #4. VSA Reepailve	0
Step #5. Configuration-sync and vsx-sync reatureGroup settings	
Step #6: VSX split-recovery	
Step #7: VSX linkup-delay-timer	
Step #8: VLANs configuration	
Step #9: Downstream VSX LAG (MCLAG) configuration	
Step #10: Access Switches configuration	
Task 3 - Resiliency tests	15
Test #1: Laver2 connectivity between HostA and HostB	
Test #2: resiliency on shutting down interfaces	17
Test #2: VSX split resilioney on ISL out	19
Appendix Deference Configurations	
Appendix –Reference Configurations	

•••

# Lab Objective

This lab will enable the reader to gain hands-on experience with VSX basic Layer2 configuration.

## Lab Overview

This lab guide explains how to configure a VSX cluster of a pair of AOS-CX switches following the <u>VSX Configuration Best</u> <u>Practices (https://support.hpe.com/hpesc/public/docDisplay?docId=a00094242en\_us</u>), for the Layer2 aspects.

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

In this lab, you'll be able to:

- Configure VSX and VSX LAG (MCLAG) for Layer2 traffic
- Test L2 connectivity between clients: HostA and HostB that are part of the same subnet
- Test resiliency by shutting down interfaces
- Test a VSX split

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:

- <u>on the VSX nodes</u>: remove ports from VSX LAGs, shut all ports, write mem, reboot, no shut all ports and finally reassign 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.



# Lab Tasks

#### 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 (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 (for initial copy/paste):

SW1	SW2	
hostname SW1	hostname SW2	
! •••	• • • • • • • • • • • • • • • • • • •	
vlan 1	vlan 1	
interface mgmt	interface mgmt	
no shutdown	no shutdown	
ip dhcp	ip dhcp	
interface 1/1/1	interface 1/1/1	
no shutdown	no shutdown	
description to SW3	description to SW3	
interface 1/1/2	interface 1/1/2	
no shutdown	• no• shutdown • • • • • • • • • • • • • • • • • • •	
description to SW4	description to SW4	• • .
interface 1/1/7	interface 1/1/7	
no shutdown	no shutdown • • • • • • • • • • • • • • • • • • •	
description keepalive link	description keepalive link	
interface 1/1/8	interface 1/1/8	
no shutdown	no shutdown	
description ISL link	description ISL link	
interface 1/1/9	interface 1/1/9	
no shutdown	no shutdown	
description ISL link	description ISL link	
SW3	SW4	
hostname SW3	hostname SW4	
!		
vlan 1	vlan 1	
interface mgmt	interface mgmt	
no shutdown	no shutdown	
ip dhcp	ip dhcp	
interface 1/1/1	interface 1/1/1	
no shutdown	no shutdown	
interface 1/1/8	interface 1/1/8	
no shutdown	no shutdown	
description to SW1	description to SW1	
interface 1/1/9	interface 1/1/9	
no shutdown	no shutdown	
description to SW2	description to SW2	

• Verify the connectivity through LLDP neighbor information as follows:

SW1		.,	,
SW1#	show	lldp	neighbor-info

```
LLDP Neighbor Information
```

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

LOCAL-PORT	CHASSIS-ID	PORT-ID	PORT-DESC	TTL	SYS-NAME
1/1/1 1/1/2 1/1/7 1/1/8 1/1/9	08:00:09:5b:7e:2d 08:00:09:ed:b5:6e 08:00:09:54:97:83 08:00:09:54:97:83 08:00:09:54:97:83	1/1/8 1/1/8 1/1/7 1/1/8 1/1/9	to SW1 to SW1 keepalive link ISL ISL	120 120 120 120 120 120	SW3 SW4 SW2 SW2 SW2 SW2
SW2					
SW2# show l	ldp neighbor-info				
LLDP Neighb	or 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 1/1/2 1/1/7 1/1/8 1/1/9	08:00:09:5b:7e:2d 08:00:09:ed:b5:6e 08:00:09:d7:5f:0f 08:00:09:d7:5f:0f 08:00:09:d7:5f:0f	1/1/9 1/1/9 1/1/7 1/1/8 1/1/9	to SW2 to SW2 keepalive link ISL ISL	120 120 120 120 120 120	SW3 SW4 SW1 SW1 SW1 SW1

## Task 2 – Configure VSX

Prerequisite: same firmware release

Both CX switches SW1 and SW2 must run the same version (the version exposed here is an example):

SW1	SW2	, 	•••	•••
SW1# show version	SW2# show version	) • • ) • •		•••
ArubaOS-CX (c) Copyright Hewlett Packard Enterprise Development LP	ArubaOS-CX (c) Copyright Hewlett Packard Enterprise Development LP			
Version : Virtual.10.06.0110 Build Date : Build ID : ArubaOS- CX:Virtual.10.06.0110:bc56d8a669a9:202103161859 Build SHA : bc56d8a669a92c8fe9d946e01c7a791c538b3bdd Active Image :	Version : Virtual.10.06.0110 Build Date : Build ID : ArubaOS- CX:Virtual.10.06.0110:bc56d8a669a9:202103161859 Build SHA : bc56d8a669a92c8fe9d946e01c7a791c538b3bdd Active Image :	I       0       0         I       0       0         I       0       0         I       0       0         I       0       0         I       0       0         I       0       0         I       0       0         I       0       0         I       0       0         I       0       0         I       0       0         I       0       0		
Service OS Version : BIOS Version :	Service OS Version : BIOS Version :			

#### Step #1: create LAG for ISL

It is highly recommended to have two physical interconnectivity links for ISL redundancy. See the VSX Best Practices for ISL bandwidth recommendation.

SW1(config)#	SW2(config)#
interface lag 256	interface lag 256
no shutdown	no shutdown
description ISL	description ISL
no routing	no routing
vlan trunk allowed all	vlan trunk allowed all
lacp mode active	lacp mode active
interface 1/1/8	interface 1/1/8
no shutdown	no shutdown
mtu 9198	mtu 9198
description ISL link	description ISL link
lag 256	lag 256
interface 1/1/9	interface 1/1/9
no shutdown	no shutdown
mtu 9198	mtu 9198
description ISL link	description ISL link
lag 256	lag 256

<u>Note</u>: Since 10.4, regardless of the MTU value configured on ports used for ISL, the MTU will be fixed internally to 9198 bytes for the ports used for ISL. This is however not applicable to the AOS-CX Switch Simulator yet.

#### Check that the ISL LAG is UP.

SW1	SW2
SW1# show interface lag256	SW2# show interface lag 256
Aggregate lag256 is up	Aggregate lag256 is up

	• • • •	• • • • • • • • • • • • • • • • • • • •	Layer2
	• • • •		
The second second			
Admin state is up		Admin state is up	
Description : ISL		Description : ISL	
MAC Address	: 08:00:09:d7:5f:0f	MAC Address	54:97:83
Aggregated-interfaces	: 1/1/8 1/1/9	• • Aggregated-interfaces • • • • • : 1/1/8 1/1	/9
Aggregation-key	: 256	Aggregation-key : 256	
Aggregate mode	: active	Aggregate mode : active	
Speed	: 2000 Mb/s	Speed • • • • • • • • • • • • • • • • • •	
L3 Counters: Rx Disabled, T	x Disabled 📃	L3 Counters: Rx Disabled, Tx Disabled	
gos trust none	•	gos trust none	
VLAN Mode: native-untagged		VLAN Mode: native-untagged	
Native VLAN: 1		Native VLAN: 1	
Allowed VLAN List all		Allowed VLAN List all	•
Rx		Ry	
7// total packets	97355 total bytes	815 total packets 106239	total bytes
0 uniquet packets	STSSS COLAI DYLES	0 unicast packets	cocar byces
0 unicast packet	5	0 unicast packets	
U multicast pack	ets	• 0 • multicast packets • • • • • •	
U broadcast pack	ets	0 broadcast packets	
0 errors	0 dropped	0 errors 0 ·	dropped
0 CRC/FCS	0 pause	0 CRC/FCS	pause
Τx		Tx	<mark></mark> .
778 total packets	1450 total bytes	779 total packets 101814	total bytes
0 unicast packets		0 unicast packets	
0 multicast packe	ts	0 multicast packets	
0 broadcast packe	ts	0 broadcast packets	
0 errors	0 dropped	0 errors	dropped
0 collision		0 collision	
0 0011131011	0 pause	0 0011131011	pause

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

#### SW1 / SW2

SW1# show lacp interfaces								
State abbre A - Active S - Short- C - Collec X - State n	eviations : P - timeout L - ting D - m/c expired	Passiv Long-1 Distr:	ve timeout ibuting	F - Ago N - Ins J E - Def	gregable I - Indiv Sync O - Outofs Fault neighbor stat	idual Sync		
Actor deta:	ils of all :	interfa	aces:					
Intf	Aggr Name	Port Id	Port Pri	State	System-ID	System Pri	Aggr Key	Forwarding State
1/1/8 1/1/9	lag256 lag256	9 10	1 1	ALFNCD ALFNCD	08:00:09:d7:5f:0f 08:00:09:d7:5f:0f	65534 65534	256 256	up up
Partner det	tails of all	l inte	rfaces	:				
Intf	Aggr Name	Port Id	Port Pri	State	System-ID	System Pri	Aggr Key	
1/1/8 1/1/9	lag256 lag256	9 10	1 1	ALFNCD ALFNCD	08:00:09:54:97:83 08:00:09:54:97:83	65534 65534	256 256	

Note: If ISL LAG is not UP while LLDP information are properly received, it might be useful to bounce interfaces 1/1/8 and 1/1/9 on both SW1 and SW2 (shut/no shut interfaces). This issue may happen sometime on AOS-CX Switch Simulator.

Note: at this stage VLAN mode is native-untagged (it will change when ISL function is associated to this LAG).

#### Step #2: VSX keepalive preparation

Create the dedicated KeepAlive VRF and associated interface Although KA VRF is optional, it is a good practice.

SW1(config)#	SW2(config)#
vrf KA	vrf KA

VSX Lab1

interface 1/1/7
 no shutdown
 vrf attach KA
 description VSX keepalive
 ip address 192.168.0.0/31

no shutdown vrf attach KA description VSX keepalive ip address 192.168.0.1/31

interface 1/1/7

Check IP connectivity between future VSX nodes inside this dedicated "KA" VRF.

#### SW1 / SW2

SW1# ping 192.168.0.1 vrf KA
PING 192.168.0.1 (192.168.0.1) 100(128) bytes of data.
108 bytes from 192.168.0.1: icmp\_seq=1 ttl=64 time=18.9 ms
108 bytes from 192.168.0.1: icmp\_seq=2 ttl=64 time=2.52 ms
108 bytes from 192.168.0.1: icmp\_seq=3 ttl=64 time=2.72 ms
108 bytes from 192.168.0.1: icmp\_seq=4 ttl=64 time=2.40 ms
108 bytes from 192.168.0.1: icmp\_seq=5 ttl=64 time=6.51 ms
--- 192.168.0.1 ping statistics --5 packets transmitted, 5 received, 0% packet loss, time 4003ms
rtt min/avg/max/mdev = 2.406/6.633/18.999/6.371 ms

#### Step #3: VSX Cluster creation

The **best practice for system-mac** is to set the system-mac manually on the VSX primary switch. Please refer to the VSX Best Practices for system-mac values. Here, 02:01:00:00:01:00 is used. The main advantage to set VSX system-mac (and not to leave it blank with default HW system-mac being used) is to be independent from the physical hardware MAC address. In case of hardware replacement of the VSX primary, the new switch can be configured exactly with the same configuration than the previous unit and there will be no impact on the secondary which will remain in the same cluster ID. HW replacement is hitless for the VSX secondary.

The **best practice for inter-switch-link timers** (dead-interval, hello-interval, hold-time, peer-detect-interval) is to keep the default timers (i.e. no specific configuration).

The **best practice for role** (primary or secondary) is to have a meaningful relationship with the switch hostname/identification. Example: SW1 is the VSX primary and SW2 is the VSX secondary.

SW1(config)#	SW2(config)#
vsx system-mac 02:01:00:00:01:00 inter-switch-link lag 256 role primary vsx-sync vsx-global	vsx inter-switch-link lag 256 role secondary

At this stage few aspects can be highlighted:

- Best practice for vsx-sync includes vsx-global. Thanks to this vsx-sync FeatureGroup parameter, the VSX managementplane will synchronize the following VSX settings: inter-switch-link hello-interval, dead-interval, hold-time, peer-detectinterval, keepalive udp-port, hello-interval, keepalive dead-interval, system-mac, split-recovery, linkup-delay-timer.
- VSX automatically tags the native VLAN configured on the LAG used for ISL.

SW1			SW2		
SW1# show vsx statu: VSX Operational Sta	s te		SW2# show vsx statu VSX Operational Sta	is ite	
ISL channel ISL mgmt channel Config Sync Statu: NAE HTTPS Server	: In-Sync : operational s : In-Sync : peer_reacha : peer_reacha	: In-Sync : operational : In-Sync : peer_reachable : peer_reachable		: In-Sync : operatio s : In-Sync : peer_rea : peer_rea	nal chable chable
Attribute	Local	Peer	Attribute	Local	Peer
ISL link ISL version	lag256 2	lag256 2	ISL link ISL version	lag256 2	lag256 2

VSX Lab1	
Laver2	
System MAC 02:01:00:00:01:00 System MAC 02:01:00:00:01:00	
02:01:00:00:01:00	
Platform X86-64 X86-64 Platform X86-64 X86-64	
Software Version Virtual.10.06.0110 ••• Software Version • Virtual.10.06.0110	
Virtual.10.06.0110	
Device Role primary secondary Device Role secondary primary	
SW1# show running-config   begin 5 vsx SW2# show running-config   begin 5 vsx	
vsx vsx	
system-mac 02:01:00:00:01:00 system-mac 02:01:00:00:01:00 synchronized	
inter-switch-link lag 256 inter-switch-link lag 256	
role primary role secondary	
vsx-sync vsx-global vsx-sync vsx-global synchronized	
SW1# show running-config vsx SW2# sh running-config vsx	
VSX VSX	
system-mac 02:01:00:00:01:00 system-mac 02:01:00:00:01:00	
inter-switch-link lag 256 inter-switch-link lag 256	
role primary role secondary	
vsx-sync vsx-global vsx-sync vsx-global	
interface lag 256	
description ISL description ISL	
no shutdown	
no routing no routing	
vlan trunk native 1 tag vlan trunk native 1 tag	
vlan trunk allowed all vlan trunk allowed all	
lacp mode active	
interface 1/1/9 interface 1/1/8	• • • • • • • •
no shutdown no shutdown	
lag 256 lag 256	
interface 1/1/8 interface 1/1/9	
no shutdown no shutdown	
lag 256 lag 256	

At this stage, the <u>VSX cluster is created</u> and show command vsx-peer can be used from any VSX node. However, the cluster is not protected yet against a split. (keepalive is not yes established). Step#3 and Step#4 can be merged into a single step (here they are separated for educational purpose).

#### SW1 / SW2

SW1# show vsx brief	
ISL State	: In-Sync
Device State	: Peer-Established
Keepalive State	: Keepalive-Init
Device Role	: primary
Number of Multi-chassis LAG interfaces	: 0
SW1# show vsx brief <mark>vsx-peer</mark>	
ISL State	: In-Sync
Device State	: Peer-Established
Keepalive State	: Keepalive-Init
Device Role	: secondary
Number of Multi-chassis LAG interfaces	: 0

#### Step #4: VSX keepalive

The **best practice for VSX keepalive timers** (dead-interval, hello-interval) is to keep the default timers (i.e. no specific configuration).

SW1(config)#		SW2(config)#
VSX keepalive peer 192.168.0.1 source 192.168.0.0 vr	f KA	VSX keepalive peer 192.168.0.0 source 192.168.0.1 vrf KA
SW1 / SW2		
SW1# show vsx brief ISL State Device State Keepalive State Device Role	: In-Sync : Peer-Estal : <mark>Keepalive</mark> : primary	olished <mark>-Established</mark>

			• • • • •	• • • •		• • •	• • •	•								
			• • • • •	• • • •		• • •	• • •	•								
												VS	X Lab	1		
														0		
													Layer	2		
Number of Multi-chassis LA	G interfaces .	0														
Number of Hufer enabers in																
SWl# show vsx brief <mark>vsx-pe</mark>	eer															
ISL State	:	In-Sync••				• • •	• • •									
Device State	:	Peer-Estab	lished	• • • •		• • •	• • •	• • • •	•							
Keepalive State	•	Keepalive-	Establi	shed		• • •	• • •	• • • •	•							
Dowigo Polo		Rocondaru	2000022	onou												
Device Kore	•	secondary														
Number of Multi-chassis LA	AG interfaces :	0														
SW2# show vsx status keepa	alive															
Keepalive State	: Keepalive-Es	tablished		• • • •		• • •	• • •	• • • •			••					
Last Established	• Thu Apr 11	8:55:53 202	1 • • • •	• • • •		• • •	• • •	• • • •	• • • •	• • •	• • •					
Last Failed	• Thu Apr 1 1	8.55.27 202	1	• • • •		• • •	• • •	• • • •		• • •	•••					
	. Inu Apr I I	0.00.202														
Peer System 1d	: 02:01:00:00:	01:00														
Peer Device Role	: secondary															
Keepalive Counters																
Keepalive Packets Tx	· 230					• • •	• • •				• • •				• • •	
Keepalive Tackets IX	. 201			•		• • •	• • •	• • • •		• • •	• • •		• • • •		• • •	• • •
Reeparive Fackets RX	: 201					• • •	• • •	• • • •		• • •	• • •	• • •	• • • •		• • •	• • •
Keepa⊥ive Timeouts	: 0				•••	•••	•••		••••		•••	•••	••••			•••
Keepalive Packets Dropped	: 0															

## Step #5: Configuration-sync and vsx-sync FeatureGroup settings

The **best practice for VSX configuration-sync** is to keep the default enabled configuration-synchronization (no configuration change).

SW1			SW2			
SW1# show vsx status con	nf	ig-sync	SW2# show vsx status con	nfi	ig-sync 🔹 • •	•
Admin state	:	Enabled	Admin state	:	Enabled	
Operational State	:	Operational	Operational State	:	Operational	
Error State	:	None	Error State	:	None	
Recommended remediation	:	N/A	Recommended remediation	:	N/A	
Current time	:	Thu Nov 28 15:54:17 2019	Current time	:	Thu Nov 28 16:05:31 2019	
Last sync time	:	Thu Nov 28 15:40:27 2019	Last sync time	:	Thu Nov 28 15:40:27 2019	

The **best practice for vsx-sync global settings** is to use as much as possible the automatic synchronization, to avoid human errors. From the list below, the best practice settings are highlighted (in the context of a traditional Aggregation layer for IPv4). According to specific requirements, more or less parameters can be used from the list.

SW1(config)#	
SW1(config-vsx)# vsx-sync ?	
aaa Sync all AAA instances	
acl-log-timer Sync access-list log timer instance	
arp-security Sync all ARP security configurations	
bfd-global Sync all BFD global configuration	
bgp Sync all BGP, ip aspath list, community list, prefix	
list, route map configurations	
control-plane-acls Sync all Control-plane Access-list instances	
copp-policy Sync all CoPP instances	
dhcp-relay Sync all DHCP RELAY instances	
dhcp-server Sync all DHCPv4-Server and DHCPv6-Server instances	
dhcp-snooping Sync all DHCPv4-Snooping and DHCPv6-Snooping instances.	
dns Sync all DNS instances	
evpn Sync all evpn configurations	
icmp-tcp Sync all icmp and tcp instances	
keychain Sync all keychain configurations	
lldp Sync all LLDP instances	
loop-protect-global Sync all Loop-protect global configuration	
mac-lockout Sync all mac lockout configurations	
macsec Sync all MACsec and MKA policies	
mclag-interfaces Sync QoS, LACP, Loop-Protect, LAG description, sFlow,	
STP, Rate-Limits, Vlans, ACLs, MACsec and Portfilters	
for MCLAG interface instances	
mgmd-global Sync all MGMD global instances	
nd-snooping Sync all ND-Snooping instances.	
neighbor Sync all IPv4 and IPv6 static neighbor entries	
ospf Sync all OSPF instances	
qos-global 🔪 Sync all QoS global instances	
rip Sync all RIP configurations	

route-map	Sync all ip aspath list, community list, prefix list,
	route map configurations
sflow-global	Sync all sFlow global instances
snmp	Sync all SNMP instances
ssh	Sync all SSH instances
static-routes	Sync all Static Routes instances
stp-global	Sync all STP Global Configuration
time	Sync all time instances
udp-forwarder	Sync all UDP FORWARDER instances
vrrp	Sync all VRRP instances
vsx-global	Sync all VSX global configuration

In addition of the current vsx-global feature synchronization, the other settings are added to the VSX primary. No configuration for this synchronization features is required on VSX secondary switch which will inherit from VSX primary.

Configure the following vsx-sync parameters on SW1:

SW1(config)#	SW2(config)#	•••	•••	•••
vsx vsx-sync aaa acl-log-timer bfd-global bgp control-plane-acls copp-policy dhcp-relay <del>dhcp-</del> server dhcp-snooping dns icmp-tcp lldp <del>loop-</del> <del>protect-global</del> mac-lockout mclag-interfaces neighbor ospf qos-global route-map sflow-global snmp ssh stp-global time vsx-global	synchronized		<ul> <li>•</li> <li>•&lt;</li></ul>	<ul> <li>•</li> <li>•&lt;</li></ul>
SW1 / SW2		•••	•••	•••
<pre>SW1# show running-config vsx-sync <snipped> vsx system-mac 02:01:00:00:01:00 vsx-sync aaa acl-log-timer bfd-global bgp contr cp-snooping dns icmp-tcp lldp loop-protect-global m l route-map sflow-global snmp ssh stp-global time v</snipped></pre>	ol-plane-acls copp-policy dhcp-relay dhcp-server dh ac-lockout mclag-interfaces neighbor ospf qos-globa sx-global		<ul> <li>•</li> <li>•&lt;</li></ul>	
<pre>SW2# show running-config   begin 0 vsx vsx vsx-sync aaa acl-log-timer bfd-global bgp contr cp-snooping dns icmp-tcp lldp loop-protect-global m l route-map sflow-global snmp ssh stp-global time v</pre>	ol-plane-acls copp-policy dhcp-relay dhcp-server dh ac-lockout mclag-interfaces neighbor ospf qos-globa sx-global			

#### Step #6: VSX split-recovery

The **best practice for VSX split-recovery** is to keep the default split-recovery enabled (no configuration change). This best practice might be revisited in case of VSX and VXLAN VTEP.

SW1	SW2
SW1# show vsx configuration split-recovery Split Recovery Mode : Enabled	<pre>SW2# show vsx configuration split-recovery Split Recovery Mode : Enabled</pre>

#### Step #7: VSX linkup-delay-timer

The best practice for VSX linkup-delay-timer for mid-size network (<10k MAC/ARP) is to keep the default timer.

Linkup-delay timer is the period of time to wait before the VSX LAG ports are enabled during a VSX cluster join event.

SW1 / SW2 (mid-size network)	
SW1# show vsx status linkup-delay	
Configured linkup delay-timer	: 180 seconds
Initial sync status	: Completed
Delay timer status	: Completed
Linkup Delay time left	:
Interfaces that will be brought up after delay time	er expires :
Interfaces that are excluded from delay timer	:
Sw2# show vsx status linkup-delay	
Configured linkup delay-timer	: 180 seconds
Initial sync status	: Completed
Delay timer status	: Completed

Linkup Delay time left Interfaces that will be brought up after delay timer expires Interfaces that are excluded from delay timer

#### Step #8: VLANs configuration

The **best practice for VLANs configuration** is to <u>configure the VLANs on the VSX primary with the **vsx-sync** attribute</u> and let the VSX config-sync automatically synchronize the VLANs on the VSX secondary. Here, VLAN 100 is the endpoints VLAN.

SW1(config)#		AGG-2(config)#
SW1(config)# vlan 100		
SW1(config-vlan-<100>)# vsx-sync		synchronized
SW1# show running-config vsx-sync   beg 1	vlan	SW2# show run vsx-sync   beg 1 vlan
vlan 100		vlan 100
vsx-sync		vsx-sync
SW1 /SW2		
SW2# show vlan		
VLAN Name St	tatus Re	eason Type Interfaces
1 DEFAULT VLAN 1 ur	o ok	default lag256
100 VLAN100 ur	o ok	static lag256
1	-	

<u>Note</u>: if vsx-sync attribute is removed from the configuration element on the VSX primary the configuration item will stay on the VSX secondary without the vsx-sync keyword. Consequently, if an item is then removed from the VSX primary, it will stay on the VSX secondary.

#### Step #9: Downstream VSX LAG (MCLAG) configuration

The **best practice for VSX LAG** is to create the multi-chassis lag interface on the VSX primary with all settings and then create the mirrored lag interface on the VSX secondary. LAG interface settings (including description) will be synchronized automatically. Only "no shut" in the lag interface context has to be performed on the VSX secondary. Once the multi-chassis lag interface is created, it can be assigned to the physical port.

The **best practice for allowed VLANs** is to exclude the native VLAN 1 from being propagated. This is a very robust method to avoid Layer2 storm propagation due to potential loop initiated on an access switch. In case of access switch Zero-Touch-Provisioning use-case., this trunking exclusion is performed after the ZTP process.

The **best practice for LACP timers** on the VSX LAG is to keep the default long timer (30s = lacp rate slow). The LACP fast rate does not provide any acceleration of link failure detection as link-state reported by transceivers takes care of this. Consequently, as LACP fast rate does not provide much benefits in most of the scenarios, and as it may increase CPU load, it is recommended to keep the slow rate.

The **best practice for MTU** is to configure on all devices the appropriate size to support features such as Dynamic Segmentation or server jumbo frame. Care should be taken to ensure that the IP path from the access devices (switches or APs) can provide a MTU of at least 1564 bytes to the mobility controllers and that the server jumbo packet of 9000 bytes can be encapsulated. Flexibility should be anticipated to perform VXLAN encapsulation from the access switch (9000+50) or VXLAN encapsulation from the aggregation layer MTU+50. So the recommended Ethernet MTU is 9100 bytes for the downstream VSX LAG to the access layer and a MTU of 9000 bytes for endpoints or servers. The SVI IP MTU should match the MTU size on the aggregation layer, so the recommended IP MTU is 9100 bytes.

The **best practice for hashing algorithm** on the VSX LAG is to keep the default I3-src-dst (alternative being I2-src-dst). This option has an effect only if at least 2 ports per VSX node are members of the same VSX LAG.

<u>Note</u>: Most of the time the VSX LAG includes only two links: one link from the primary and one link from the secondary. Consequently, hashing algorithm selection has no effect on the traffic path as it is forwarded to the local port of the VSX LAG on the switch receiving the traffic.

SW1(config)#	SW2(config)#	
interface lag 1 multi-chassis	interface lag 1 multi-chassis	
description SW3 VSX LAG	no shutdown	
no shutdown		
vlan trunk allowed 100		
interface 1/1/1	interface 1/1/1	
no shutdown	no shutdown	
mtu 9100	mtu 9100	
description to SW3	description to SW3	
lag 1		•
		•
interface lag 2 multi-chassis	interface lag 2 multi-chassis	
description SW4 VSX LAG	no shutdown	
no shutdown	no bildedowi	. •
vian trunk allowed 100		
Vian clunk allowed 100		
interface 1/1/2	interface 1/1/2	
interiace 1/1/2		
no shuldown	no shuldown	
mtu 9100	mtu 9100	
description to SW4	description to SW4	
Lag 2	Lag 2	

#### Step #10: Access Switches configuration

SW3(config)#	SW4(config)#
vlan 100	vlan 100
interface lag 1	interface lag 1
no shutdown	no shutdown
no routing	no routing
vlan trunk native 1	vlan trunk native 1
vlan trunk allowed 100	vlan trunk allowed 100
lacp mode active	lacp mode active
interface 1/1/8	interface 1/1/8
no shutdown	no shutdown
description to SW1	description to SW1
lag 1	lag 1
interface 1/1/9	interface 1/1/9
no shutdown	no shutdown
description to SW2	description to SW2
lag 1	lag 1
interface 1/1/1	interface 1/1/1
no shutdown	no shutdown
no routing	no routing
vlan access 100	vlan access 100

#### Check the LAGs state:

#### SW3 / SW4

SW3# 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:

Inti	Aggr Name	Port Id	Port Pri	State	System	LD	System Pri	Aggr Key	Forwarding State
1/1/8 1/1/9	lag1 lag1								up up

VSX Lab1 Layer2

Partner de	etails of al	l inter	faces:			••••	• • • • • • • • • •	• • • • • • • •	• • • •									
Intf	Aggr Name	Port Id	Port S Pri	State	e Syst	em-ID	Syst Pri	em Aggr Key										
1/1/8	lag1						• • • • • • • •		• • • • • •	•								
1/1/9	lagi					•••	• • • • • • • • •			••								
_, _, _						• •				• • • .								
SW3# show	interface l	ag 1				•••	• • • • • • • • •			• • • •								
		2									• • •							
Aggregate	lag1 is up						• • • • • • • •	• • • • • • •	• • • • • •	• • • •	• • •	•						
Admin sta	ate is up						· • • • • • • • •											
Descript	ion :								• • • • • •	• • • •	• • •	• • •	• • •					
MAC Addre	ess		: 08:0	00:09	9:5b:7e:	2d				• • • •	•••	•••	•••					
Aggregate	ed-interface	S	: 1/1,	/8 1/	/1/9										• •			
Aggregat:	lon-key		: 1	0.10	,			• • • • • • •	• • • • • •	• • • •	• • •	• • •	• • •	•••	••	• • •	• • •	
Speed	D D	la la al m	: 2000	UMD/	S													
L3 Counte	ers: KX Disa	plea, T	x Disa	piea			• • •		• • • • • •		• • •	• • •	• • •		• •	• • •		
YIN Mod	. notivo-un	taggod					•		• • • • • •	• • • •	•••	•••	•••			•••		
Native VI	z. nacive-un LAN· 1	cayyeu																
Allowed V	VIAN List. 1	0.0							• • • • • •	• • • •	• • •	• • •	•••	• • •	••	• • •	•••	
Rx	·							••••				•••						
	1458 total n	ackets			1893	45 tota	al bytes	• •				• • •	• • •			• •		
-	0 unica	st pack	ets				- 2		• • • • • •		• • •	• • •	•••	•••		•••	•••	1
	0 multi	cast pa	ckets															i
	0 broad	cast pa	ckets						• •		• • •	• • •	• • •	• • •	• •	• • •	•••	•
	0 errors					0 drop	pped			· · · ·	•••	•••	•••					
	0 CRC/FCS					0 paus	se								• •			•
Tx												• • •	•••	•••		•••	•••	1
	894 total p 0 unica 0 multi	ackets st pack cast pa	ets ckets		1195	04 tota	al bytes							• • •	• •	• • •		•
	0 broad	cast pa	ckets			0 -1												
	0 collisi	0.0				0 010	opea											
CN14 / CN12	0 COIIISI	011				0 pau	se											
SWI/SWZ	leen intenf		المان مام.		_													
SWI# SNOW	lacp interi	aces mu	lti-cha	assis	5													
State abbr A - Active S - Short C - Collec X - State	reviations : P - -timeout L - cting D - m/c expired	Passiv Long-t Distri	e imeout buting	F - N - E -	Aggrega InSync Default	ble I · O · neighl	- Individual - OutofSync											
	·, · ···			_														
Actor deta	ails of all 	interfa	ces:															
Intf	Aggregate	Port	Port		State	Syster	n-ID	System	Aggr									
	name	id	Prio	rity				Priority	Кеу									
1/1/1	lag1(mc)	1	1		ALFNCD	02:01	:00:00:01:00	65534	1									
1/1/2	lag2(mc)	2	1		ALFNCD	02:01	:00:00:01:00	65534	2									
Do ut no no alu		1 20200																
Partner de	etalls of al	1 inter	Iaces:															
Tntf	Addregate	Partne	r Port		State	Sustar	 n-ТП	Sustem	Agar									
INCI	name	Port-i	d Prior	ritv	State	bybter		Priority	Kev									
1/1/1	lag1(mc)	9	1		ALFNCD	08:00	:09:5b:7e:2d	65534	1									
1/1/2	lag2(mc)	9	1		ALFNCD	08:00	:09:ed:b5:6e	65534	1									
_, _, _	j ( ,	-	-						-									
<mark>Remote</mark> Act	tor details	of all	interfa	aces:	:													
Intf	Aggregate	Port	Port		State	Syster	n-ID	System	Aggr									
	name	id	Prio	rity				Priority	Кеу									
1 /1 /1		1000																
1/1/1	Lagl(mc)	1001	1		ALFNCD	02:01	:00:00:01:00	65534	1									
1/1/2	Lag2(mc)	T005	Ţ		ALF'NCD	02:01	:00:00:01:00	65534	2									
Remote Par	rtner detail	sofal	] inte	rface	-5:													

															_										
Intf	Aggregate name	Partner Port-id	Port Priority	State	Sys <sup>.</sup>	tem-I	D			Syst Pric	cem prit	A Y K	lggr Cey	• •	_										
1/1/1	lag1(mc)	10	1	ALFNCD	08:	00:09	:5b	:7e:	2d	6553	34	1	•••	•••											
1/1/2	lag2(mc)	10	1	ALFNCD	08:	00:09	:ed	:b5:	6e	6553	34	• 1	• •	• • •	• •										
									::				•••	•••		•									
GW2# show	intor log1								• •			• •	• •	• • •	• •	• •									
SWZ# SHOW	incer lagi					• • •	• • •	•••	••	•••	•••	••	•••	•••	•••	•••	•								
Addredate	ladis un								•••			•••	•••	• • •		•••	•••	•••							
Admin st	ate is up					• •	• • •	• • •	• •	• • •	• •	• •	• •	• • •	• •	• •	• •	• •	•						
Descript	ion : SW3 VS	X LAG					•••			•••			•••	•••		•••	•••	•••		•••					
MAC Addr	ess		: 08:00:0	9:54:97:8	33				• •			• •	• •	• • •	• •	• •	• •	• •	• • •	• •	• •	•			
Aggregat	ed-interface	s	: 1/1/1					•••	•••	•••	•••	•••	•••	•••	•••	•••	•••	•••	•••		•••	•••			
Aggregat	ion-key		: 1						•••			•••	•••	• • •		•••	•••	•••			•••	•••			
Aggregat	e mode		: active					•••	• •	• • •	• • •	• •	• •	• • •	• •	• •	• •	• •	• • •		• •	• •	• •	٠	
Speed			: 1000 Mb	/s					•••			•••	•••	•••		•••	•••	•••	•••		•••				
L3 Count	ers: Rx Disa	bled, Tx	Disabled						• •			• •	• •		• •	• •	• •	• •	(		• •	• •	• •		
qos trus	t none								•	• • •	• • •	••	• •	• • •	• •	• •	••	• •	• • •	•	• •	• •	• •	•	
VLAN Mod	e: native-un	tagged											•••	•••		•••	•••	•••	•••						
Native V	LAN: 1											• •	• •	• • •	• •	• •	• •	• •	• • •	• •	• •	• •	• •	•	
Allowed	VLAN List: 1	00									•••	•••	•••	•••	•••	•••	•••	•••	•••		•••	•••			
Rx													•••	• • •		•••	•••	•••	/		•••	•••			
	4/9 total p	ackets		6404	13 t(	stal .	byte	es					• •	• • •	• •	• •	• •	• •	• • •		• •	• •	• •	•	
	0 unica	st packe	LS boto											•••			•••	•••							
		cast pac	kets												• •	• •	• •	• •	• • •	• •	• •	• •	• •	•	
	0 prors	icasi pac	Kels		0 4	ronne	d									••	•••	•••	•••		•••	•••		•	
	0 CRC/FCS				0 n	ause	a														•••	• •			
Τx	0 0100/100				5 P	~~~~													. • (		• •	• •	• •	•	
	765 total r	ackets		9950	)4 t	otal	bvte	es													•	••			
	0 unica	st packe	ts				1																		
	0 multi	cast pac	kets																						
	0 broad	cast pac	kets																						
	0 errors	-			0 d	roppe	d																		
	0 collisi	on			0 p	ause																			

The "show lacp interfaces multi-chassis" command is very useful to get a complete status of the local LACP partnership as well as the VSX peer partnership details. Actor = local node, Partner = LACP neighbor (the access switch), Remote Actor = the VSX peer, Remote Partner = LACP neighbor of the VSX peer.

Note that the port id of the VSX secondary is equal to 1000+ID\_of\_the\_primary (in the example 1001). ALFNCD LACP state-flags should appear on all entries.

The **best practice for LACP fallback** feature is to enable it on the VSX LAGs for the <u>following use-cases</u>: PXE boot, access switch ZTP, server NIC driver migration from active/standby to LACP. When applied to the VSX primary, LACP fallback is automatically synced on the VSX secondary.

Configure LACP fallback on LAG2 only (just for educational purpose in this lab):

SW1(config)#	SW2(config)#
interface lag 2 multi-chassis lacp fallback	synchronized
SW1# show running-config interface lag2 interface lag 2 multi-chassis no shutdown description SW4 VSX LAG	<pre>SW2# show running-config interface lag2 interface lag 2 multi-chassis no shutdown description SW4 VSX LAG</pre>
no routing vlan trunk native 1	no routing vlan trunk native 1
vlan trunk allowed 100	vlan trunk allowed 100
lacp mode active	lacp mode active
lacp fallback	lacp fallback
exit	exit

Further on in this document, LACP fallback is no longer shown as this is reserved for the previous indicated use-cases.

The main configuration for Layer2 on VSX cluster is completed. More advanced configuration might be exposed in other labs, showing VSX interaction with other features like Spanning-tree for instance.

#### 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 SW1/SW2/SW3/SW4 switch configuration, adjust the LACP timer to short on all LAG interfaces:

SW1	SW2	
SW2(config)# interface lag 256	SW2(config)# interface lag 256	
SW2(config-lag-if)# <pre>lacp rate fast</pre>	SW2(config-lag-if)# lacp rate fast	• • •
SW1(config)# interface lag 1-2	synchronized	•••
SW1(config-lag-if-<1-2>)# fact fact fact SW1(config-lag-if-<1-2>)# exit	0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0	•••
Similarly on SW3 and SW4		•••
C/M/3	C/W/A	•••
	5004	
W3(config)# int lag 1	SW4(config)# int lag 1	• • •
SW3(config-lag-if)# lacp rate fast	SW4(config-lag-if)# lacp rate fast	•••

As a reference, configuration of SW1/SW2/SW3/SW4 should look like: (you may copy/paste easily from Appendix if needed)

SW1	SW2
hostname SW1	hostname SW2
!	!
vrf KA	vrf KA
!	!
vlan 1	vlan 1
vlan 100	vlan 100
vsx-sync	vsx-sync
interface mgmt	interface mgmt
no shutdown	no shutdown
ip dhcp	ip dhcp
interface lag 1 multi-chassis	interface lag 1 multi-chassis
no shutdown	no shutdown
description SW3 VSX LAG	description SW3 VSX LAG
no routing	no routing
vlan trunk native 1	vlan trunk native 1
vlan trunk allowed 100	vlan trunk allowed 100
lacp mode active	lacp mode active
lacp rate fast	lacp rate fast
interface lag 2 multi-chassis	interface lag 2 multi-chassis
no shutdown	no shutdown
description SW4 VSX LAG	description SW4 VSX LAG
no routing	no routing
vlan trunk native 1	vlan trunk native 1
vlan trunk allowed 100	vlan trunk allowed 100
lacp mode active 🔍	lacp mode active
lacp fallback	lacp fallback
lacp rate fast	lacp rate fast
interface lag 256	interface lag 256
no shutdown	no shutdown
description ISL	description ISL
no routing	no routing
vlan trunk native 1 tag	vlan trunk native 1 tag
vlan trunk allowed all	vlan trunk allowed all
lacp mode active	lacp mode active
lacp rate fast	lacp rate fast
interface 1/1/1	interface 1/1/1
no shutdown	no shutdown
mtu 9100	mtu 9100
description to SW3	description to SW3
lag 1	lag 1

interface 1/1/2	interface 1/1/2	
no shutdown	no shutdown	
mtu 9100	mtu 9100	
description to SW4	••••description to SW4 •••••	
lag 2	lag 2	
interface 1/1/7	interface 1/1/7	
no shutdown	••••no shutdown •••••••••	
vrf attach KA	vrf attach KA	
description keepalive link	description keepalive link	
ip address 192.168.0.0/31	ip address 192.168.0.1/31	
interface 1/1/8	interface 1/1/8	
no shutdown	no snutdown	
description ISI	description TSI	
lag 256		
interface 1/1/0	interface 1/1/9	
no shutdown	no shutdown	
mtu 9198	mtu 9198	
description ISL	description ISL	
lag 256	lag 256	
VSX	vsx	
system-mac 02:01:00:00:01:00	system-mac 02:01:00:00:01:00	
inter-switch-link lag 256	inter-switch-link lag 256	
role primary	role secondary	
keepalive peer 192.168.0.1 source 192.168.0.0	keepalive peer 192.168.0.0 source 192.168.0.1	
vrf KA	vrf KA	
vsx-sync aaa acl-log-timer bfd-global bgp	vsx-sync aaa acl-log-timer bfd-global bgp	
control-plane-acls copp-policy dhcp-relay dhcp-	control-plane-acls copp-policy dhcp-relay dhcp-	
server dhcp-snooping dns icmp-tcp lldp loop-	server dhcp-snooping dns icmp-tcp lldp loop-	
protect-global mac-lockout mclag-interfaces	protect-global mac-lockout mclag-interfaces	
neighbor ospi qos-global route-map silow-global	neighbor ospi qos-global route-map silow-global	
snmp ssh stp-global time vsx-global	snmp ssh stp-global time vsx-global	
: SM/3	: SWA	
hostname SW3	hostname SW4	
1		
vlan 1,100	vlan 1,100	
interface mgmt	interface mgmt	
no shutdown	no shutdown	
ip dhcp	ip dhcp	
interface lag 1	interface lag 1	
no shutdown	no shutdown	
no routing	no routing	
vlan trunk native 1	vlan trunk native 1	
vlan trunk allowed 100	vlan trunk allowed 100	
lacp mode active	lacp mode active	
Lacp rate fast	lacp rate fast	
interface 1/1/1	interface 1/1/1	
no snutdown	no shuldown	
vlan access 100	vlan access 100	
interface 1/1/8	interface 1/1/8	
no shutdown	no shutdown	
description to SW1	description to SW1	
lag 1	lag 1	
interface 1/1/9	interface 1/1/9	
no shutdown	no shutdown	
description to SW2	description to SW2	
lag 1	lag 1	

# Test #1: Layer2 connectivity between HostA and HostB

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

HostA	HostB
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	<pre>VPCS&gt; ip 10.10.100.12/24 10.10.100.1 Checking for duplicate address VPCS : 10.10.100.12 255.255.255.0 gateway 10.10.100.1</pre>
VPCS> show ip	VPCS> show ip

			0       0       0       0       0       0         0       0       0       0       0       0       0         0       0       0       0       0       0       0       0         0       0       0       0       0       0       0       0         0       0       0       0       0       0       0       0         0       0       0       0       0       0       0       0         0       0       0       0       0       0       0       0         0       0       0       0       0       0       0       0         0       0       0       0       0       0       0       0         0       0       0       0       0       0       0       0         0       0       0       0       0       0       0       0							VSX	Lab Layer	1 2		
NAME IP/MASK GATEWAY DNS MAC LPORT RHOST:PORT MTU	: VPCS[1] : 10.10.100.11/24 : 10.10.100.1 : : 00:50:79:66:68:07 : 20000 : 127.0.0.1:30000 : 1500		NAME IP/MASK GATEWAY DNS MAC LPORT RHOST:PO MTU	: : : RT	VPCS[1 10.10. 10.10. 20000 127.0. 1500	] 100.12, 100.1 79:66:0	/24 58:05							
Ping HostB fro	m HostA		• • • • • • • •	• • • •			• • • •	•••	•					
HostA														
VPCS> ping	10.10.100.12		000				• • •	• • •	• • • •				•	
84 bytes fr 84 bytes fr 84 bytes fr 84 bytes fr 84 bytes fr 84 bytes fr	om 10.10.100.12 icmp_seq=1 om 10.10.100.12 icmp_seq=2 om 10.10.100.12 icmp_seq=3 om 10.10.100.12 icmp_seq=4 om 10.10.100.12 icmp_seq=4	ttl=64 time=4 2 ttl=64 time=4 8 ttl=64 time=3 4 ttl=64 time=3 5 ttl=64 time=3	.431 ms .459 ms .723 ms .759 ms .633 ms			0     0     0     0       0     0     0     0     0       0     0     0     0     0       0     0     0     0     0       0     0     0     0     0       0     0     0     0     0       0     0     0     0     0       0     0     0     0     0								

#### Test #2: resiliency on shutting down interfaces

Right-click on SW3, select Graph Menu 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:

	Gr	aph Aruba	aCX-SW	3 1/1/8																	*	4) / X
k ilo bits/s	2.00 0.00 -2.00	ArubaCX-SW3	1/1/8	14:55:30	14:58:00	14:58:30	14:57:00	14:57:30	14:58:00	14:58:30	14:59:00	14-59:30	15:00:00	 15:01:00	15-01:30	15:02:00	15:02:30	15:03:00	15:03:30	Tue, A 15:04 	pr 06, 2021 :10 :eived nt	kilobits/s 0.99 -1.91
L																						
	Gr	aph Aruba	aCX-SW	3 1/1/9										 							*	♦ ,* ×
ilobits/s	Gr 2.00 0.00	aph Aruba ArubaCX-SW3	aCX-SW	3 1/1/9			/	\	_/					 	/					Tue, A 15:04 	# pr 08, 2021 10 ceived nt	kilobits/s 0.99 -0.99

The nominal traffic (and peaks) are due to LACP. If LACP rate slow is used instead of fast, the graphs would look like:

VSX Lab1
Layer2

	L (	Fraph Aru	ibaCX-	SW3 1/1	/8																									* 4	6 🥜 X
k ito bits/s	1.00 0.00 -1.00	ArubaCX-S	W3 1/1/8	13:05:00	13:05:30	D 13:1	00:00	13:08:30	13:07	:00	13:07:30	13:08:00	13:08	30 1	3:09:00	13:09:30	13:10:00	13:1	10:30	13:11:00	13:11:	30 13	8:12:00	13:12:30	13:13	3:00	13:13:30		Tue, Apr 06 13:13:54 — receive — sent • • •	, 2021 k ed ≫ + -	lobits/s 0.00 0.00
	▲ G	Graph Aru	ibaCX-S	SW3 1/1	/9																									* <	) 🧨 🗙
k ilo bits/a	1.00 0.00 -1.00	ArubaCX-5	W3 1/1/9	13:05:00	13:05:30	D 13:1	08:00	13:08:30	13:07	:00	13:07:30	13:08:00	13:08	:30 1	3:09:00	13:09:30	13:10:00	13:1	10:30	13:11:00	13:11:	30 13	8:12:00	13:12:30	13:13	3:00	13:13:30		Tue, Apr 06 13:13:54 receive sent	, 2021 k ed ⊪ + -	lobits/s 0.00 0.00
																••		••	••	••			••	• • •		•••	••	•••	••	• •	
Sta pin	art a Ig 10	gain 1 0.10.1	the p 100.1	oing to 12 –I	o Ho 1400	ostB D -t	froi	m Ho	ostA	wi	th la	rge c	lataç	gran	n siz	e like	): •	• • • • • •	• • • • • •	• • • • • •		<ul> <li>•</li> <li>•&lt;</li></ul>	• • • • • •			• • • • • •	• • • • • •			• • •	
Yo	u sł	ould	now	see	one	of th	ne li	ink b	eing	g lo	adeo	d fror	n thi	s pi	ng, ł	nere	nter	ace	1/1	/9	•••	•••	•••	•••	•••	•••	•••	•••	•••	•••	

	Graph Ar	ubaCX-S	N3 1/1/8																		* (	/ <b>/</b> X
k ilobils/s	ArubaCX- 2.00 -2.00 14:57:30	SW3 1/1/8	14:58:30	14:59:00	A 14:59:30	15:00:00	▲ 	15:01:00	15:01:30	15:02:00	15:02:30	15:03:00	15:03:30	15:04:00	15:04:30	15:05:00	15:05:30	15:08:00	15:08:30	Λ - //γ 15:07:00	Tue, Apr 08, 2021 15:07:17 — received — sent ≪ ▶ ≫ + -	ilobits/s 0.99 -0.99
	Graph Ar	ubaCX-S\	N3 1/1/9																		* <	∥
k ilo bits/s	ArubaCX- 10.0 -10.0 14:57:30	SW3 1/1/9	14:58:30	14:59:00	14:50:30	15:00:00	15:00:30	15:01:00	15:01:30	15:02:00	15:02:30	15:03:00	15:03:30	15:04:00	15:04:30	15:05:00	15:05:30	15:08:00	15:08:30	15:07:00	Tue, Apr 08, 2021 15:07:17 received sent	ilobits/s 1.0 -1.0

So based on the L2 hashing mechanism ping traffic from 10.10.100.11 to 10.10.100.12 is sent over interface 1/1/9.

Repeat the ping with -t option for continuous ping (until interrupted with CTRL+C): ping 10.10.100.12 -I 1400 -t

#### Shutdown interface 1/1/9 on SW3.

You should now see traffic moved from interface 1/1/9 to interface 1/1/8.

	Graph ArubaCX-SW3 1/1/8	* 🔷 🦯 X
k ilo bits/s	ArubaCX:SW3 1/1/8	Tue. Apr 08, 2021 15:09:09 Freceived sent 08:00 ≪ ▶ ▶ + - ♀
	Graph ArubaCX-SW3 1/1/9	* 🔹 🦨

This is expected as standard LAG ports mechanism. In production, failover of traffic from one link to another link of the same VSX LAG should happen in less than 200 milliseconds.

Restore (no shut) interface 1/1/9 on SW3 and interface 1/1/1 on SW2.

#### Test #3: VSX split resiliency on ISL cut

Make sure that the continuous ping is running between HostA and HostB, and that all links are in nominal state (up).

Create a VSX split event by shutting down ISL on SW1 (int lag 256, then shutdown).

SW1		SW2	
SW1 <b># show vsx status</b> VSX Operational State		SW2 <b># show vsx status</b> VSX Operational State	
ISL channel ISL mgmt channel	: In-Sync : operational	ISL channel ISL mgmt channel	: In-Sync : operational

		• • • •		• • • • • •	
Config Sync Statu	is : In-Sync		Config Sync Statu	is : In-Sync	
NAE	: peer_reac	hable	NAE	: peer_reacl	nable
HTTPS Server	: peer_reac	hable	HTTPS Server	: peer_reac	nable
Attribute	Local	Peer	Attribute	Local	Peer
TSL link	lag256	lag256	ISL link	lag256	 lag256
ISL version	2	2	ISL version	2	2
System MAC	02:01:00:00:01:00	•	System MAC	02:01:00:00:01:00	
02:01:00:00:01:00			02:01:00:00:01:00		
Platiorm Software Version	X86-64 Virtual 10 06 011	X86-64	Platiorm Software Version	X86-64 Virtual 10 06 011	X86-64
Virtual.10.06.0110	viicuai.i0.00.011	0	Virtual.10.06.0110	VIICUAI.10.00.011	
Device Role	primary	secondary	Device Role	secondary	• primary• • •
			· · · · · · · · · · · · · · · · · · ·	• • • • • • • • • • • • • • • •	) 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
SW1# show ver statu	s inter-switch-lin	k	SW2# show ver statu	inter-switch-lin	, , , , , , , , , , , , , , , , , , ,
State	: In-Sync	~	State	: In-Sync	
Last Connect Time	: Fri Åpr	2 13:29:41	Last Connect Time	: Fri Apr	2 13:29:40
2021			2021	• • • • • • • • • • • • • • • •	) • • • • • • • • • • • • • • • • • • •
Last Disconnect Tin	ne :		Last Disconnect Tim	ne	
Mgmt state	. up : operation	al	Mgmt state	•••••• <u>••up</u> ••••••	
<u></u>	. 0201401011		<u></u>	. Speracion	
Inter-switch link S	Statistics		Inter-switch link S	Statistics	, , , , , , , , , , , , , , , , , , ,
Hello Dackota Tr	. 340001		Hello Dackoto Tr		)
Hello Packets Rx	: 340981 : 340981		Hello Packets Rx	: 341031	
Data Packets Tx	: 11859		Data Packets Tx	: 11750	,
Data Packets Rx	: 11748		Data Packets Rx	: 11884	
Mgmt Packets Tx	: 2883972		Mgmt Packets Tx	: 1527779	· · · · · · · · · · · · · · · · · · ·
Mgmt Packets Rx	: 1526638		Mgmt Packets Rx	: 2883846	
Mgmit Packet Drops	: 0		Mgmit Packet Drops	: 0	
SW1(config-lag-if) SW1(config-lag-if) SW1(config-lag-if) SW1# show vsx statu VSX Operational Sta	shut end ns		SW2# <b>show vsx statu</b> VSX Operational Sta	15 Ate	
		_			
ISL channel	: Out-Of-Sy	nc tah link dawa	ISL channel	: Out-Of-Syn	10 - ch link down
Config Sync Stati	: <u>Inter_swi</u> is : Out-Of-Sv	nc	Config Sync Statu	: Inter_SWI : Out-Of-Svi	
NAE	: peer unre	achable	NAE	: peer unrea	achable
HTTPS Server	: peer_unre	achable	HTTPS Server	: peer_unrea	achable
Attribute	Local	Peer	Attribute	Local	Peer
TSL link	 lag256		ISL link	lag256	
ISL version	2		ISL version	2	
System MAC	02:01:00:00:01:00		System MAC	02:01:00:00:01:00	
02:01:00:00:01:00			02:01:00:00:01:00		
Platiorm Software Version	X86-64 Virtual 10 06 011	0	Platform Software Version	X86-64 Wirtual 10 06 011	
Device Role	primary	0	Device Role	secondarv	
SW1# sh vsx status	inter-switch-link		SW2# sh vsx status	inter-switch-link	
State Last Connect Time	: Out-Of-Sy	nc 2 13•29•41	State Last Connect Time	: Out-Of-Syr	1C 2 13•29•40
2021	. LII ADI	2 13.23.41	2021	• LIT WHI	- +3.23.70
Last Disconnect Tim	ne : Tue Apr	6 13:11:58	Last Disconnect Tim	ne : Tue Apr	5 13:12:03
2021	-		2021	-	
Link Status	: down	بال المراجع	Link Status	: down	
Mgmt state	: inter_swi	tcn_link_down	Mgmt state	: inter_swi	cn_link_down
Inter-switch link S	Statistics		Inter-switch link S	Statistics	
Hello Packets Tx	: 344536		Hello Packets Tx	: 344541	
Hello Packets Rx	: 344536		Hello Packets Rx	: 344540	
Data Packets Tx	: 12003		Data Packets Tx	: 11992	
Data Packets Rx	: 11991		Data Packets Rx	: 12027	

Mgmt Packets Tx : 2924516		Mqmt Packets Tx • • • • • • • • 1552118
Momt Packets Rx : 1551344		Momt Packets Rx : 2923783
Manda Davida Davida		
Mgmt Packet Drops : U	• • • •	Mgmt Packet Drops : U
SW1 # show wer brief		SW2 the show way brief
SWIT SHOW VSA DITEL	• • •	SW2 # SHOW VSA DITEL
ISL State	: Out-Of-\.	ISL State : Out-Of-
Sync	•	Sync
Device State	: Split-	Device State : Split-
System-Primary		System-Secondary
Keepalive State	:	Keepalive State Keepalive-
Keepalive-Established		Established
Device Role	: Primary	Device Role
Number of Multi-chassis LAG interfaces	: 2	Number of Multi-chassis LAG interfaces : 2 · · · ·

#### Check interface state:

SW1									• • • •
SW1# show	interfa	ce brie:	£				• • • • • •	•       •	••••
Port	Native VLAN	Mode	Туре	Enabled	Status	Reason	Speed (Mb/s)	Description	
1/1/1 1/1/2 1/1/3 1/1/4 1/1/5 1/1/6 1/1/7 1/1/8 1/1/9	1 1   1	trunk trunk routed routed routed trunk trunk	      	yes no no no yes yes	up up down down down up down	Administratively down Administratively down Administratively down Administratively down Administratively down	1000 1000   1000 1000	to SW3 to SW4   keepalive link ISL	
lag1 lag2 lag256	1 1 1	trunk trunk trunk		yes yes yes no	up up down		1000 1000 auto	SW3 VSX LAG SW4 VSX LAG ISL	~ • • • •

#### SW2

SW2# show interface brief								
Port	Native VLAN	Mode	Туре	Enabled	Status	Reason	Speed (Mb/s)	Description
1/1/1 1/1/2 1/1/3 1/1/4 1/1/5 1/1/6 1/1/7 1/1/8 1/1/9 lag1 lag2	1 1   1 1 1 1	trunk trunk routed routed routed trunk trunk trunk trunk		yes no no no yes yes yes yes yes	down down down down down up up up down down	Administratively down Administratively down Administratively down Administratively down	1000 1000   1000 1000 1000 auto auto	to SW1 to SW2   keepalive link ISL ISL SW3 VSX LAG SW4 VSX LAG
Lag256	T	trunk		yes	blocked		auto	ISL

SW2# show interface lag 1

Aggregate lag1 is down	
Admin state is up	
State information : Disable	d by aggregation
Description : SW3 VSX LAG	
MAC Address	: 08:00:09:54:97:83
Aggregated-interfaces	: 1/1/1
Aggregation-key	: 1
Aggregate mode	: active
Speed	: 0 Mb/s
L3 Counters: Rx Disabled, T:	x Disabled
qos trust none	
VLAN Mode: native-untagged	
Native VLAN: 1	
Allowed VLAN List: 100	

		VSX Lab1 Layer2
Rv		
Tx	29090 total packets 0 unicast packets 0 multicast packets 0 broadcast packets 0 errors 0 CRC/FCS 29245 total packets	9183765 total bytes 0 dropped 0 pause 9204402 total bytes
	0 unicast packets 0 multicast packets 0 broadcast packets 0 errors 0 collision	0 dropped 0 pause

<u>Note</u>: On the VSX secondary (SW2) the Reason for Down Status should be "Disabled by VSX". (CX Simulator shows none due to a limitation).

The traffic from HostA to HostB is not interrupted and moved from interface 1/1/9 to interface 1/1/8.



Restore the ISL LAG (no shutdown of interface lag 256 on SW1)

C\W/1			C/W/D							
3001			3002							
SWI# show vsx statu	S		SW2# show vsx status							
VSX Operational Sta	te		VSX Operational State							
ISL channel ISL mgmt channel Config Sync Statu NAE HTTPS Server	: In-Sync : operational s : In-Sync : peer_unreac : peer_reacha	hable ble	ISL channel : In-Sync ISL mgmt channel : operational Config Sync Status : In-Sync NAE : peer_unreachable HTTPS Server : peer_reachable							
Attribute	Local	Peer	Attribute	Local	Peer					
ISL link	 lag256	lag256	ISL link	lag256	lag256					
Suston MAC	2 02.01.00.00.01.00	2	Suston MAC	2	2					
02.01.00.00.01.00	02.01.00.00.01.00		02.01.00.00.01.00	02.01.00.00.01.00						
Dlatform	X86-64	X86-64	Dlatform	X86-64	X86-64					
Software Version	Virtual 10 06 0110	200 04	Software Version	Virtual 10 06 0110	200 04					
Virtual 10.06.0110	VII Cuul. 10.00.0110		Virtual. 10.06.0110	VII CUUI . 10.00.0110						
Device Role	primary	secondary	Device Role	secondary	primary					
SW1# sh vsx status Configured linkup d : 180 seconds Initial sync status : Completed Delay timer status : Completed Linkup Delay time l : Interfaces that wil timer expires : Interfaces that are :	linkup-delay elay-timer eft l be brought up afte excluded from delay	r delay timer	<pre>SW2# sh vsx status linkup-delay Configured linkup delay-timer : 180 seconds Initial sync status : In-progress Delay timer status : Waiting-to-start Linkup Delay time left : Interfaces that will be brought up after delay timer expires : lag1-lag2 Interfaces that are excluded from delay timer .</pre>							

As soon as the ISL is restored the VSX status shows both CX nodes in the VSX cluster. However, the VSX secondary is not yet

forwarding the traffic as the MAC/ARP are not yet synchronized. The VSX secondary will be forwarding as soon as the VSX linkup delay timer expires (timer to guarantee that MAC/ARP area synchronized between VSX primary and VSX secondary).

-

SW2																						
W2# show l	acp interfa	ice							• • •	•••	•											
State abbr	eviations :						• • •			• •	•••											
A - Active	P -	• Passi	ve	F - Ag	gregable I - Indiv	idual				•••	•••											
S - Short-	timeout L -	- Long-	timeou	ıt N - In	Sync 0 - Outof	Sync																
C - Collec	tina D-	Distr	ibutin	a						• •	• •	• •	• •	•								
X - State	m/c expired	1			fault neighbor sta	+ -				• •	• •	• •	• •	• •								
n beace	m/ e enpiree			<u> </u>	idaite neignbor bea		• • • •		• • •	• •	• •	• •	• •	• •	•••							
Natan data	41 a a E a 11					• • • • •	•••				•••		•••									
ACTOR deta	IIS OI AII	Interi	aces:																			
					• •					• •	• •	• •		• •		• •					• •	
Intf	Aggr	Port	Port	State	System-ID 🥚	System	Aggr	Forwa	rdi	ng 🛛	• •	• •	• •	• •	• •	• •	• •	• • •	•	• •	• •	•
	Name	Id	Pri			Pri	Key	State		• •	• •	• •	• •	• •	• •	• •	• •	•	•	• •	• •	•
										÷:	•••	•••	•••	•••	•••	•••				•••	•••	
1/1/1	lag1(mc)							down														
1/1/2	lag2(mc)					• •		down		• •	• •	• •	• •	• •	• •						• •	•
1/1/8	lag256	9	1	ASENCD	08:00:09:54:97:83	65534	256	າມກ		• •	• •	• •	• •	• •	• •	• •		• • •	•	• •	• •	•
1/1/9	lag256	10	1	ASENCD	08.00.09.54.97.83	65534	256	110 110		• •	• •	• •	• •	• •	• •	• •	• •	• • •	•	• •	• •	•
1/1/2	1ag250	10	Ŧ	ASPNCD	00.00.09.94.97.09	05554	200	up	•••	•••	•••	•••	••	•••	•••	•••	•••			• •	•••	•
Partner de	tails of al	l inte.	erfaces	:				•		• •	• •	• •	• •	• •	• •	• •					• •	•
										• •	• •	• •	• •	• •	• •	• •	• •	• • •	•	• •	• •	•
Intf	Aggr	Port	Port	State	System-ID	System	Aggr		•	• •	• •	• •	• •	• •	• •	• •	••	• • •	•	• •	• •	•
	Name	Id	Pri			Pri	Kev				•••	•••	•••	•••	•••	•••	•••			•••	•••	
1/1/1	lag1(mc)																					•
1/1/2	1 ag2 (mg)														• •	• •					• •	•
1/1/0	1 a y 2 (IIIC)	0	1	ACENCE	00 00 00 17 55 05	65524	050										•	• • •	•	• •	• •	•
1/1/8	1ag256	9	1	ASFNCD	u8:00:09:d/:51:01	05534	256													• •	• •	•
1/1/9	Lag256	10	1	ASFNCD	U8:00:09:d7:5f:0f	65534	256															

If the linkup delay-timer status is stuck to "Waiting-to-start", it might be necessary to bounce (shut/no shut) the ISL LAG on SW2 to unfreeze the timer (AOS-CX Simulator intermittent issue).



									• •													
									• •													
Configured	d linkup del	lay-tim	ler			• : 180	seco	nds 🔹	• •													
Initial sy	ync status					: Com	plete	d	•••													
Delay timer status							nlete	d	•••													
Linkun Del	lav time let	F+					p1000															
Tatawfaaa	ray crime rei	LL las lasas			A . 7																	
interlaces	s that will	be bro	ugnt t	ip aller	delay timer expire	S .				• •												
Interfaces	s that are e	exclude	d from	n delay t	imer					• •	•											
						••••	• • •	• • • •	• • •	• •	• •											
						• • • • •	• • •	• • • •	• • •	• •	•••											
SW2# show	lacp int																					
	1																					
State abby	reviations .	•											•••									
		· Dogoj		E 7.00	gragable T Today					• •	• •	• •	• •	• •								
A - ACLIVE	e P •	- Passi	ve	F - Ag	gregable i - indiv	Idual				• •	• •	• •	• •	• •								
S - Short-	-timeout L -	- Long-	timeou	it N - In	Sync 0 - Outor	Sync 🔹	• • •	• • • •	• • •	• •	• •	• •	• •	• •								
C - Collec	cting D -	- Distr	ibutir	ıg	• • •	• • • • •	• • •	• • • •	• • •	•••	•••	•••	•••				• • •	•				
X - State	m/c expired	b		E - De	fault neighbor sta	te		••••												1		
Actor deta	ails of all	interf	aces:									• •	• •	•			• •		• •	• •		
necor acco	arro or arr									• •	• •	• •	• •	• •			• •	• •	• •	• •	• •	
Tote	7	Dent	Dent	0++++	Count and TD			• • •	••••	• •	• •	• •	• •	• •	. • •	. • •	• •	• •	• •	• •	• •	. • .
INCL	Aggr	Port	POLL	State	System-ID	system	Aggr	FOIW	arar	ng	• •	• •	• •	• •		. • •	• •	•••	• •	• •	• •	
	Name	Id	Pri			Pri	Key	State	9		•••		•••				•••	•••	•••	•••	•••	
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	ASENCD	08:00:09:54:97:83	65534	256	ຳເກົ	• • •	• •	• •	• •	• •	• •			• •	• •	• •	• •	• •	. • .
1/1/0	120256	10	1	ASENCD	08.00.09.54.97.83	65534	256	up up	• • •	• •	• •	• •	• •	• •	. • •		• •	• •	• •	• •	• •	. • .
1/1/2	Tag200	10	Ŧ	ASPINCE	00.00.09.94.97.03	05554	200	up 🍵	•••	•••	•••	•••	•••	•			• • •	•••	•••	• •	•••	
Partner de	etails of al	ll inte	rfaces	:													• •			• •		
													• •	• •			• •	• •	• •	• •	• •	
Intf	Aqqr	Port	Port	State	System-ID	System	Aqqr							• •		. • •	• •	• •	• •	• •	• •	. • .
	Name	ЪТ	Pri			Pri	Kev								. • •		• •	• •	• •	• •	• •	. • •
																	•	••	•••	•••	•••	
1 /1 /1		1.0	1	ACENCE	00.00.00.Eb.7.0.1	CEE24	1			_												
$\perp / \perp / \perp$	iagi(mc)	TÜ	1	ASFNCD	u8:00:09:56:/e:2d	65534	1															
1/1/2	⊥ag2(mc)	10	1	ASFNCD	U8:00:09:ed:b5:6e	65534	1															
1/1/8	lag256	9	1	ASFNCD	08:00:09:d7:5f:0f	65534	256															
1/1/9	lag256	10	1	ASFNCD	08:00:09:d7:5f:0f	65534	256															

As soon as the VSX linkup delay timer on the VSX secondary is "Completed", the traffic is now restored on the link between SW3 and SW2 as is was in nominal situation (before VSX split).

	Graph ArubaCX-SW3 1/1/8	* 🔷 🧨 X
k ilo bits/s	ArubaCX-SW3 11/18	Tue. Apr 06, 2021 15:32:56 ■ received 1.0 = sent -1.0 ◄ ▷ ▷ + - ↓
	Graph ArubaCX-SW3 1/1/9	* 🗘 🧨 X

When the VSX peer joins back the VSX cluster, there is no reboot required.

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 KA	0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0	
!		
vlan 1	\ 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0	
Vian 100		
interface momt		
no shutdown		
in dhen	<b>1 1 1 1 1 1 1 1 1 1</b>	
interface lag 1 multi-chassis		
no shutdown	` • • • • • • • • • • • • • • • • • • •	, <b></b> .
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 snutdown deceription SMA NEX LAC		
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		
vian trunk native i tag		
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		
Mcu 9100 description to SM4		
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		
Mcu 9198 description ISI		
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-IINK 1ag 256		
LULE PILHALY keenalive neer 192 168 0 1 source 192 168 0	0 vrf KA	
vsx-svnc aaa acl-log-timer bfd-global bon co	ontrol-plane-acls copp-policy dhep-relay dhep-server	
dhcp-snooping dns icmp-tcp lldp loop-protect-alc	obal mac-lockout mclag-interfaces neighbor ospf gos-	
global route-map sflow-global snmp ssh stp-globa	al time vsx-global	

VSX Lab1 Layer2

	) 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0	
SW2		
hostname SW2		
!		
vrf KA		
!		
vlan 1		
vlan 100		
vsx-sync	$\ \ \ \ \ \ \ \ \ \ \ \ \ \ \ \ \ \ \ $	
interface mgmt		
no shutdown	` 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0	
ip dhcp		
interface lag 1 multi-chassis		• • .
no shutdown		
description SW3 VSX LAG	· · · · · · · · · · · · · · · · · · ·	
no routing		
vian trunk native i		• • • • • • • • •
Vian trunk allowed 100		
lacp mode active		
lacp fale last		
Interface lag 2 multi-chassis		
deceription CHA NEX LAC		
description SW4 VSA LAG		
ulan trunk natiwo 1	***************************************	
vian trunk allowed 100		• • • • • • • •
laco mode active		
lach 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		
vri attach KA		
in 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 c	ontrol-plane-acls copp-policy dhcp-relay dhcp-server	
dhcp-snooping dns icmp-tcp lldp loop-protect-gl	obal mac-lockout mclag-interfaces neighbor ospf qos-	
global route-map sflow-global snmp ssh stp-globa	al time vsx-global	

	VSX Lab1	
	l aver?	
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		
mt 11 9100		
deceription to SW1		
les 1		
Idg I		
interiace 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



#### 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