

Contents

[Introduction](#)

[Prerequisites](#)

[Requirements](#)

[Components Used](#)

[Configure](#)

[Network Diagram](#)

[Configurations](#)

[Verify](#)

[Troubleshoot](#)

[Related Information](#)

Introduction

This document describes how to deploy an Application Virtual Switch (AVS) switch with an Adaptive Security Virtual Appliance (ASAv) single firewall in Routed/GOTO mode as a L4-L7 Service Graph between two End Point Groups (EPGs) to establish client-to-server communication using ACI 1.2(x) Release.

Prerequisites

Requirements

Cisco recommends that you have knowledge of these topics:

- Access Policies configured and interfaces up and in service
- EPG, Bridge Domain (BD) and Virtual Routing and Forwarding (VRF) already configured

Components Used

The information in this document is based on these software and hardware versions:

Hardware & Software:

- UCS C220 - 2.0(6d)
- ESXi/vCenter - 5.5
- ASAv - asa-device-pkg-1.2.4.8
- AVS - 5.2.1.SV3.1.10
- APIC - 1.2(1i)
- Leaf/Spines - 11.2(1i)
- Device packages *.zip already downloaded

Features:

- AVS
- ASAv

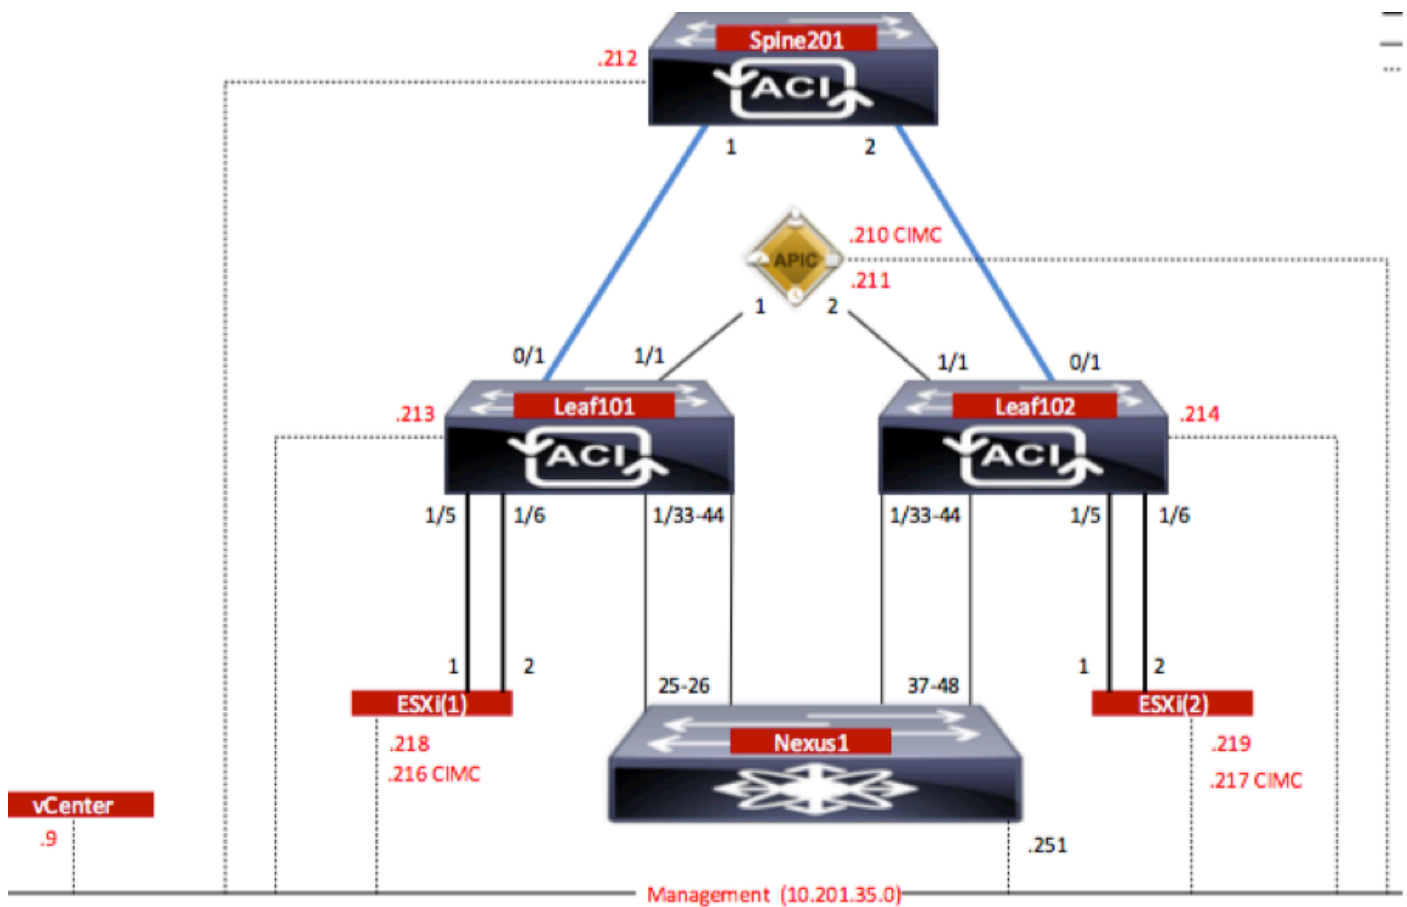
- EPGs, BD, VRF
- Access Control List (ACL)
- L4-L7 Service Graph
- vCenter

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, make sure that you understand the potential impact of any command.

Configure

Network Diagram

As shown in the image,



Configurations

AVS Initial Setup creates a VMware vCenter Domain (VMM integration)2

Note:

- You can create multiple datacenters and Distributed Virtual Switch (DVS) entries under a single domain. However, you can have only one Cisco AVS assigned to each datacenter.
- Service graph deployment with Cisco AVS is supported from Cisco ACI Release 1.2(1i) with Cisco AVS Release 5.2(1)SV3(1.10). The entire service graph configuration is performed on

the (Cisco APIC).

- Service Virtual Machine (VM) deployment with Cisco AVS is supported only on Virtual Machine Manager (VMM) domains with Virtual Local Area Networks (VLAN) encapsulation mode. However, the compute VMs (the provider and consumer VMs) can be part of VMM domains with Virtual Extensible LAN (VXLAN) or VLAN encapsulation.
- Also note that if local switching is used, Multicast address and pool are not required. If no local switching is selected, then Multicast pool has to be configured and the AVS Fabric-wide multicast address should not part of the Multicast pool. All traffic originated from the AVS will be either VLAN or VXLAN encapsulated.

Navigate to **VM Networking > VMWare > Create vCenter Domain**, as shown in the image:

Create vCenter Domain

Specify vCenter domain users and controllers

Virtual Switch Name: AVS

Virtual Switch: VMware vSphere Distributed Switch **Cisco AVS**

Switching Preference: No Local Switching **Local Switching**

Encapsulation: VLAN VXLAN

Associated Attachable Entity Profile: AEP-AVS

VLAN Pool: VlanPool-AVS(dynamic)

Security Domains: × +

Name	Description
------	-------------

vCenter Credentials: × +

Profile Name	Username	Description
vCenterCredentials	root	

vCenter: × +

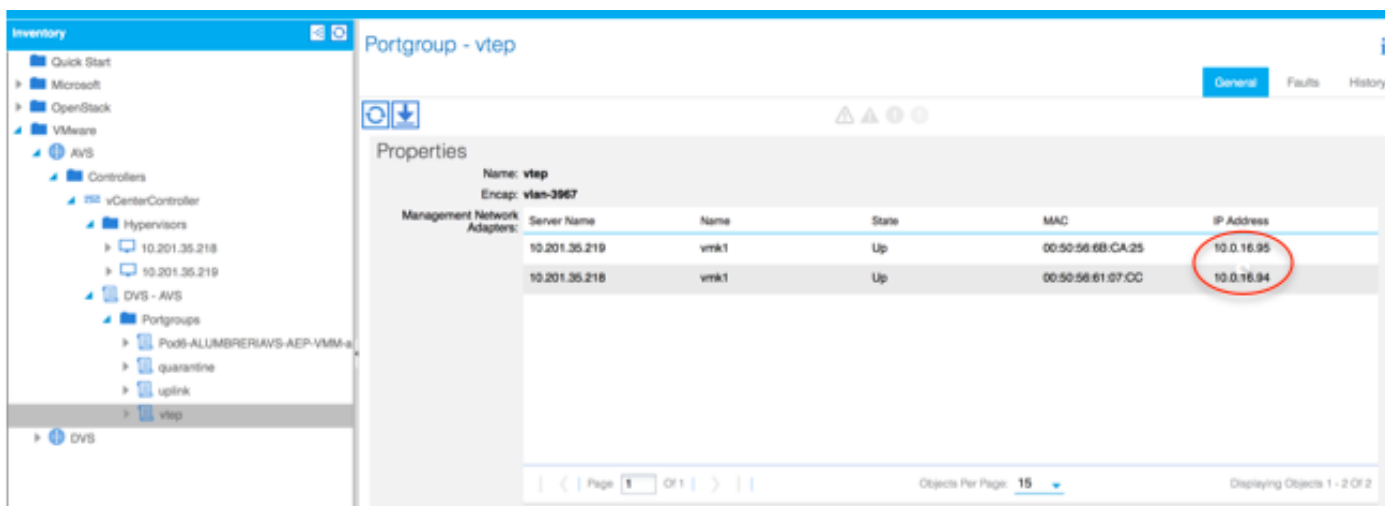
Name	IP	Type	Stats Collection
vCenterController	10.201.35.9	vCenter	Disabled

If you're using Port-channel or VPC (Virtual Port-channel) it is recommended to set the vSwitch policies to use Mac Pinning.

After this, APIC should push AVS switch configuration to vCenter, as shown in the image:



On APIC you can notice that a VXLAN Tunnel Endpoint (VTEP) address is assigned to the VTEP port-group for AVS. This address is assigned no matter what Connectivity mode is used (VLAN or VXLAN)



Install the Cisco AVS software in vCenter

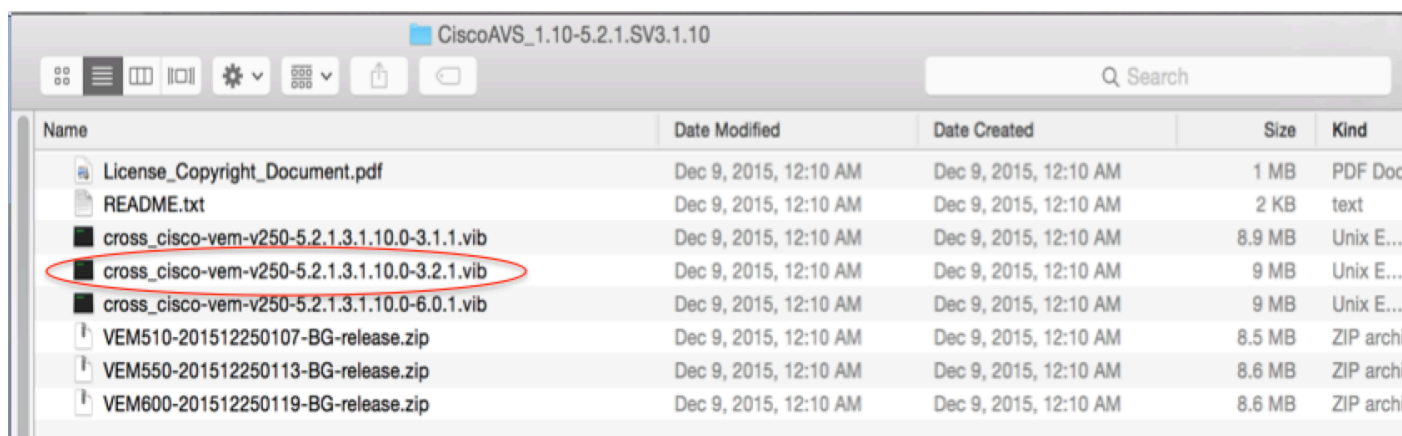
- Download (VIB) from CCO using this [link](#)

Note: In this case we are using ESX 5.5, Table 1, shows the Compatibility matrix for ESXi 6.0, 5.5, 5.1, and 5.0

Table 1 - Host Software Version Compatibility for ESXi 6.0, 5.5, 5.1, and 5.0

VMware	VIB	VEM Bundle	Windows VC Installer	Linux vCenter Server Appliance
ESXi 6.0	cross_cisco-vm-v250-5.2.1.3.1.10.0-6.0.1.vib	VEM600-201512250119-BG-release.zip (Offline) VEM600-201512250119-BG (Online)	6.0	6.0
ESXi 5.5	cross_cisco-vm-v250-5.2.1.3.1.10.0-3.2.1.vib	VEM550-201512250113-BG-release.zip (Offline) VEM550-201512250113-BG (Online)	5.5	5.5
ESXi 5.1	cross_cisco-vm-v250-5.2.1.3.1.10.0-3.1.1.vib	VEM510-201512250107-BG-release.zip (Offline) VEM510-201512250107-BG (Online)	5.1	5.1
ESXi 5.0	cross_cisco-vm-v250-5.2.1.3.1.10.0-3.0.1.vib	VEM500-201512250101-BG-release.zip (Offline) VEM500-201512250101-BG (Online)	5.0	5.0

Within the ZIP file there are 3 VIB files, one for each of the ESXi host versions, select the one appropriate for ESX 5.5, as shown in the image:



- Copy the VIB file to ESX Datastore - this can be done via CLI or directly from vCenter

Note: If a VIB file exists on the host, remove it by using the **esxcli software vib remove**

command.

esxcli software vib remove -n cross_cisco-vem-v197-5.2.1.3.1.5.0-3.2.1.vib

or by browsing the Datastore directly.

- Install the AVS software using the following command on the ESXi host:

esxcli software vib install -v /vmfs/volumes/datastore1/cross_cisco-vem-v250-5.2.1.3.1.10.0-3.2.1.vib --maintenance-mode --no-sig-check

```
~ # esxcli software vib install -v /vmfs/volumes/datastore1/cross_cisco-vem-v250-5.2.1.3.1.10.0-3.2.1.vib --maintenance-mode --no-sig-check
Installation Result
Message: Operation finished successfully.
Reboot Required: false
VIBs Installed: Cisco_bootbank_cisco-vem-v250-esx_5.2.1.3.1.10.0-3.2.1
VIBs Removed: Cisco_bootbank_cisco-vem-v197-esx_5.2.1.3.1.5.0-3.2.1
VIBs Skipped:
~ # vem status

VEM modules are loaded


Switch Name      Num Ports  Used Ports  Configured Ports  MTU  Uplinks
vSwitch0         5632       8           128               1500 vmnic0
DVS Name          Num Ports  Used Ports  Configured Ports  MTU  Uplinks
DVS               5632       10          512               9000 vmnic5,vmnic4

VEM Agent (vemdpa) is running

~ #
```

- Once the Virtual Ethernet module (VEM) is up, you can add Hosts to your AVS:

In the Add Host to vSphere Distributed Switch dialog box, choose the virtual NIC ports that are connected to the leaf switch (In this example you move only vmnic6), as shown in the image:



- Click **Next**
- In the Network Connectivity dialog box, click **Next**
- In the Virtual Machine Networking dialog box, click **Next**
- In the Ready to Complete dialog box, click **Finish**

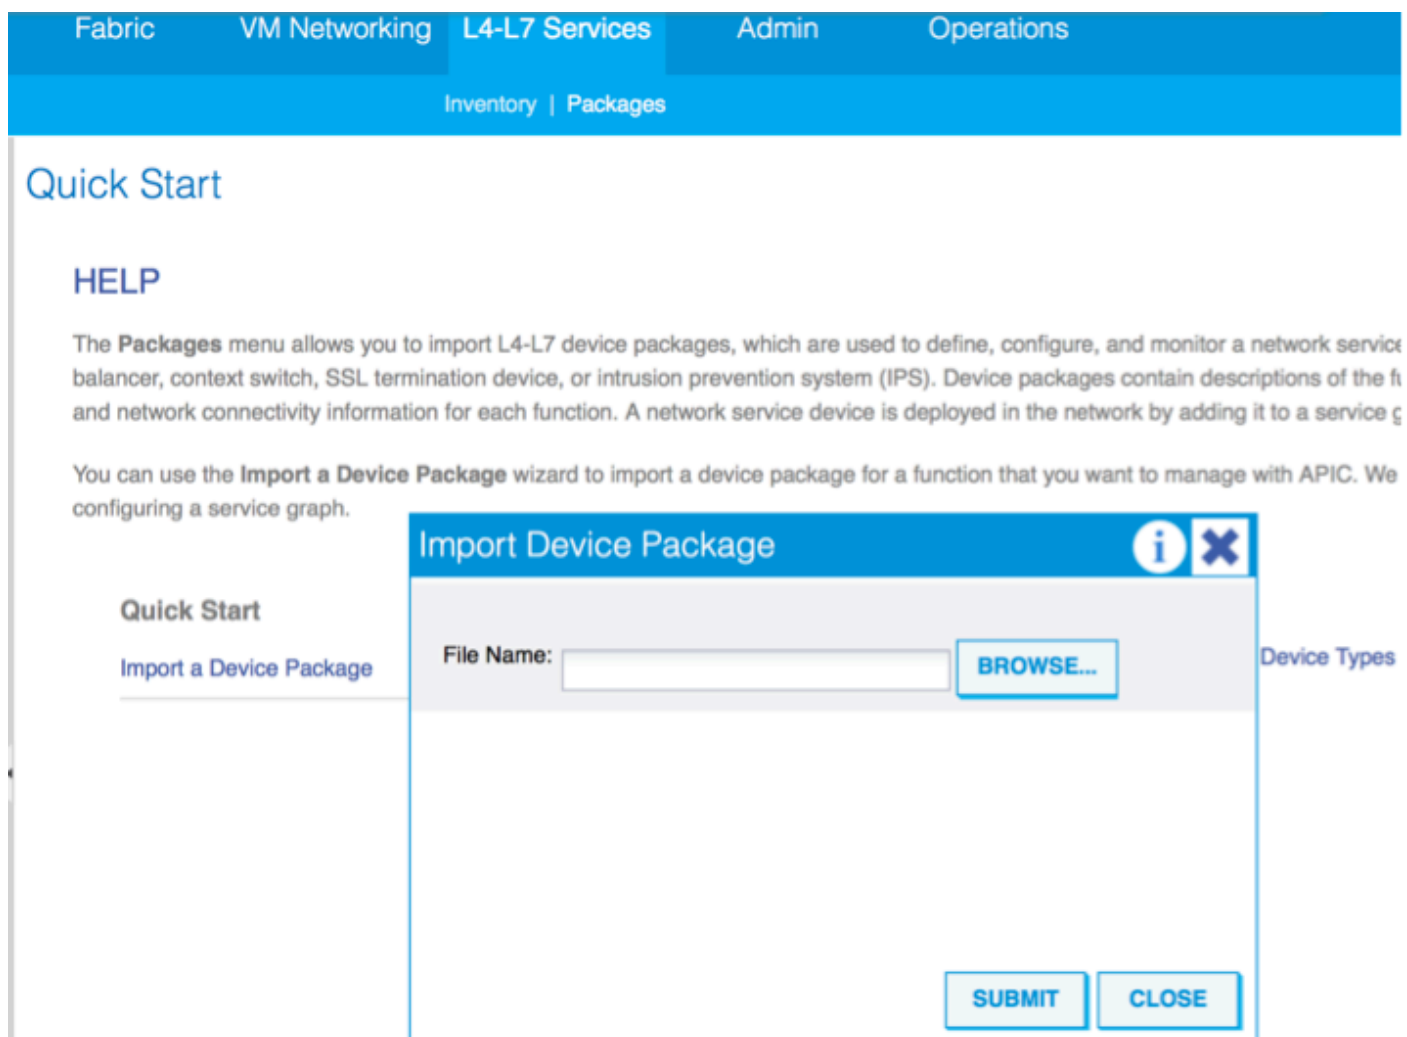
Note: If multiple ESXi hosts are used, all of them need to run the AVS/VEM so they can be managed from Standard switch to DVS or AVS.

With this, AVS integration has been completed and we are ready to continue with L4-L7 ASAv deployment:


ASAv Initial Setup

- Download Cisco ASAv Device Package and import it into APIC:

Navigate to **L4-L7 Services > Packages > Import Device Package**, as shown in the image:



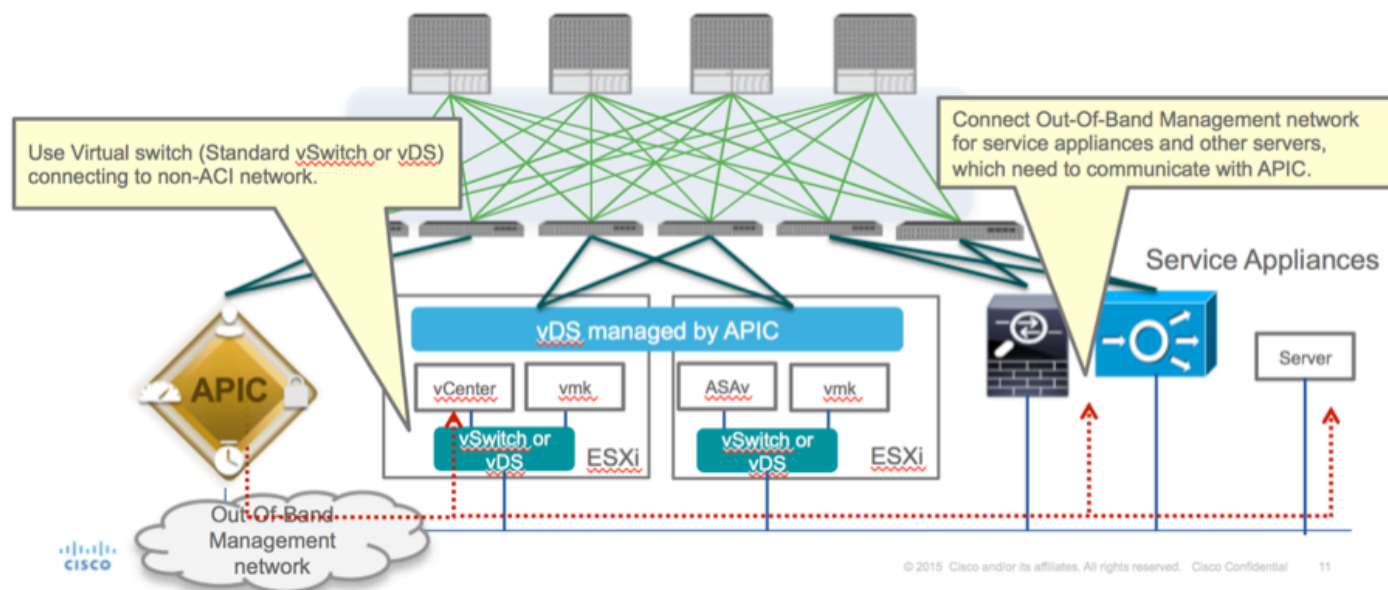
- If everything works well, you can see the imported device package expanding L4-L7 Service Device Types folder, as shown in the image:



Before you continue, there are few aspects of the installation that need to be determined before the actual L4-L7 integration is performed:

There are two types of Management networks, In-Band Management and Out-Of-Band (OOB), these can be used to manage devices that are not part of the basic Application Centric Infrastructure (ACI) (leaf, spines nor apic controller) which would include ASAv, Loadbalancers, etc.

In this case, OOB for ASAv is deployed with the use of Standard vSwitch. For bare metal ASA or other service appliances and/or servers, connect the OOB Management port to the OOB switch or Network, as shown in the image.



ASAv OOB Mgmt Port management connection needs to use ESXi uplink ports to communicate with APIC via OOB. When mapping vNIC interfaces, Network adapter1 always matches the Management0/0 interface on the ASAv, and the rest of the data plane interfaces are started from Network adapter2.


The Table 2 shows the concordance of Network Adapter IDs and ASAv interface IDs:

Table 2

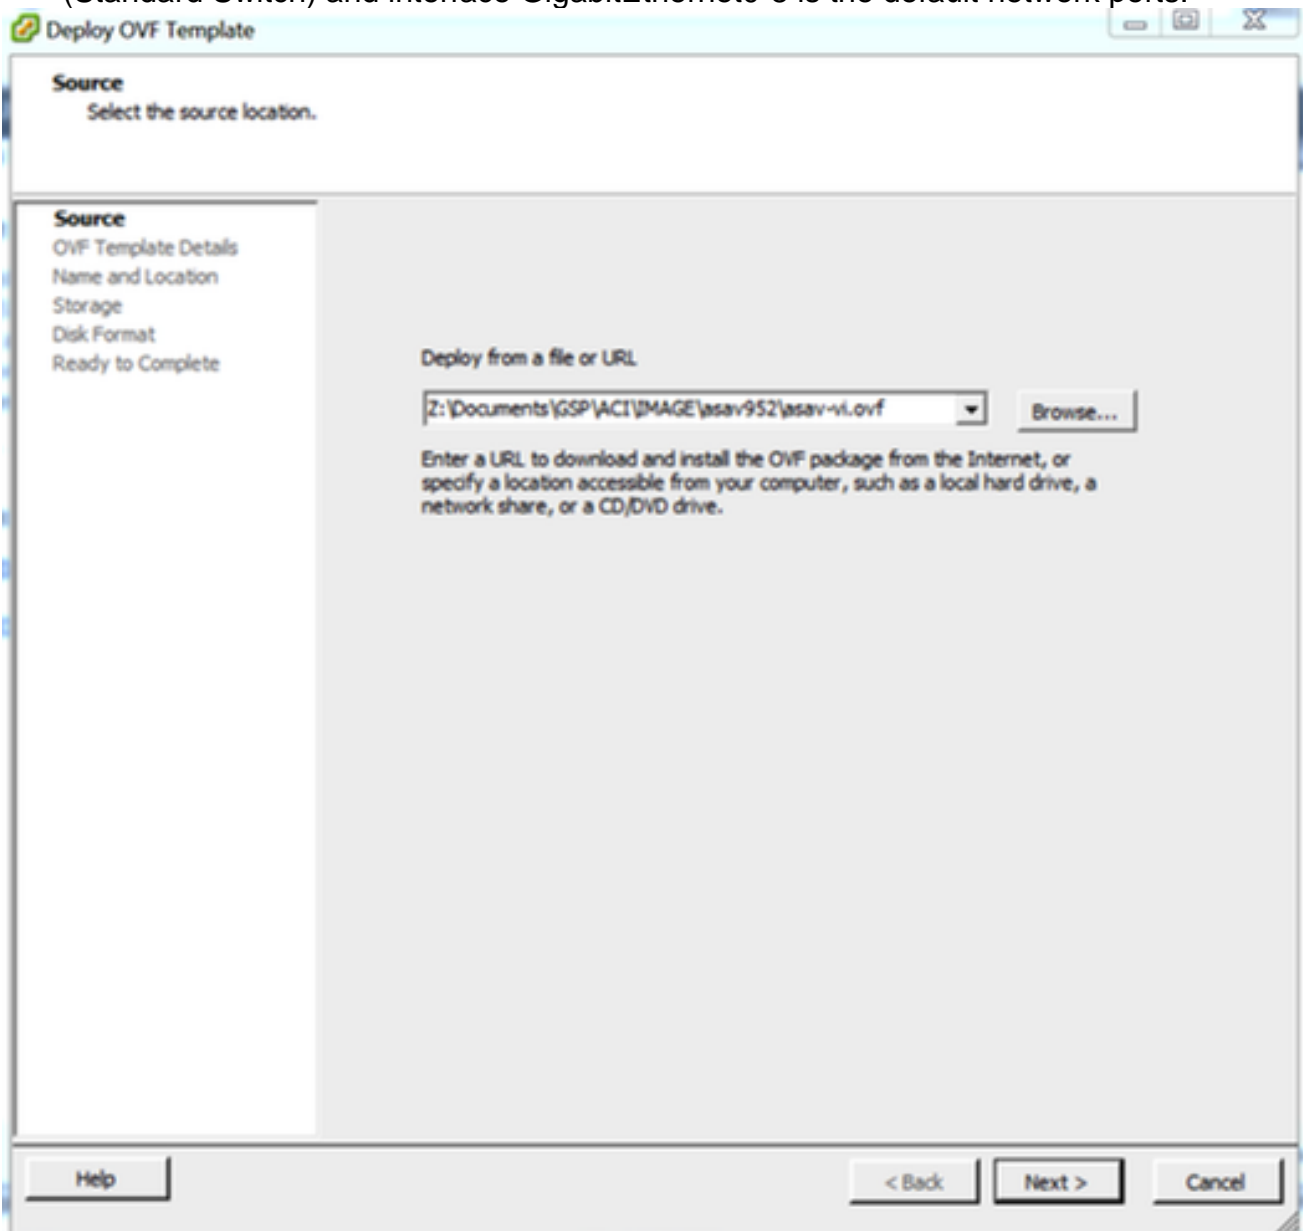
Network Adapter ID	ASAv Interface ID
Network Adapter 1	Management0/0
Network Adapter 2	GigabitEthernet0/0
Network Adapter 3	GigabitEthernet0/1
Network Adapter 4	GigabitEthernet0/2
Network Adapter 5	GigabitEthernet0/3
Network Adapter 6	GigabitEthernet0/4
Network Adapter 7	GigabitEthernet0/5
Network Adapter 8	GigabitEthernet0/6
Network Adapter 9	GigabitEthernet0/7
Network Adapter 10	GigabitEthernet0/8

- Deploy the ASAv VM through the wizard from **File>Deploy OVF (Open Virtualization Format) Template**
- Select **asav-esxi** if you want to use standalone ESX Server or **asav-vi** for vCenter. In this

case, vCenter is used.



- Go through the installation wizard, accept terms and conditions. In the middle of the wizard you can determine several options like hostname, management, ip address, firewall mode and other specific information related to ASA. Remember to use OOB management for ASA, as in this case you need to keep interface Management0/0 while you use the VM Network (Standard Switch) and interface GigabitEthernet0-8 is the default network ports.



- Back to the ASA GUI, the deployment is finished, initial configuration is done. Configure admin username and password. This username and password is used by

username admin password <device_password> encrypted privilege 15

```
ASAv-w-AVS(config)# username admin password C1sc0123 privilege 15
ASAv-w-AVS(config)# wr mem
Building configuration...
Cryptochecksum: d491b980 86fa522f 6f937baf b5bfb318

7977 bytes copied in 0.250 secs
[OK]
ASAv-w-AVS(config)# ping 10.201.35.211
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.201.35.211, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/10 ms
ASAv-w-AVS(config)# _
```

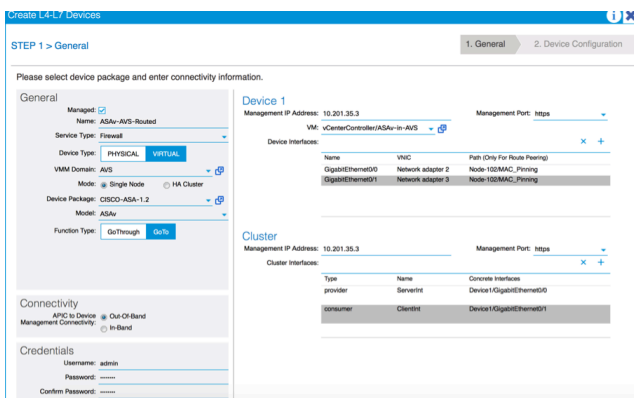
Additionally, from Global configuration mode enable http server:

http server enable

http 0.0.0.0 0.0.0.0 management

L4-L7 for ASAv Integration in APIC:

- Log in to the ACI GUI, click on the Tenant where the service graph will be deployed. Expand L4-L7 services at the bottom of the navigation pane and right click on **L4-L7 Devices** and click on **Create L4-L7 devices** to open the wizard
- For this implementation, following settings will be applied:
 - Managed Mode
 - Firewall Service
 - Virtual Device
 - Connected to AVS domain with a Single Node
 - ASAv Model
 - Routed mode (GoTo)
 - Management Address (has to match the previously address assigned to Mgmt0/0 interface)
- Use HTTPS as APIC by default uses the most secure protocol to communicate with ASAv



- The correct definition of the Device Interfaces and the Cluster Interfaces is critical for a successful deployment

For the first part, use Table 2 showed in the previous section to properly match the Network Adapter IDs with the ASAv interface IDs that you'd like to use. The Path refers to the physical Port or Port-channel or VPC that enables the way in and out of the Firewall interfaces. In this case, ASA is located in an ESX host, where in and out are same for both interfaces. In a Physical appliance, Inside and Outside of the Firewall (FW) would be different physical ports.

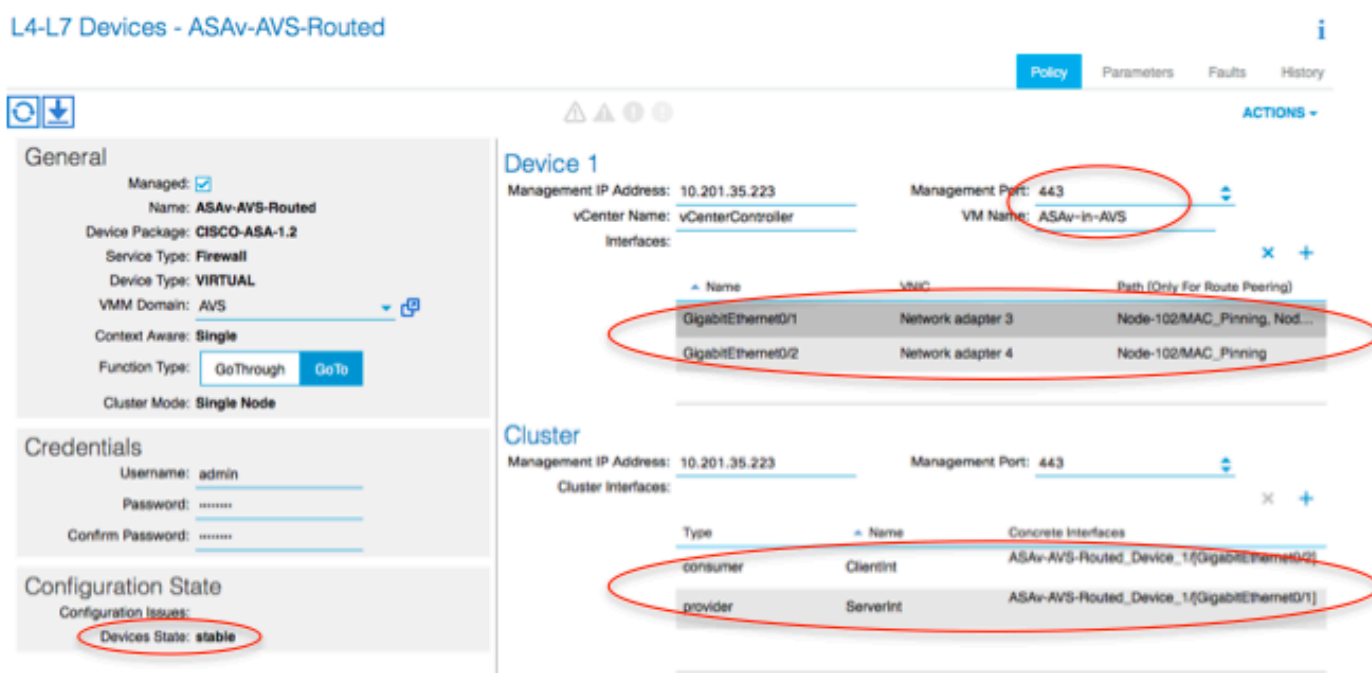
For the second part, the Cluster interfaces have to be defined always with not exceptions (even if Cluster HA is not used), this is because the Object Model has an association between the **mIf** interface (meta interface on the Device Package), the **LIf** interface (leaf interface such as e.g., external, internal, inside, etc.) and the **CIf** (concrete interface). The L4-L7 concrete devices have to be configured in a device cluster configuration and this abstraction is called a logical device. The logical device has logical interfaces that are mapped to concrete interfaces on the concrete device.

For this example, the following association will be used:

Gi0/0 = vmnic2 = ServerInt/provider/server > EPG1

Gi0/1 = vmnic3 = ClientInt/consumer/client > EPG2

L4-L7 Devices - ASA-AVS-Routed

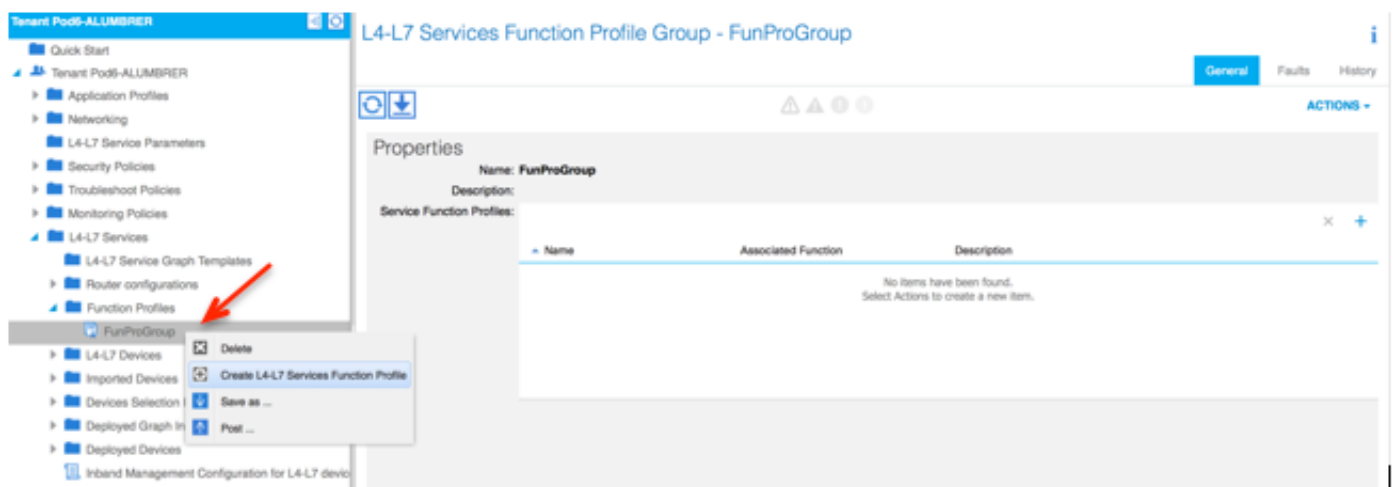
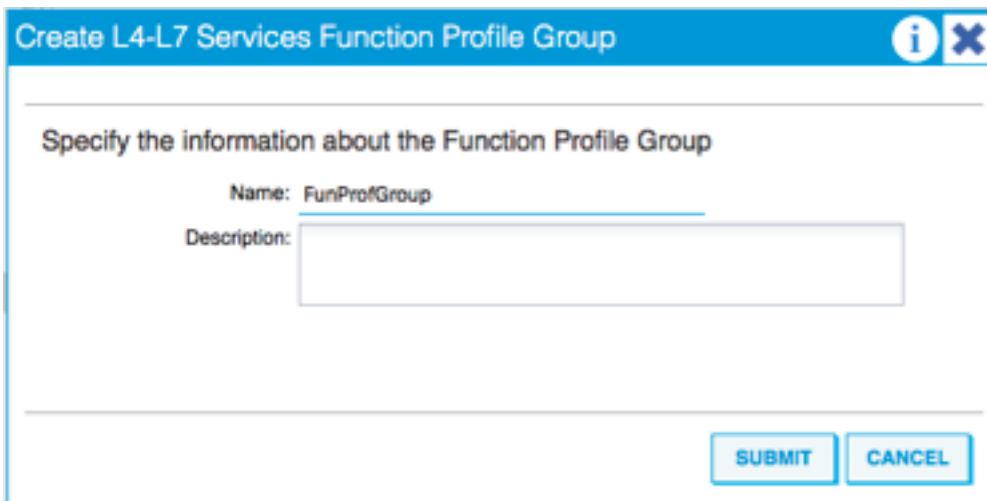


Note: For failover/HA deployments, GigabitEthernet 0/8 is pre-configured as the failover interface.

Device state should be Stable and you should be ready to deploy the Function Profile and Service Graph Template

Service Graph Temple

Firstly, create a Function Profile for ASAv but before that you need to create Function Profile Group and then L4-L7 Services Function Profile under that folder, as shown in the image:



- Select the **WebPolicyForRoutedMode** Profile from the drop down menu and proceed to configure the interfaces on the firewall. From here on, the steps are optional and can be implemented/modified later. These steps can be taken at a few different stages in the deployment depending on how reusable or custom the Service Graph could be.

For this exercise, a routed firewall (GoTo mode) requires that each interface has a unique IP address. Standard ASA configuration also has a interface security level (external interface is less secure, internal interface is more secure). You can also change the name of the interface as per your requirement. Defaults are used in this example.

- Expand Interface Specific Configuration, add IP address and security level for ServerInt with the following format for the IP address **x.x.x.x/y.y.y.y** or **x.x.x.x/yy**. Repeat the process for the ClientInt interface.

Note: You can also modify the default Access-List settings and create your own base

template. By default, the RoutedMode template will include rules for HTTP & HTTPS. For this exercise, SSH and ICMP will be added to the allowed outside access-list.

Create Function Profile

Name: FunProf-ASA
 Description: optional

Copy Existing Profile Parameters:
 Profile: CISCO-ASA-1.2/WebPolicyForRoutedMode

Features and Parameters

In order to auto apply new values to the parameters of existing graph instance when users modify function profiles, the name of top folder must be ended with -Default.

Basic Parameters All Parameters

Folder/Param	Name	Value	Mandatory	Locked	Shared
Destination Service	destination_service				
High Port					
Low Port	low_port	22		false	
Operator	operator	eq		false	
ICMP					
Logging					
Protocol					
Source Address					
Source Service					
Action	action	permit		false	
Order	order	30		false	

SUBMIT CANCEL

- Then click **Submit**
- Now, create the Service Graph Template

Tenant Pod6-ALUMBRER

Quick Start

- Tenant Pod6-ALUMBRER
 - Application Profiles
 - Networking
 - L4-L7 Service Parameters
 - Security Policies
 - Troubleshoot Policies
 - Monitoring Policies
 - L4-L7 Services
 - L4-L7 Service Graph Templates
 - Router configurations

L4-L7 Service Graph Templates

Create L4-L7 Service Graph Template

- Drag and Drop the Device Cluster to the right to form the relationship between Consumer and Provider, select Routed Mode and the previously created Function Profile.

Graph Name: Graph1-akumbler

Graph Type: Create A New One Clone An Existing One

Consumer

Provider

ASA+AVS-...
ASA+V

Please drag a device from devices table and drop it here to create a service node.

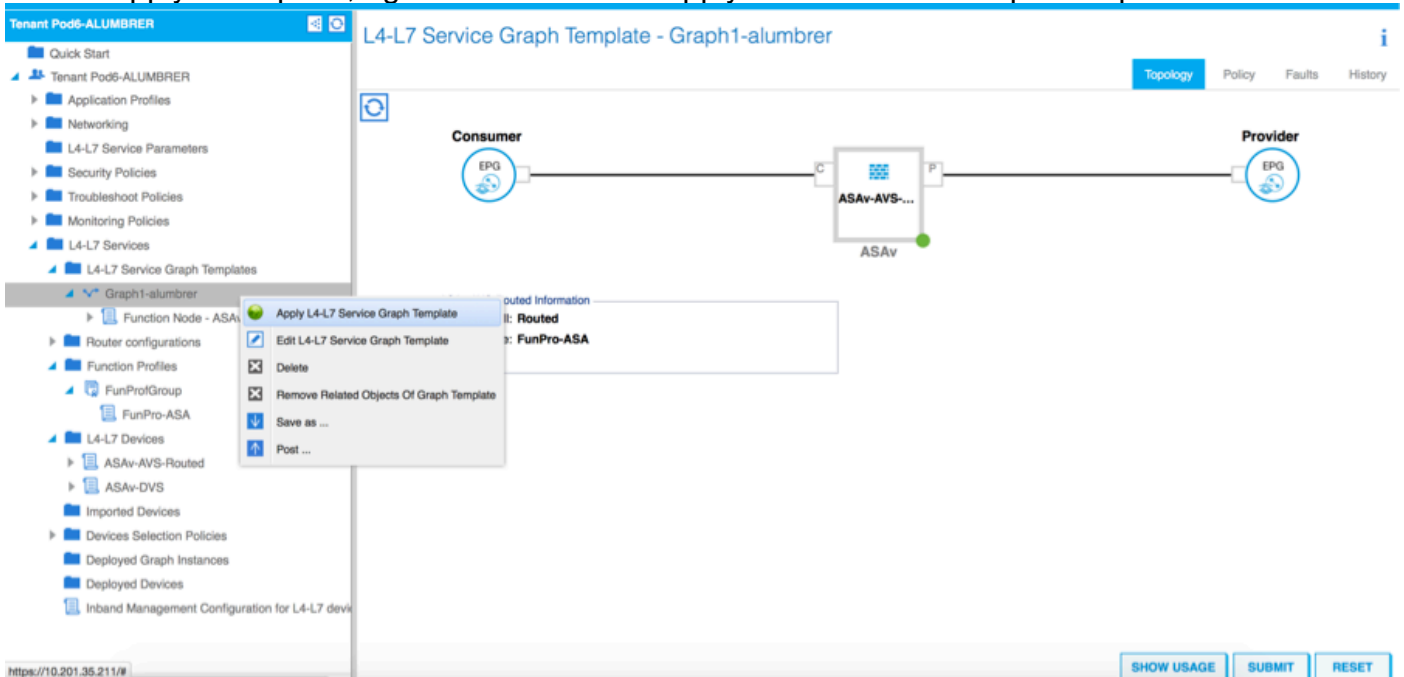
ASA+AVS-Routed Information

Firewall: Routed Transparent

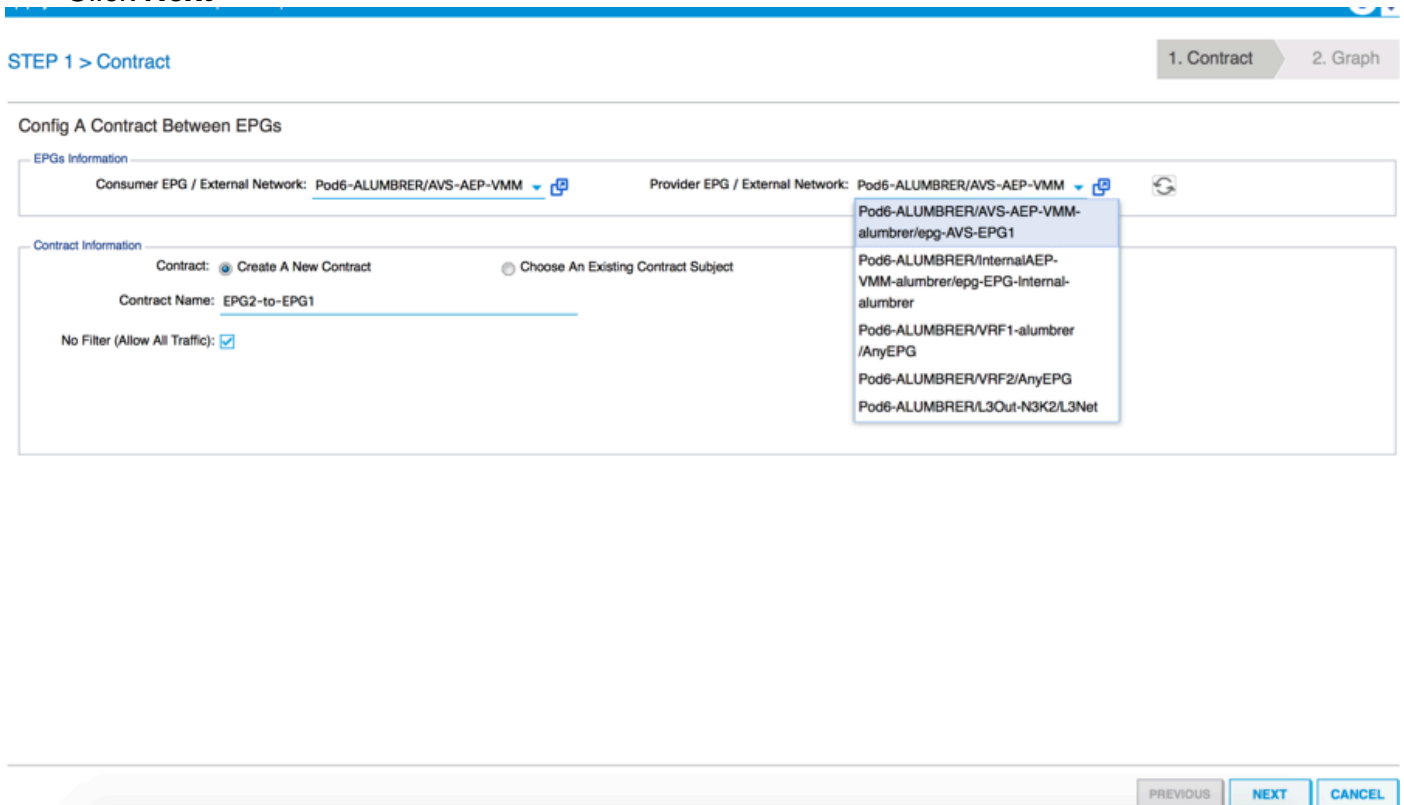
Profile: Pod6-ALUMBRER/FunProfGroup/FunPro

SUBMIT CANCEL

- Check template for faults. The templates are created to be reusable, they must then be applied to particular EPGs etc.
- To apply a template, right click and select Apply L4-L7 Service Graph Template




- Define which EPG will be on the Consumer side and Provider side. In this exercise, AVS-EPG2 is the Consumer (Client) and AVS-EPG1 is the Provider (server). Remember that no Filter is applied, this will allow the firewall to do all the filtering based on the access-list defined in the last section of this wizard.
- Click **Next**



- Verify the BD information for each of the EPGs. In this case, EPG1 is the Provider on the IntBD DB and EPG2 is the Consumer on BD ExtBD. EPG1 will connect on firewall interface

ServerInt and EPG2 will be connected on interface ClientInt. Both FW interfaces will become the DG for each of the EPGs so traffic is forced to cross the firewall at all times.

- Click **Next**



ASAv-AVS-Routed Information

Firewall: routed
Profile: FunPro-ASA

Consumer Connector

Type: General Route Peering

BD: Pod6-ALUMBRER/ExtBD-alumbrer

Cluster Interface: ClientInt

Provider Connector

Type: General Route Peering

BD: Pod6-ALUMBRER/IntBD-alumbrer

Cluster Interface: ServerInt

PREVIOUS NEXT CANCEL

- In the Config Parameters section, click on **All Parameters** and verify if there are RED indicators that need to be updated/configured. In the output as shown in the image, it can be noticed that the order on the access-list is missed. This is equivalent to the line order you'll see in a show ip access-list X.

STEP 3 > ASAv-AVS-Routed Parameters

1. Contract 2. Graph 3. ASAv-AVS-Routed Parameters

config parameters for the selected device

Profile Name: FunPro-ASA

Features: Interfaces, Access Lists, NAT, TrafficSelectorObjects, All

Required Parameters All Parameters

Folder/Param	Name	Value	Write Domain
Access List	access-list-inbound		
Access Control Entry	ICMP		
Access Control Entry	SSH-2		
Access Control Entry	SSH		
Destination Address			
Destination Service	destination_service		
ICMP			
Logging			
Protocol	protocol		
Source Address			
Source Service			
Action	action	permit	
Order	order	30	select asa domain
Access Control Entry			
Access Control Entry			

UPDATE RESET CANCEL

RED indicators parameters needed to be updated and GREEN indicates parameters will be submitted to the provider EPG.

PREVIOUS FINISH CANCEL

- You can also verify the IP addressing assigned from the Function Profile defined earlier, here is a good chance to change information if required. Once all parameters are set, click **Finish**, as shown in the image:

STEP 3 > ASA- AVS-Routed Parameters

1. Contract > 2. Graph > 3. ASA- AVS-Routed Parameters

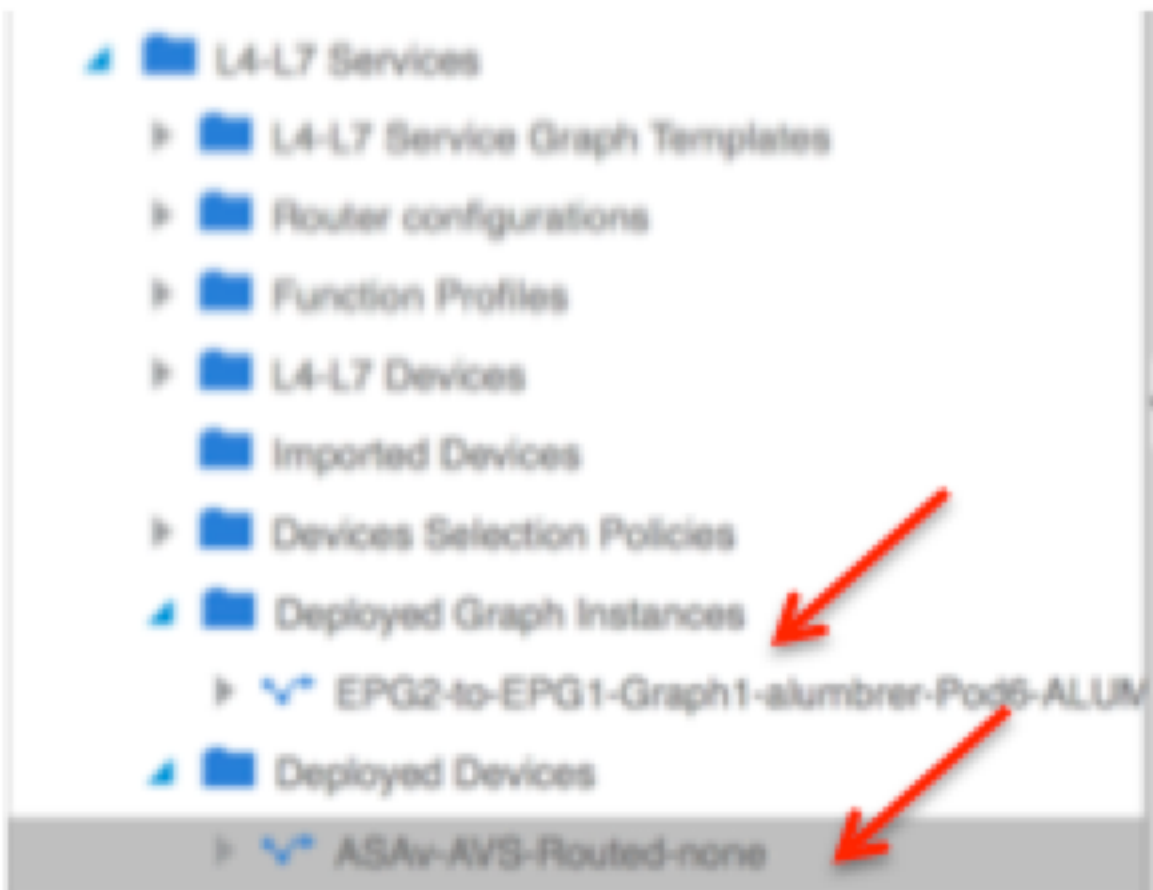
config parameters for the selected device

Profile Name: FunProf-ASA

Folder/Param	Name	Value	Write Domain
Device Config	Device		
Access List	access-list-inbound		
Bridge Group Interface			
Interface Related Configuration	externalIf		
Access Group	ExtAccessGroup	access-list-inbound	
Inbound Access List	name	access-list-inbound	
Outbound Access List			
IPv6 Enforce EUI-64			
Interface Specific Configuration	externalIfCtg		
IPv4 Address Configuration	IPv4Address		
IPv4 Address	ipv4_address	192.168.10.1/24	
IPv4 Standby Address			
IPv6 Address Configuration			
IPv6 Link Local Address Configuration			
IPv6 Router Advertisements			

RED indicators parameters needed to be updated and GREEN indicates parameters will be submitted to the provider EPG.

- If everything goes fine, a new Deployed device and Graph Instance should appear.

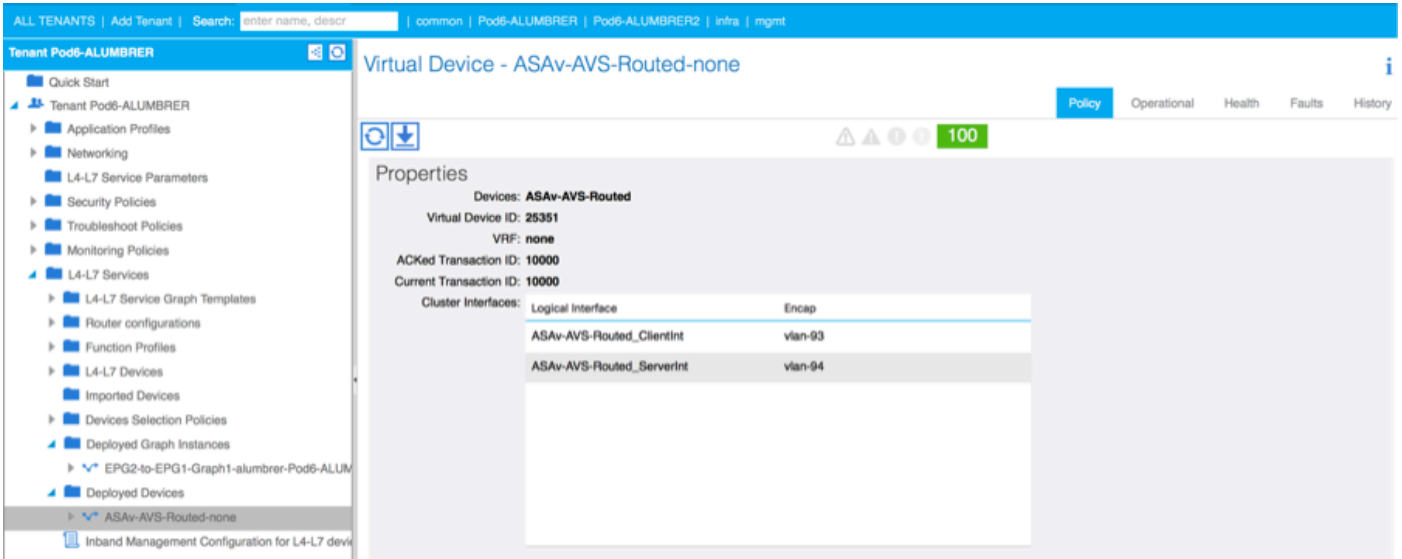


Verify

- One important thing to verify after creating the Service graph is that the Consumer/Provider relationship was created with proper Meta Connector. Verify under the Function Connector Properties.

Note: Each interface of the Firewall will be assigned with an encap-vlan from the AVS

Dynamic Pool. Verify there are no faults.

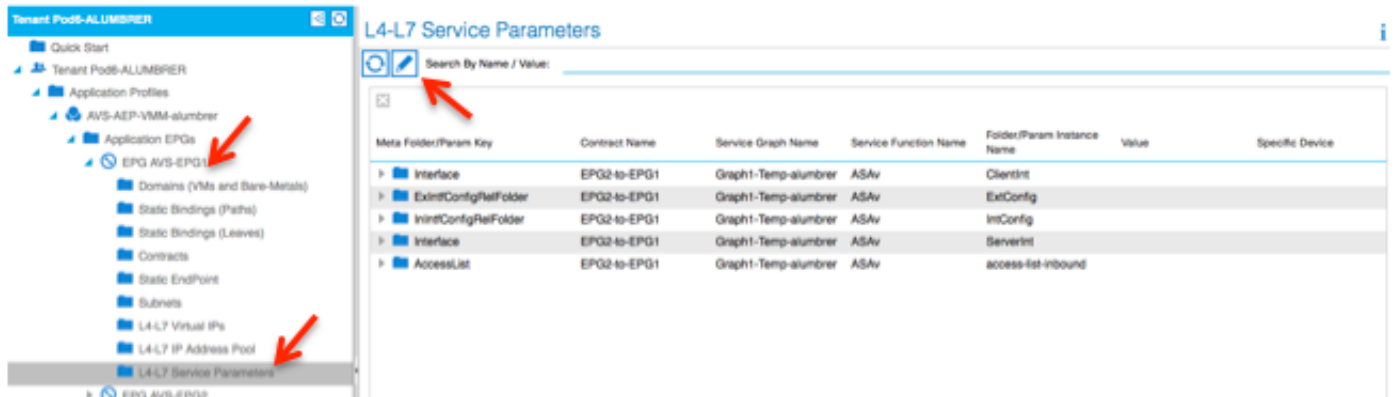


- Now, you can also verify the information that was pushed to the ASAv

```

ASAV-w-AUS# show interface ip brief
Interface                               IP-Address      OK? Method Status Prot
-----                               -
GigabitEthernet0/0                     192.168.10.1   YES manual up    up
GigabitEthernet0/1                     172.16.1.1    YES manual up    up
GigabitEthernet0/2                     unassigned     YES unset  administratively down up
GigabitEthernet0/3                     unassigned     YES unset  administratively down up
GigabitEthernet0/4                     unassigned     YES unset  administratively down up
GigabitEthernet0/5                     unassigned     YES unset  administratively down up
GigabitEthernet0/6                     unassigned     YES unset  administratively down up
GigabitEthernet0/7                     unassigned     YES unset  administratively down up
GigabitEthernet0/8                     unassigned     YES unset  administratively down up
Management0/0                          10.201.35.223 YES CONFIG up    up
ASAV-w-AUS# show run access-list
access-list access-list-inbound extended permit tcp any any eq www
access-list access-list-inbound extended permit tcp any any eq https
access-list access-list-inbound extended permit tcp any any eq ssh
access-list access-list-inbound extended permit icmp any any
ASAV-w-AUS#
    
```

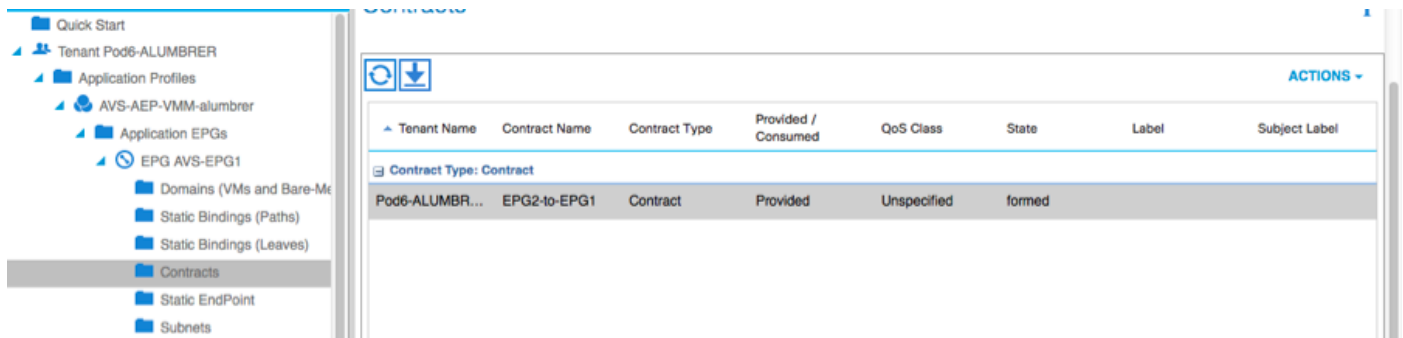
- A new Contract is assigned under the EPGs. From now on, if you need to modify anything on the access-list, the change has to be done from the L4-L7 Service parameters of the Provider EPG.



- On vCenter, you can also verify the Shadow EPGs are assigned to each of the FW interfaces:



For this test, I had the 2 EPGs communicating with standard contracts, these 2 EPGs are in different Domains and different VRFs, so route leaking between them was previously configured. This simplifies a bit after you insert the Service Graph as the FW sets up the routing and filtering in between the 2 EPGs. The DG previously configured under the EPG and BD can now be removed same as the contracts. Only the contract pushed by the L4-L7 should remain under the EPGs.



As the standard contract is removed, you can confirm that traffic is now flows through the ASA, the command show access-list should display the hit count for the rule incrementing every time the client sends a request to the server.

```
ASA-V-W-AUS#
ASA-V-W-AUS# show access-list
access-list cached ACL log flows: total 0, denied 0 (deny-flow-max 4096)
  alert-interval 300
access-list access-list-inbound; 4 elements; name hash: 0xcb5bd6c7
access-list access-list-inbound line 1 extended permit tcp any any eq www (hitcnt=0) 0xc873a747
access-list access-list-inbound line 2 extended permit tcp any any eq https (hitcnt=0) 0x48bedbdd
access-list access-list-inbound line 3 extended permit tcp any any eq ssh (hitcnt=4) 0x532fd57a
access-list access-list-inbound line 4 extended permit icmp any any (hitcnt=4) 0xe4b5a75d
ASA-V-W-AUS#
```

On the leaf, endpoints should be learned for client and server VMs as well as the ASA interfaces

```
leaf2# show endpoint
Legend:
0 - peer-attached   H - vtep           a - locally-aged   S - static
V - vpc-attached   p - peer-aged     L - local          M - span
s - static-arp     B - bounce

+-----+-----+-----+-----+-----+
| VLAN/ | Encap | MAC Address | MAC Info/ | Interface |
| Domain| VLAN  | IP Address  | IP Info   |            |
+-----+-----+-----+-----+-----+
Pod6-ALUMBRER:VRF1-aalumbler          50.50.50.50 L
14/Pod6-ALUMBRER:VRF1-aalumbler      vxlan-14778359 5897.bda4.f9bc L          eth1/13
30                                     vlan-98         0050.5689.1a08 L          eth1/7
Pod6-ALUMBRER:VRF1-aalumbler      Server IP & MAC  vlan-98         192.168.10.10 L
25                                     vlan-94         0050.5689.ca89 L          po4
Pod6-ALUMBRER:VRF1-aalumbler      vlan-94         192.168.10.1 L
mgmt:inb                             192.168.2.11 S
21                                     vlan-97         0050.5689.3fca L          eth1/7
Pod6-ALUMBRER:VRF2      Client IP & MAC  vlan-97         172.16.1.10 L
26                                     vlan-93         0050.5689.e7dd L          po4
Pod6-ALUMBRER:VRF2      vlan-93         172.16.1.1 L
overlay-1                    10.0.104.93 L
overlay-1                    10.0.96.67 L
13                               vxlan-16777209 0050.5677.18a5 H          unspecified
overlay-1                    vxlan-16777209 10.0.32.93 H          unspecified
13                               vxlan-16777209 0050.5660.ddab H          unspecified
overlay-1                    vxlan-16777209 10.0.32.64 H          unspecified
```

see both firewall interfaces attached to the VEM.

ESX-1

```
~ # vemcmd show port vlan
```

LTL	VSM Port	Admin	Link	State	Cause	PC-LTL	SGID	ORG	svcp	Type	Vem Port
22	Eth1/5	UP	UP	FWD	-	1040	4	0	0		vmnic4
23	Eth1/6	UP	UP	FWD	-	1040	5	0	0		vmnic5
50		UP	UP	FWD	-	0	4	0	0		vmk1
51		UP	UP	FWD	-	0	4	0	0		ASAv-in-AVS.eth1
52		UP	UP	FWD	-	0	4	0	0		ASAv-in-AVS.eth2
1040	Po1	UP	UP	FWD	-	0		0	0		

ESX-2

```
~ # vemcmd show port vlan
```

LTL	VSM Port	Admin	Link	State	Cause	PC-LTL	SGID	ORG	svcp	Type	Vem Port
24	Eth1/7	UP	UP	FWD	-	1040	6	0	0		vmnic6
50		UP	UP	FWD	-	0	6	0	0		vmk1
51		UP	UP	FWD	-	0	6	0	0		Client1-AVS.eth0
52		UP	UP	FWD	-	0	6	0	0		Server1-AVS.eth0
1040	Po1	UP	UP	FWD	-	0		0	0		

Finally, the Firewall rules can be verified at the leaf level too if we know the PC Tags for source and destination EPGs:

EPG1

Name	Description	State	Issues	QoS	Encap	PC Tag
AVS-EPG1		applied		Unspecified		17
EPG-Internal-almubrer		applied		Unspecified		32772

EPG2

Name	Description	State	Issues	QoS	Encap	PC Tag
AVS-EPG2		applied		Unspecified		5476

Filter IDs can be matched with the PC tags on the leaf to verify the FW rules.

Note: The EPG PCTags/Sclass never communicate directly. The communication is interrupted or tied together via the shadow EPGs created by the L4-L7 service graph

insertion.

And communication Client to Server works.

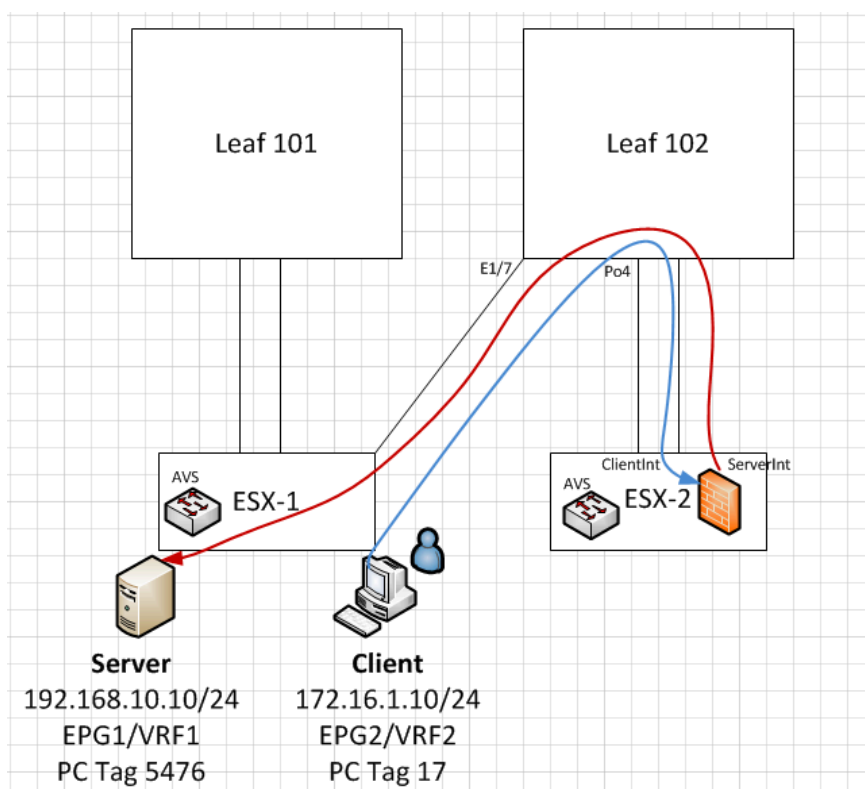
```
cisco@cisco-UbuntuClient:~$ ifconfig
eth1      Link encap:Ethernet  HWaddr 00:50:56:89:3f:ca
          inet addr:172.16.1.10  Bcast:172.16.1.255  Mask:255.255.255.0
          inet6 addr: fe80::250:56ff:fe89:3fca/64  Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:346596  errors:0  dropped:97  overruns:0  frame:0
          TX packets:533034  errors:0  dropped:0  overruns:0  carrier:0
          collisions:0  txqueuelen:1000
          RX bytes:33670388 (33.6 MB)  TX bytes:42734068 (42.7 MB)

lo        Link encap:Local Loopback
          inet addr:127.0.0.1  Mask:255.0.0.0
          inet6 addr: ::1/128  Scope:Host
          UP LOOPBACK RUNNING  MTU:65536  Metric:1
          RX packets:170350  errors:0  dropped:0  overruns:0  frame:0
          TX packets:170350  errors:0  dropped:0  overruns:0  carrier:0
          collisions:0  txqueuelen:0
          RX bytes:18739044 (18.7 MB)  TX bytes:18739044 (18.7 MB)

cisco@cisco-UbuntuClient:~$ ssh 192.168.10.10
cisco@192.168.10.10's password:
Welcome to Ubuntu 14.04 LTS (GNU/Linux 3.13.0-24-generic x86_64)

 * Documentation:  https://help.ubuntu.com/

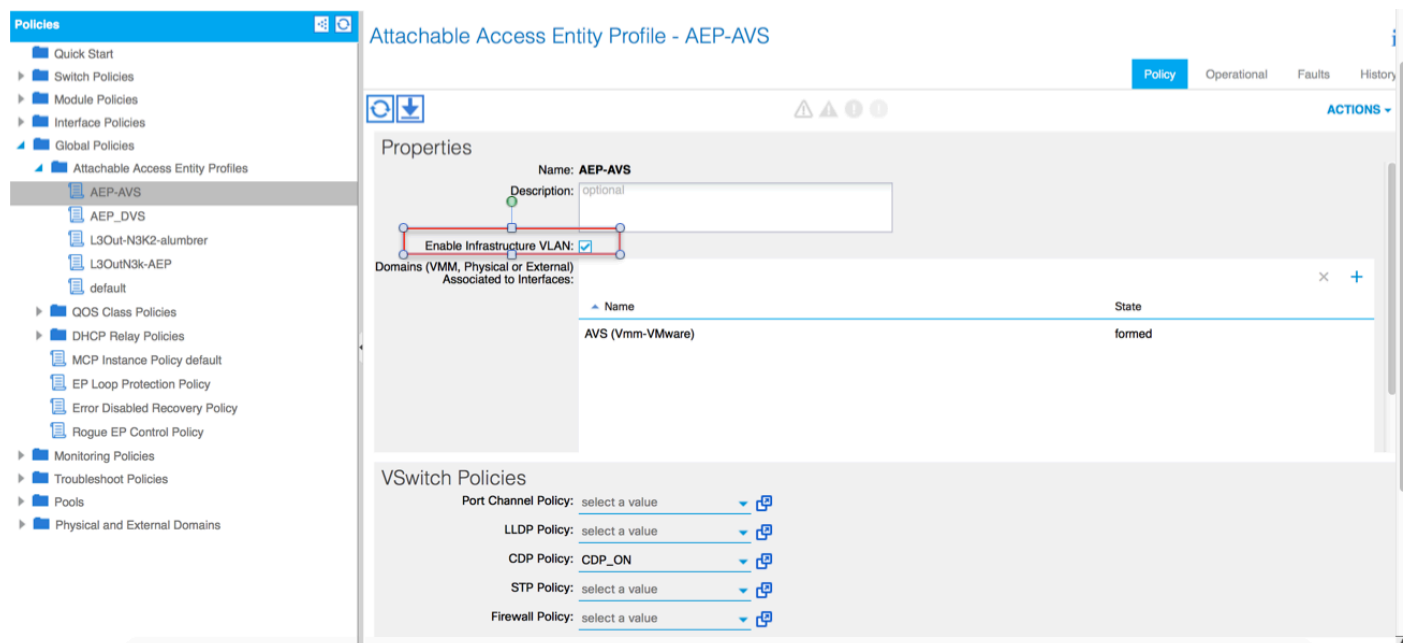
Last login: Mon Feb  1 10:14:11 2016 from 172.16.1.10
cisco@cisco-UbuntuClient:~$ █
```



Troubleshoot

VTEP address is not assigned

Verify that Infrastructure Vlan is checked under the AEP:



Unsupported Version

Verify VEM version is correct and support appropriate ESXi VMWare system.

```
~ # vem version
Running esx version -1746974 x86_64
VEM Version: 5.2.1.3.1.10.0-3.2.1
OpFlex SDK Version: 1.2(1i)
System Version: VMware ESXi 5.5.0 Releasebuild-1746974
ESX Version Update Level: 0
```

VEM and Fabric communication not working

- Check VEM status
vem status

- Try reloading or restating the VEM at the host:
vem reload
vem restart

- Check if there's connectivity towards the Fabric. You can try pinging 10.0.0.30 which is (infra:default) with 10.0.0.30 (shared address, for both Leafs)

```
~ # vmkping -I vmk1 10.0.0.30
PING 10.0.0.30 (10.0.0.30): 56 data bytes
```

```
--- 10.0.0.30 ping statistics ---
3 packets transmitted, 0 packets received, 100% packet loss
```

If ping fails, check:

- Check OpFlex status - The DPA (DataPathAgent) handles all the control traffic between AVS and APIC (talks to the immediate Leaf switch that is connecting to) using OpFlex (opflex client/agent).

```

All EPG communication will go thru this opflex connection. ~ # vemcmd show opflex Status: 0
(Discovering) Channel0: 0 (Discovering), Channel1: 0 (Discovering) Dvs name: comp/prov-
VMware/ctrlr-[AVS]-vCenterController/sw-dvs-129 Remote IP: 10.0.0.30 Port: 8000 Infra vlan: 3967
FTEP IP: 10.0.0.32 Switching Mode: unknown Encap Type: unknown NS GIPO: 0.0.0.0 you can also
check the status of the vmnics at the host level: ~ # esxcfg-vmknic -l Interface Port
Group/DVPort IP Family IP Address Netmask Broadcast MAC Address MTU TSO MSS Enabled Type vmk0
Management Network IPv4 10.201.35.219 255.255.255.0 10.201.35.255 e4:aa:5d:ad:06:3e 1500 65535
true STATIC vmk0 Management Network IPv6 fe80::e6aa:5dff:fead:63e 64 e4:aa:5d:ad:06:3e 1500
65535 true STATIC, PREFERRED vmk1 160 IPv4 10.0.32.65 255.255.0.0 10.0.255.255 00:50:56:6b:ca:25
1500 65535 true STATIC vmk1 160 IPv6 fe80::250:56ff:fe6b:ca25 64 00:50:56:6b:ca:25 1500 65535
true STATIC, PREFERRED ~ # - Also on the host, verify if DHCP requests are sent back and forth:
~ # tcpdump-uw -i vmk1 tcpdump-uw: verbose output suppressed, use -v or -vv for full protocol
decode listening on vmk1, link-type EN10MB (Ethernet), capture size 96 bytes 12:46:08.818776 IP
truncated-ip - 246 bytes missing! 0.0.0.0.bootpc > 255.255.255.255.bootps: BOOTP/DHCP, Request
from 00:50:56:6b:ca:25 (oui Unknown), length 300 12:46:13.002342 IP truncated-ip - 246 bytes
missing! 0.0.0.0.bootpc > 255.255.255.255.bootps: BOOTP/DHCP, Request from 00:50:56:6b:ca:25
(oui Unknown), length 300 12:46:21.002532 IP truncated-ip - 246 bytes missing! 0.0.0.0.bootpc >
255.255.255.255.bootps: BOOTP/DHCP, Request from 00:50:56:6b:ca:25 (oui Unknown), length 300
12:46:30.002753 IP truncated-ip - 246 bytes missing! 0.0.0.0.bootpc > 255.255.255.255.bootps:
BOOTP/DHCP, Request from 00:50:56:6b:ca:25 (oui Unknown), length 300

```

At this point it can be determined that Fabric communication between the ESXi host and the Leaf does not work properly. Some verification commands can be checked at the leaf side to determine root cause.

```
leaf2# show cdp ne
```

```

Capability Codes: R - Router, T - Trans-Bridge, B - Source-Route-Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater,
                  V - VoIP-Phone, D - Remotely-Managed-Device,
                  s - Supports-STP-Dispute

```

Device-ID	Local Intrfce	Hldtme	Capability	Platform	Port ID
AVS:localhost.localdomainmain	Eth1/5	169	S I s	VMware ESXi	vmnic4
AVS:localhost.localdomainmain	Eth1/6	169	S I s	VMware ESXi	vmnic5
N3K-2 (FOC1938R02L)	Eth1/13	166	R S I s	N3K-C3172PQ-1	Eth1/13

```
leaf2# show port-c sum
```

```

Flags: D - Down          P - Up in port-channel (members)
        I - Individual    H - Hot-standby (LACP only)
        s - Suspended     r - Module-removed
        S - Switched      R - Routed
        U - Up (port-channel)
        M - Not in use. Min-links not met
        F - Configuration failed

```

```

-----
Group Port-      Type      Protocol  Member Ports
Channel
-----
5      Po5(SU)     Eth       LACP      Eth1/5(P)  Eth1/6(P)

```

There are 2 Ports used in the ESXi connected via a Po5

```
leaf2# show vlan extended
```

VLAN Name	Status	Ports
13 infra:default	active	Eth1/1, Eth1/20
19 --	active	Eth1/13
22 mgmt:inb	active	Eth1/1

```

26  --                active    Eth1/5, Eth1/6, Po5
27  --                active    Eth1/1
28  ::                active    Eth1/5, Eth1/6, Po5
36  common:pod6_BD    active    Eth1/5, Eth1/6, Po5

```

```

VLAN Type  Vlan-mode  Encap
-----
13  enet  CE        vxlan-16777209, vlan-3967
19  enet  CE        vxlan-14680064, vlan-150
22  enet  CE        vxlan-16383902
26  enet  CE        vxlan-15531929, vlan-200
27  enet  CE        vlan-11
28  enet  CE        vlan-14
36  enet  CE        vxlan-15662984

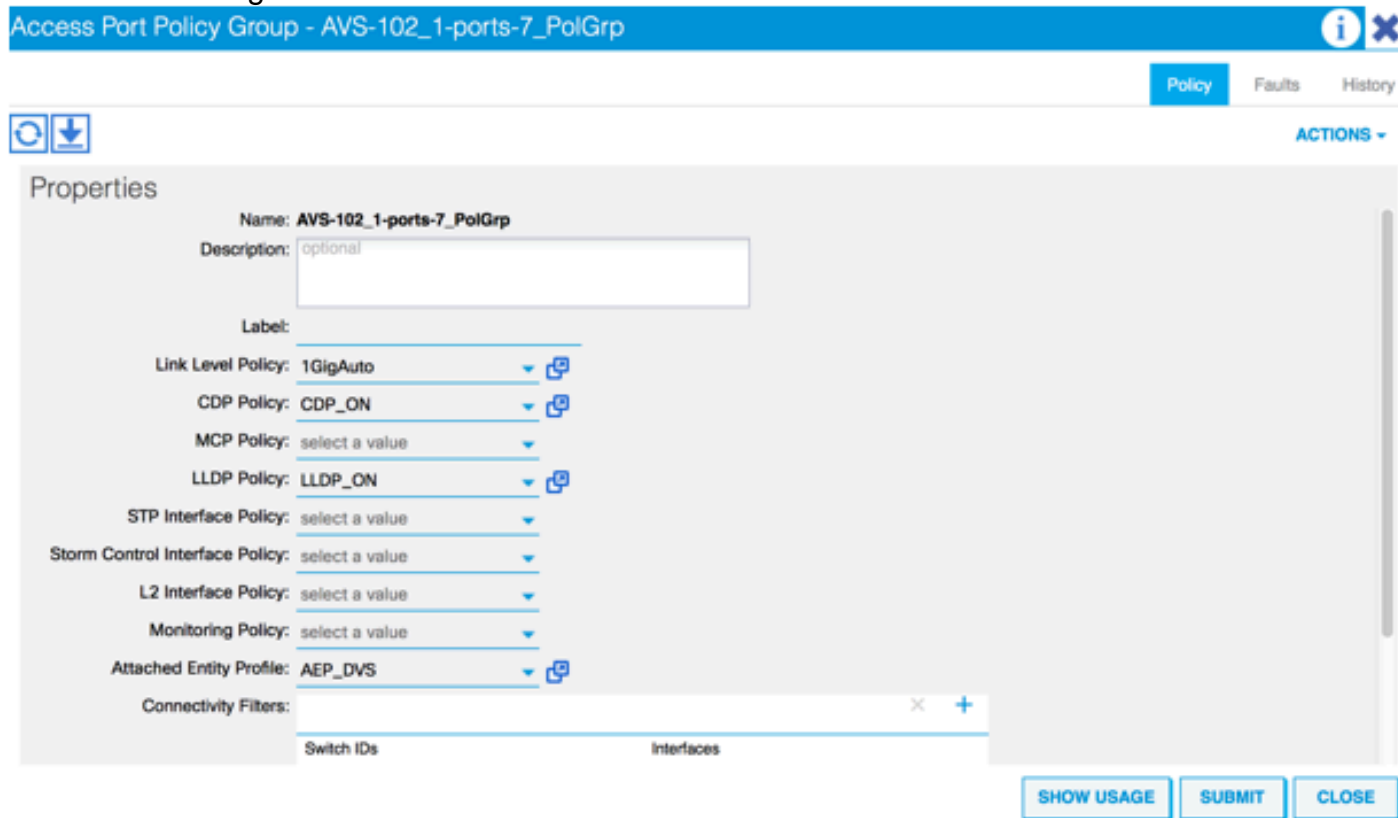
```

From the above output it can be observed that the Infra Vlan is not allowed or passed through the Uplinks ports that go to the ESXi host (1/5-6). This indicates a misconfiguration with the Interface Policy or Switch Policy configured on APIC.

Check both:

Access Policies > Interface Policies > Profiles Access Policies > Switch Policies > Profiles

In this case, the interface profiles are attached to the wrong AEP (old AEP used for DVS), as shown in the image:



After setting of the correct AEP for AVS, we can now see that the Infra Vlan is seen thru the proper Uplinks at the Leaf:

leaf2# show vlan extended

```

VLAN Name                Status    Ports
-----
13  infra:default          active    Eth1/1, Eth1/5, Eth1/6,
                                   Eth1/20, Po5
19  --                    active    Eth1/13
22  mgmt:inb              active    Eth1/1
26  --                    active    Eth1/5, Eth1/6, Po5
27  --                    active    Eth1/1
28  ::                    active    Eth1/5, Eth1/6, Po5
36  common:pod6_BD        active    Eth1/5, Eth1/6, Po5

```

VLAN	Type	Vlan-mode	Encap
13	enet	CE	vxlan-16777209, vlan-3967
19	enet	CE	vxlan-14680064, vlan-150
22	enet	CE	vxlan-16383902
26	enet	CE	vxlan-15531929, vlan-200
27	enet	CE	vlan-11
28	enet	CE	vlan-14
36	enet	CE	vxlan-15662984

and Opflex connection is reestablished after restarting the VEM module:

```

~ # vem restart
stopDpa
VEM SwISCSI PID is
Warn: DPA running host/vim/vimuser/cisco/vem/vemdpa.213997
Warn: DPA running host/vim/vimuser/cisco/vem/vemdpa.213997
watchdog-vemdpa: Terminating watchdog process with PID 213974

~ # vemcmd show opflex
Status: 0 (Discovering)
Channel0: 14 (Connection attempt), Channel1: 0 (Discovering)
Dvs name: comp/prov-VMware/ctrlr-[AVS]-vCenterController/sw-dvs-129
Remote IP: 10.0.0.30 Port: 8000
Infra vlan: 3967
FTEP IP: 10.0.0.32
Switching Mode: unknown
Encap Type: unknown
NS GIPO: 0.0.0.0

~ # vemcmd show opflex
Status: 12 (Active)
Channel0: 12 (Active), Channel1: 0 (Discovering)
Dvs name: comp/prov-VMware/ctrlr-[AVS]-vCenterController/sw-dvs-129
Remote IP: 10.0.0.30 Port: 8000
Infra vlan: 3967
FTEP IP: 10.0.0.32
Switching Mode: LS
Encap Type: unknown
NS GIPO: 0.0.0.0

```

Related Information

Application Virtual Switch Installation

[Cisco Systems, Inc. Cisco Application Virtual Switch Installation Guide, Release 5.2\(1\)SV3\(1.2\)](#)

Deploy the ASAv Using VMware

[Cisco Systems, Inc. Cisco Adaptive Security Virtual Appliance \(ASAv\) Quick Start Guide, 9.4](#)

Cisco ACI and Cisco AVS

[Cisco Systems, Inc. Cisco ACI Virtualization Guide, Release 1.2\(1i\)](#)

Service Graph Design with Cisco Application Centric Infrastructure White Paper

[Service Graph Design with Cisco Application Centric Infrastructure White Paper](#)

[Technical Support & Documentation - Cisco Systems](#)