Setting Up a Cisco ACI Fabric: Initial Deployment Cookbook

Cisco Systems, Inc.

Corporate Headquarters

170 West Tasman Drive

San Jose, CA 95134-1706 USA

http://www.cisco.com

Tel: 408 526-4000 Toll Free: 800 553-NETS (6387)

Fax: 408 526-4100

Contents

Cisco Systems. All printed copies and duplicate soft copies are considered uncontrolled
and the original online version should be referred to for the latest version.

Page 2 of 215

Introduction

This section provides an overview of the goals and prerequisites for this document.

Note: The documentation set for this product strives to use bias-free language. For the purposes of this documentation set, bias-free is defined as language that does not imply discrimination based on age, disability, gender, racial identity, ethnic identity, sexual orientation, socioeconomic status, and intersectionality. Exceptions may be present in the documentation due to language that is hardcoded in the user interfaces of the product software, language used based on RFP documentation, or language that is used by a referenced third-party product.

Goals of this document

This document describes step-by-step Cisco ACI configuration based on common design use cases.

Prerequisites

To best understand the design presented in this document, the reader must have a basic working knowledge of Cisco ACI technology. For more information, see the Cisco ACI white papers available at Cisco.com: https://www.cisco.com/c/en/us/solutions/data-center-virtualization/application-centric-infrastructure/white-paper-listing.html.

Terminology

This document uses the following terms with which you must be familiar:

- VRF: Virtual Routing and Forwarding

- BD: Bridge domain

- EPG: Endpoint group

- Class ID: Tag that identifies an EPG

- Policy: In Cisco ACI, "policy" can mean configuration in general or contract action. In the context of forwarding action, "policy" refers specifically to the Access Control List (ACL)-like Ternary Content-Addressable Memory (TCAM) lookup used to decide whether a packet sourced from one security zone (EPG) and destined for another security zone (EPG) is permitted, redirected, or dropped

Topology

This document covers multiple features up to Cisco ACI Release 4.0(1h). It discusses step-by-step configuration using the topology and IP addressing in Figure 1, Table 1 and Figure 2. The detail will be provided in each section.

Cisco Systems. All printed copies and duplicate soft copies are considered uncontrolled and the original online version should be referred to for the latest version.

Page 5 of 215

Figure 1. Cisco ACI Topology
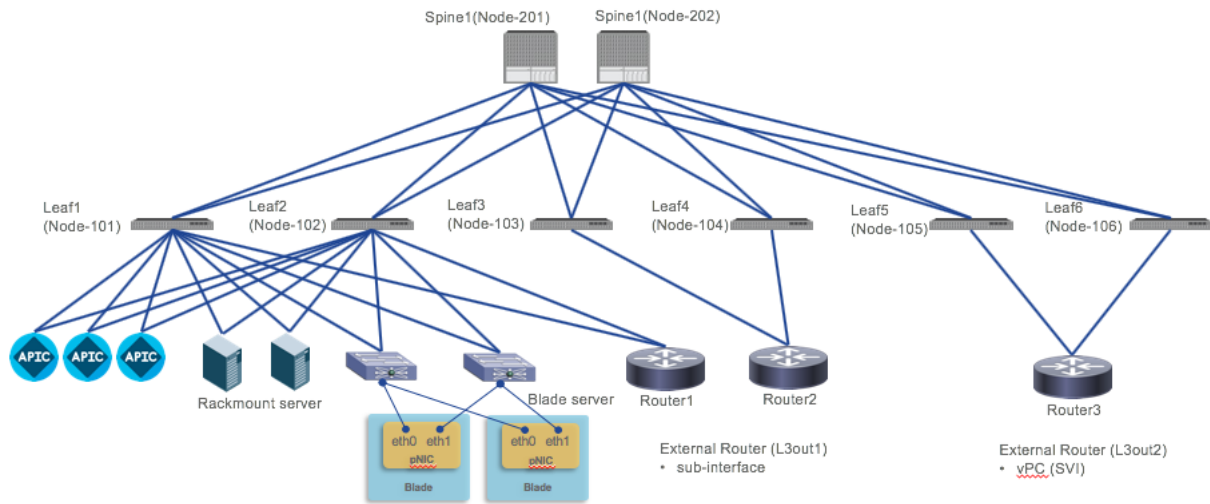
Table 1. Management IP Addressing

Device	Out of band Management IP Address/subnet mask	Out of band Default gateway	Version
APIC1	10.48.22.69/24	10.48.22.100	Release 4.1(1h)
APIC2	10.48.22.70/24	10.48.22.100	Release 4.1(1h)
APIC3	10.48.22.61/24	10.48.22.100	Release 4.1(1h)
vCenter	10.48.22.68	-	-
NTP server	10.48.35.151	-	-
DNS server	10.48.35.150	-	-

Figure 2. Tenant topology
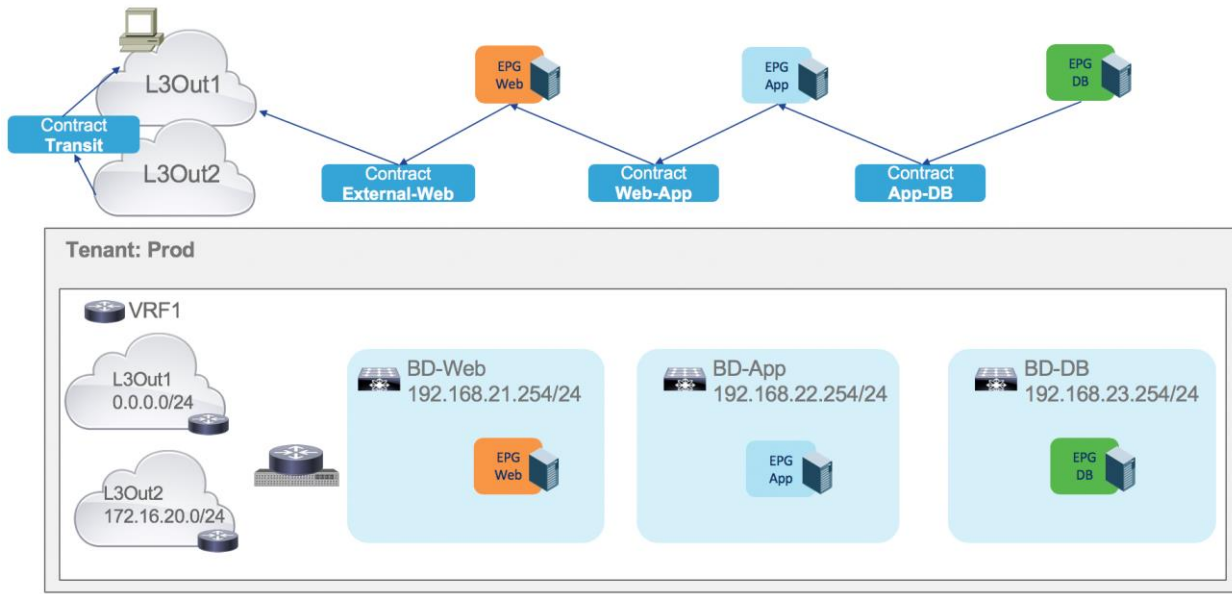
Table 2. Tenant IP Addressing Used in the Example

BD	BD Subnet	EPG	VM-Name	VM IP Address
BD-Web	192.168.21.254/24	EPG-Web	VM-Web	192.168.21.11
BD-App	192.168.22.254/24	EPG-App	VM-App	192.168.22.11
BD-DB	192.168.23.254/24	EPG-DB	VM-DB	192.168.23.11
L3Out1			External-Client1	172.16.10.1
L3Out2			External-Client2	172.16.20.1

Hardware Installation Overview

Setting up your first ACI fabric can be a bit of a challenge if you're new to Cisco server appliances and Nexus 9000 switches. The goal of this chapter is to help you navigate the process of setting up your ACI fabric from scratch. Greater details for the installation process can be found in the *Installation Guides* and *Getting Started Guides* found on Cisco.com for Cisco ACI. To begin, let's review the various components of a Cisco ACI fabric:

■ Cisco Application Policy Infrastructure Controllers (APICs) – One or more, typically three.

■ Nexus 9000 switches running a Cisco ACI software image (spines and leaf Switches)

■ Out-of-Band Management Network connectivity for Cisco APICs and switches

■ Fabric connectivity (Cisco APICs and switches)

■ Cisco APIC integrated management controller (IMC)

The following tasks will be covered

Cisco Systems. All printed copies and duplicate soft copies are considered uncontrolled and the original online version should be referred to for the latest version.

Page 7 of 215

1. Rack and cable the hardware

2. Configure each Cisco APIC's integrated management controller

3. Check APIC firmware and software

4. Check the image type (NX-OS vs. Cisco ACI) and software version of your switches

5. APIC1 initial setup

6. Fabric discovery

7. Set up the remainder of APIC cluster

Detailed Procedure

This section provides detailed configuration procedures.

Rack & Cable Hardware

In this step you configure APIC and switch connectivity.

APIC Connectivity

The first task for setting up your physical fabric devices will be racking & cabling. Depending on your rack space, it might be preferred to spread leaf pairs across racks. Note that all Leaf switches will need to connect to Spine switches, so you will want to locate Leaf & Spine switches in relative proximity to each other to limit cabling requirements. The APICs will be connected to Leaf switches. When using multiple APICs, we recommend connecting APICs to separate Leafs for redundancy purposes.

Figure 3. APIC-Server-M3 Rear Interfaces

From the diagram above at minimum you'll want to connect the following:

■ #2 – 1G/10G LAN connections for the APIC OS Management Interfaces

■ #3 – VGA Video Port. Temporarily needed for initial setup of the IMC to configure remote management (Virtual KVM access)

■ #4- Integrated Management Controller (IMC) for remote platform management. Similar to HP iLO, Dell iDRAC & IBM RAS platforms. This will serve your virtual Keyboard, Video & Mouse (vKMV) services and provide the ability mount virtual ISO (vMedia) images for manual APIC software upgrades.

■ #7- Power Supply Unit 1 & 2. These should be connected to separate power sources (ie. Blue & Red circuits)

■ #9-Fabric Connectivity Ports. These can operate at 10G or 25G speeds (depending on the model of APIC) when connected to Leaf host interfaces. We recommend connecting two fabric uplinks, each to a separate leaf and/or VPC leaf pairs.

If it's APIC M3/L3, VIC 1445 has four ports (port-1, port-2, port-3, and port-4 from left to right). Port-1 and port-2 make a single pair corresponding to eth2-1 on the APIC; port-3 and port-4 make another pair corresponding to eth2-2 on the APIC. Only a single connection is allowed for each pair. For example, you can connect one cable to either port-1 or port-2 and another cable to either port-3 or port-4, but not 2 cables to both ports on the same pair. All ports must be configured for the same speed, either 10G or 25G.

Switch Connectivity

Switch connectivity is pretty straight forward. Leaf switches only connect to Spine switches and vice-versa. This provides your fabric with a fully redundant switching fabric. In addition to the fabric network connections, you'll also connect redundant PSUs to separate power sources, Management Interface to your 1G out-of-band management network, and a console connection to a Terminal server (optional, but highly recommended). You'll want to ensure your switches are mounted correctly such that the fan air flow is in the appropriate direction with your hot/cool aisle of your datacenter. Fans can be ordered with airflow in both directions so ensure you check this prior to ordering. Switches will have two operational port types – Host Interfaces & Fabric Uplinks. Leaf switches connect to Spines using Fabric Uplinks. APICs and other host devices connect to Leaf Host interfaces. Refer to your specific switch model documentation to determine the interface ranges & default port modes.

Configure Each APIC's Integrated Management Controller

When you first connect your APIC's IMC connection marked with "mgmt" on the Rear facing interface, it will be configured for DHCP by default. Cisco recommends that you assign a static address for this purpose to avoid any loss of connectivity or changes to address leases. You can modify the IMC details by connecting up a crash cart (physical monitor, USB keyboard and mouse) to the server and powering it on. During the boot sequence, it will prompt you to press "F8" to configure the IMC. From here you will be presented with a screen similar to below – depending on your firmware version.

Figure 4. Sample IMC configuration Utility

For the "NIC mode" we recommend using *Dedicated* which utilizes the dedicated "mgmt" interface in the rear of the APIC appliance for IMC platform management traffic. Using "Shared LOM" mode which will send your IMC traffic over the LAN on Motherboard (LOM) port along with the APICs OS management traffic. This can

cause issues with fabric discovery if not properly configured and not recommended by Cisco. Aside from the IP address details, the rest of the options can be left alone unless there's a specific reason to modify them. Once a static address has been configured you will need to Save the settings & reboot. After a few minutes you should then be able to reach the IMC Web Interface using the newly assigned IP along with the default IMC credentials of admin and password. Its recommended that you change the IMC default admin password after first use.

Logging into the IMC Web Interface

To log into the IMC, open a web browser to `https://<IMC_IP>`. You'll need to ensure you have flash installed & permitted for the URL. Once you've logged in with the default credentials you'll be able to manage all the IMC features including launching the KVM console.

Figure 5. IMC Web Interface – Launching the KVM Console
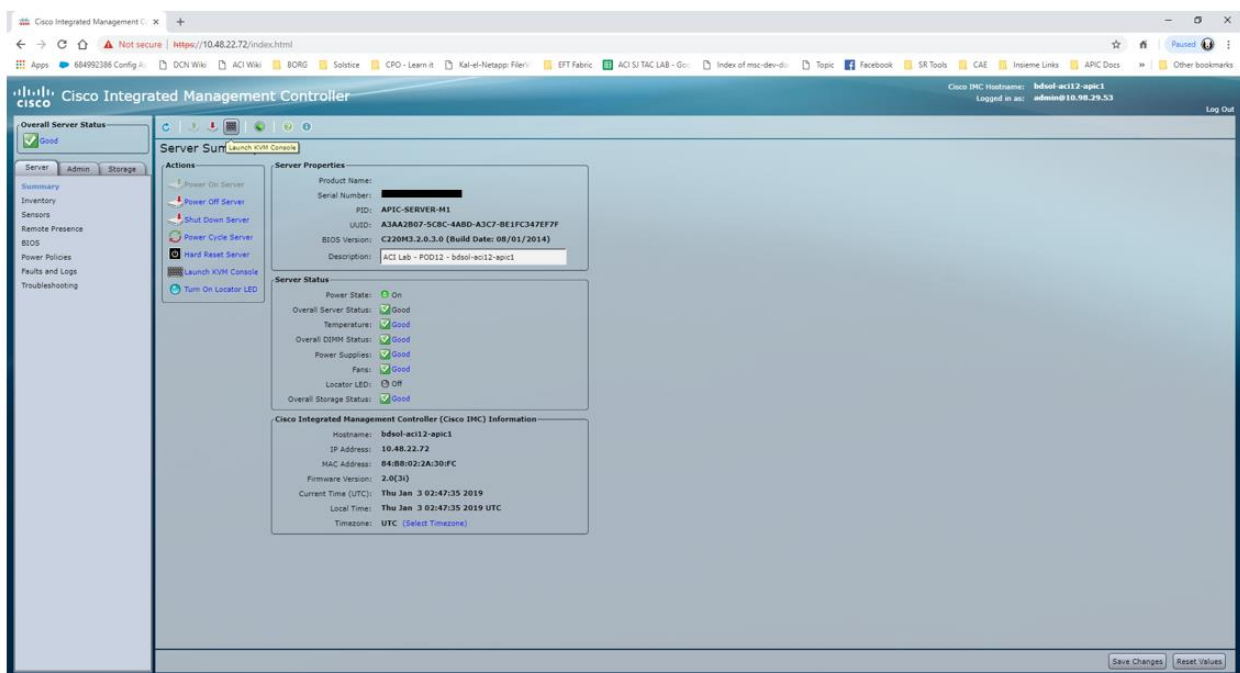
Launching the KVM console will require that you have Java version 1.6 or later installed. Depending on your client security settings, you may need to whitelist the IMC address within your local Java settings in order for the KVM applet to load. Open the KVM console and you should be at the Setup Dialog for the APIC assuming the server is powered on. If not powered up, you can do so from the IMC Web utility.

Note: The IMC will remain accessible assuming the appliance has at least one power source connected. This allows independent power up/down of the APIC appliance operating system.

Check APIC Firmware & Software

About ACI Software Versions: It's important to know which version of ACI software you wish to deploy. ACI software releases include the concept of "Long Lived Release (LLR)" versions which will be maintained and patches longer than non LLR versions. As of this writing, the current LLRs are versions 2.2 and 3.2. Unless you have specific feature requirements available in later releases, you may be best suited deploying the latest LLR as it will have a larger installation base and reduce the possibility of having to perform major software version upgrades. Typically, a new LLR is added for each Major version (ie. 2.x, 3.x etc). Equally important to note is that all your APICs require to run the same version when joining a cluster. This may require manually upgrading/downgrading your APICs manually prior to joining them to the fabric. Instructions on upgrading

standalone APICs using KVM vMedia can be found in the "*Cisco APIC Management, Installation, Upgrade, and Downgrade Guide*" for your respective version.

Switch nodes can be running any version of ACI switch image and can be upgraded/downgraded once joined to the fabric via firmware policy.

Check Image Type (NXOS vs. ACI) & Software Version of Your Switches

For a Nexus 9000 series switch to be added to an ACI fabric, it needs to be running an ACI image. Switches that are ordered as "ACI Switches" will be typically be shipped with an ACI image. If you have existing standalone Nexus 900 switches running traditional NXOS, then you may need to install the appropriate image (For example, `aci-n9000-dk9.14.0.1h.bin`). For detailed instructions on converting a standalone NXOS switch to ACI mode, please see the "*Cisco Nexus 9000 Series NX-OS Software Upgrade and Downgrade Guide*" on CCO for your respective version of NXOS.

Sample Nexus 9000 running NXOS mode, `with n9000-dk9.7.0.3.I1.1a.bin` being the NXOS image:

```
Switch1# show version | grep image

  NXOS image file is: bootflash:///n9000-dk9.7.0.3.I1.1a.bin
```
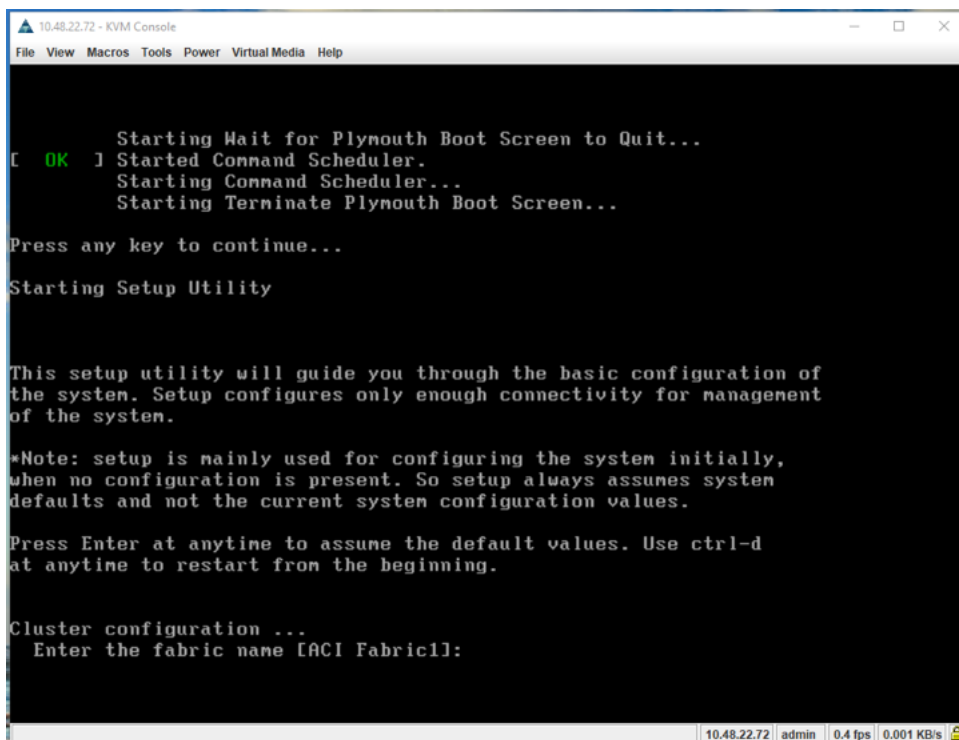
Sample Nexus 9000 running ACI mode, with `aci-n9000-dk9.14.0.1h.bin` being the NXOS image:

```
Switch2# show version | grep image

  kickstart image file is: /bootflash/aci-n9000-dk9.14.0.1h.bin

  system image file is:    /bootflash/auto-s
```

APIC1 Initial Setup

Now that you have basic remote connectivity, you can complete the setup of your ACI fabric from any workstation with network access the APIC. If the server is not powered on, do so now from the IMC interface. The APIC will take 3-4 mins to fully boot. Next thing we'll do is open up a console session via the IMC KVM console using the procedure detailed previously. Assuming the APIC has completed the boot process it should sitting at a prompt "*Press any key to continue…*". Doing so will begin the setup utility.

Figure 6. Starting the Setup Utility on APIC

From here, the APIC will guide you through the initial setup dialogue. Carefully answer each question. Some of the items configured can't be change after initial setup, so review your configuration before submitting it.

Figure 7. APIC Setup Utility using vKVM

Some of the fields are self-explanatory. Select fields are highlighted below for explanation:

■ Fabric Name: User defined, will be the logical friendly name of your fabric.

■ Fabric ID: Leave this ID as the default 1.

■ # of Controllers in fabric: Set this to the # of APICs you plan to configure. This can be increased/decreased later.

■ Pod ID: The Pod ID to which this APIC is connected to. If this is your first APIC or you don't have more than a single Pod installed, this will be always be 1. If you are located additional APICs across multiple Pods, you'll want to assign the appropriate Pod ID where it's connected.

■ Standby Controller: Beyond your active controllers (typically 3) you can designate additional APICs as standby. In the event you have an APIC failure, you can promote a standby to assume the identity of the failed APIC.

■ APIC-X: A special-use APIC model use for telemetry and other heavy ACI App purposes. For your initial setup this typically would not be applicable. *Note: In future release this feature may be referenced as "ACI Services Engine".*

■ TEP Pool: This will be a subnet of addresses used for Internal fabric communication. This subnet will NOT be exposed to your legacy network unless you're deploying the Cisco AVS or Cisco ACI Virtual Edge. Regardless, our recommendation is to assign an unused subnet of size between and /16 and /21 subnet. The size of the subnet used will impact the scale of your Pod. Most customer allocate an unused /16 and
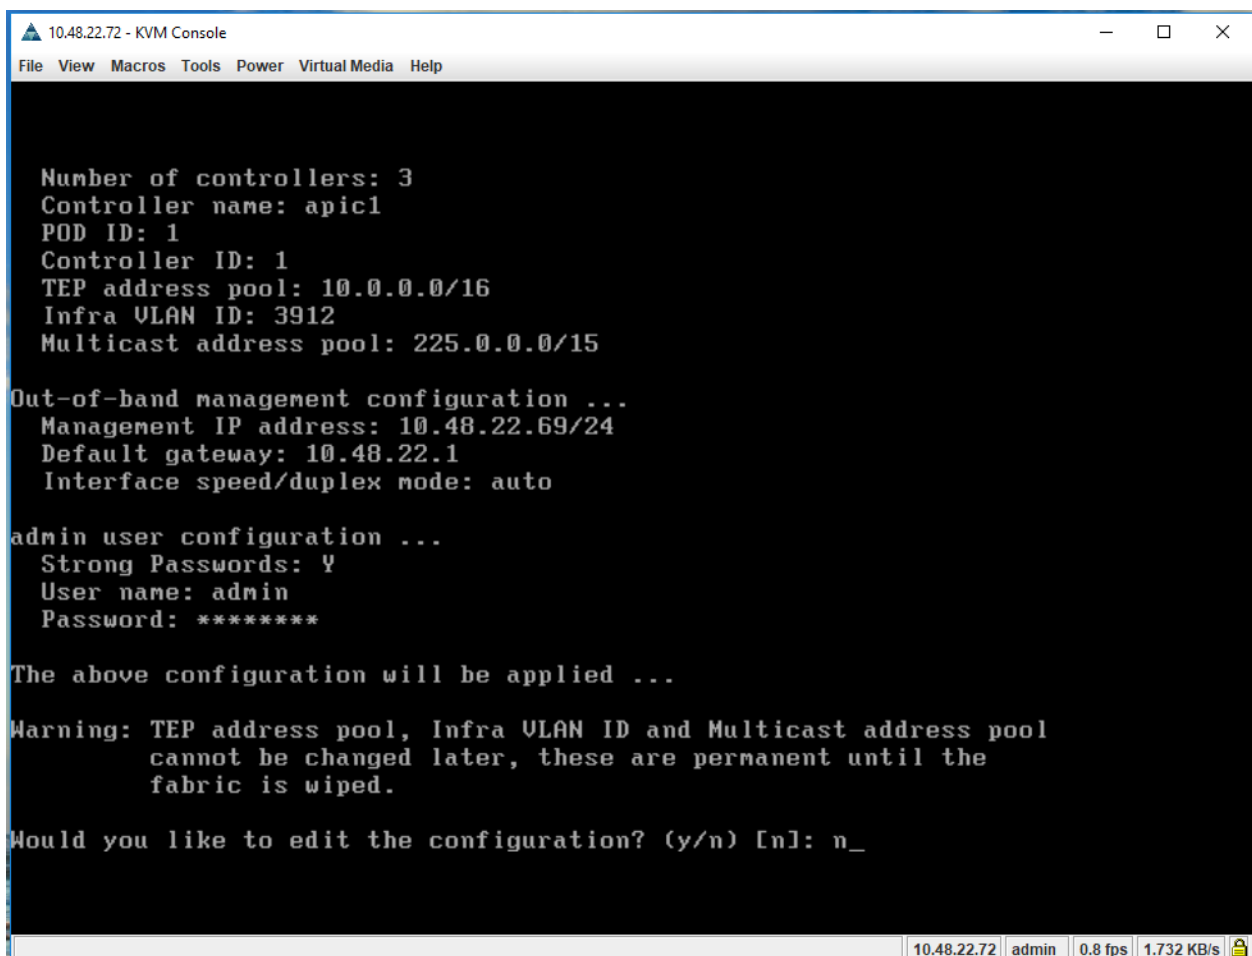
move on. This value can NOT be changed once configured. Having to modify this value requires a wipe of the fabric.

The 172.17.0.0/16 subnet is not supported for the infra TEP pool due to a conflict of address space with the docker0 interface. If you must use the 172.17.0.0/16 subnet for the infra TEP pool, you must manually configure the docker0 IP address to be in a different address space in each Cisco APIC before you attempt to put the Cisco APICs in a cluster.

- Infra VLAN: This is another important item. This is the VLAN ID for all fabric connectivity. This VLAN ID should be allocated solely to ACI, and not used by any other legacy device in your network. Though this VLAN is used for fabric communication, there are certain instances where this VLAN ID may need to be extended outside of the fabric such as the deployment of the Cisco AVS/AVE.   Due to this, we also recommend you ensure the Infra VLAN ID selected does not overlap with any "reserved" VLANs found on your networks. Cisco recommends a VLAN smaller than VLAN 3915 as being a safe option as it is not a reserved VLAN on Cisco DC platforms as of today. This value can NOT be changed once configured. Having to modify this value requires a wipe of the fabric.

- BD Multicast Pool (GIPO): Used for internal connectivity. We recommend leaving this as the default or assigning a unique range not used elsewhere in your infrastructure. This value can NOT be changed once configured. Having to modify this value requires a wipe of the fabric.

Once the Setup Dialogue has been completed, it will allow you to review your entries before submitting. If you need to make any changes enter "y" otherwise enter "n" to apply the configuration. After applying the configuration allow the APIC 4-5 mins to fully bring all services online and initialize the REST login services before attempting to login though a web browser.

Figure 8. APIC Initial Setup - Reviewing Configuration

Fabric Discovery

With our first APIC fully configured, now we will login to the GUI and complete the discovery process for our switch nodes.

1. When logging in for the first time, you may have to accept the Cert warnings and/or add your APIC to the exception list.

Figure 9. Accepting Cert Warnings for APIC Login
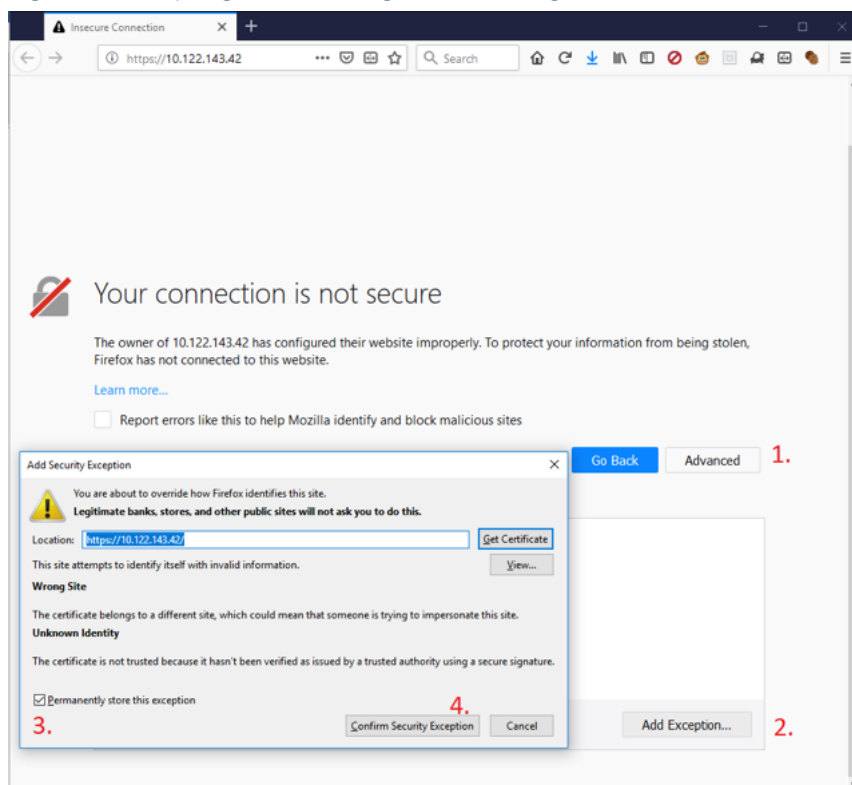
2. Go ahead and login with the admin account and password you assigned during the setup procedure.

On first login you will presented with a "What's New" Window which highlight some of the new features and videos included with this version of APIC. You can optionally click "Do not show me this again" if you wish to prevent this popping up at each login. If you wanted to re-enable this pop up you can do so from the Settings menu for your user account which will be covered next.

Figure 10. APIC Login Screen
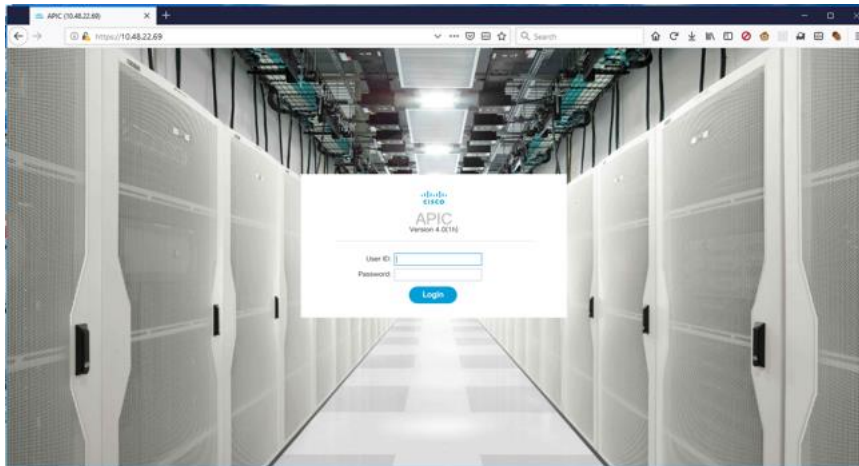
3. Before we get too far into configuration there's a few settings that will make navigation and using the APIC UI easier. You can access UI settings from the top right menu. These settings will be maintained as long as the cookie for this URL managed by your browser exists.

Figure 11. Browser GUI User Settings
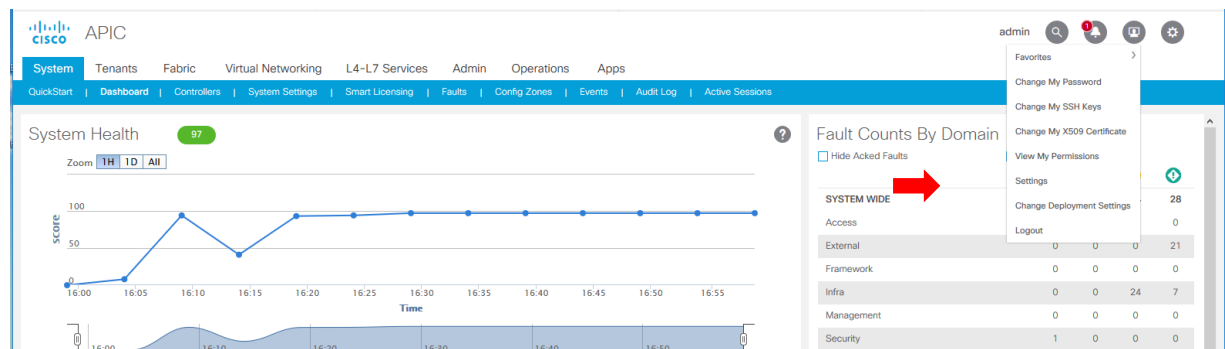
4. Some of the options you may want to enable are:

- Remember Tree Selection: Maintains the folder expansion & location when navigating back & forth between tabs in the APIC UI

- Preserve Tree Divider Position: Makes navigation pane changes persistent

- Default Page Size for Tables: Default is 10 items, but you can increase this to avoid having to click "Next page" on tables.

Figure 12. GUI Application Settings
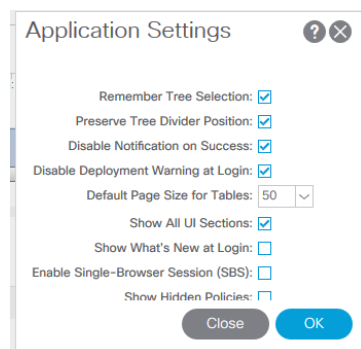
5. Now we'll proceed with the fabric discovery procedure. We'll need to navigate to *Fabric tab > Inventory sub-tab > Fabric Membership folder*.

    From this view you are presented with a view of your registered fabric nodes. Click on the *Nodes Pending Registration* tab in the work pane and we should see our first Leaf switch waiting discovery. Note this would be one of the Leaf switches where the APIC is directly connected to.

    To register our first node, click on the first row, then from the Actions menu (Tool Icon) select Register.

Figure 13. First Node Pending Discovery
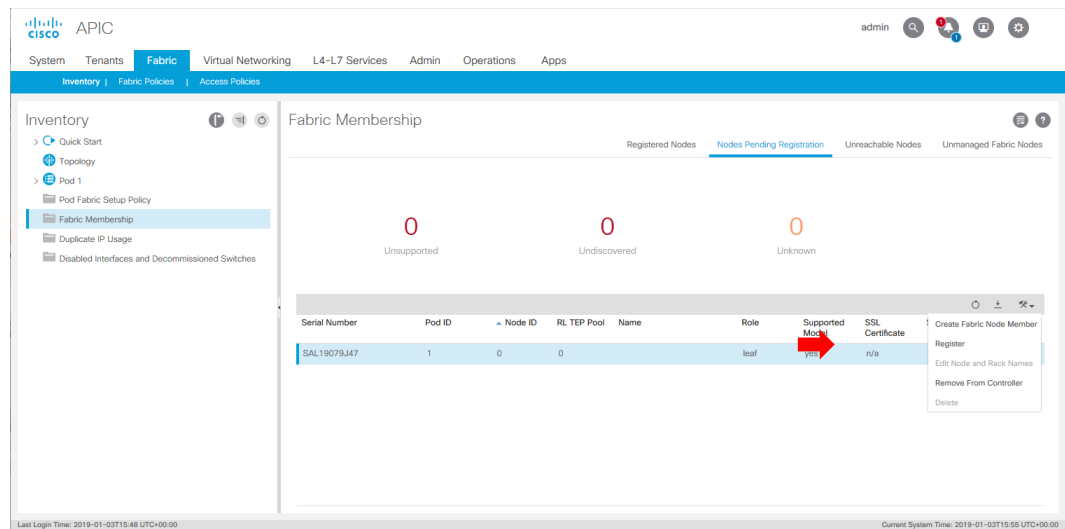
6. The Register wizard will pop up and require some details to be entered including the Node ID you wish to assign, and the Node Name (hostname).

    Hostnames can be modified, but the Node ID will remain assigned until the switch is decommissioned and remove from the APIC. This information is provided to the APIC via LLDP TLVs. If a switch was previously registered to another fabric without being erase, it would never appear as an unregistered node. It's important that all switches have been wiped clean prior to discovery.   It's a common practice for Leaf switches to be assigned Node IDs from 100+, and Spine switches to be assigned IDs from 200+. To accommodate your own numbering convention or larger fabrics you can implement your own scheme. RL TEP Pool is reserved for Remote Leafs usage only and doesn't apply to local fabric-connected Leaf switches. Rack Name is an optional field.

Figure 14. Registration Popup
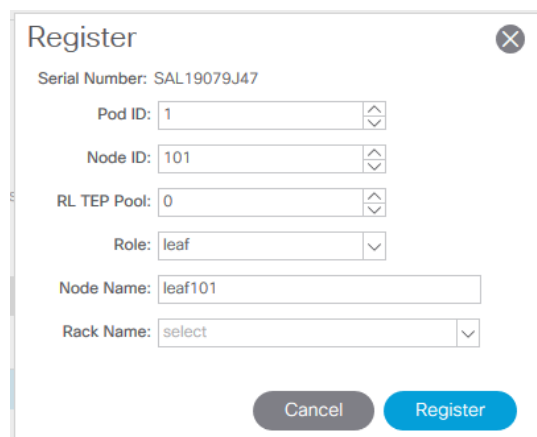
7. Once the registration details have been submitted, the entry for this leaf node will move from the *Nodes Pending Registration* tab to the *Registered Nodes* tab under Fabric Membership. The node will take 3 to 4 minutes to complete the discovery, which includes the bootstrap process and bringing the switch to an "Active" state. During the process, you will notice a tunnel endpoint (TEP) address gets assigned. This will be pulled from the available addresses in your Infra TEP pool (such as 10.0.0.0/16).

Figure 15. Registered Node Discovering
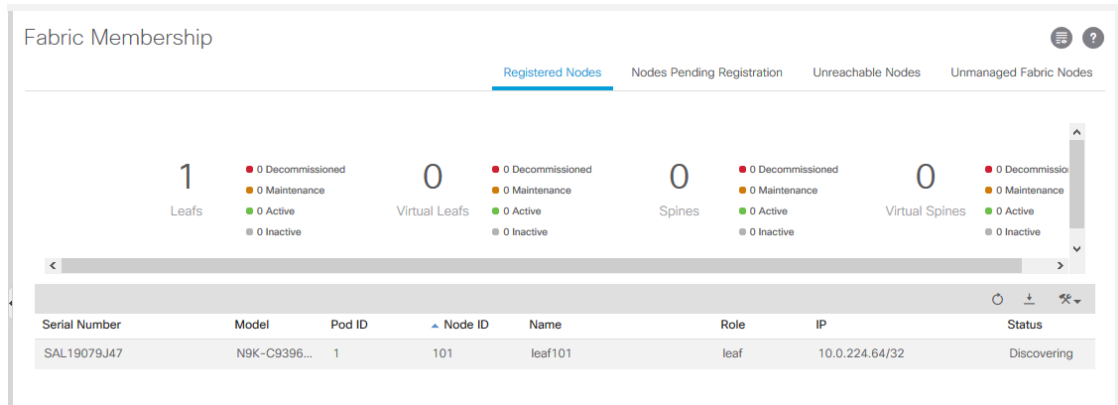
Figure 16. Node Registration Complete
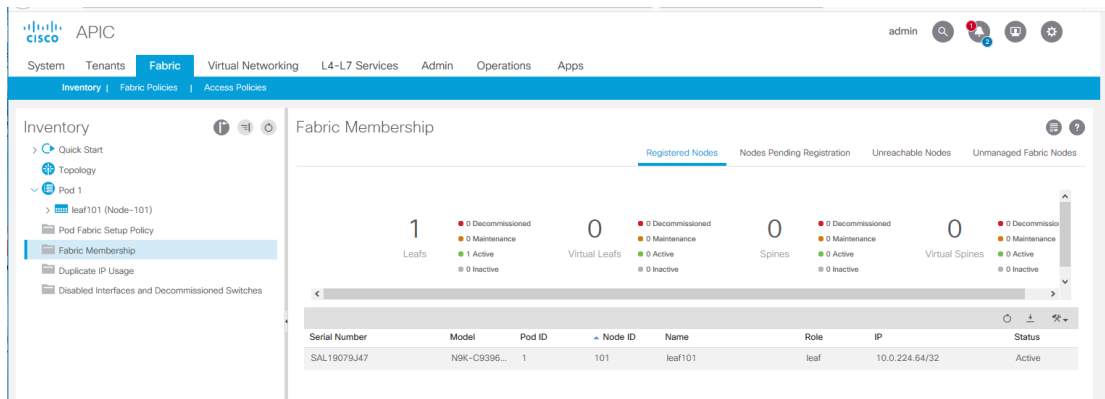
8. After the first Leaf has been discovered and move to an Active state, it will then discovery every Spine switch it's connected to. Go ahead and register each Spine switch in the same manner.

Figure 17. Spines Pending Registration
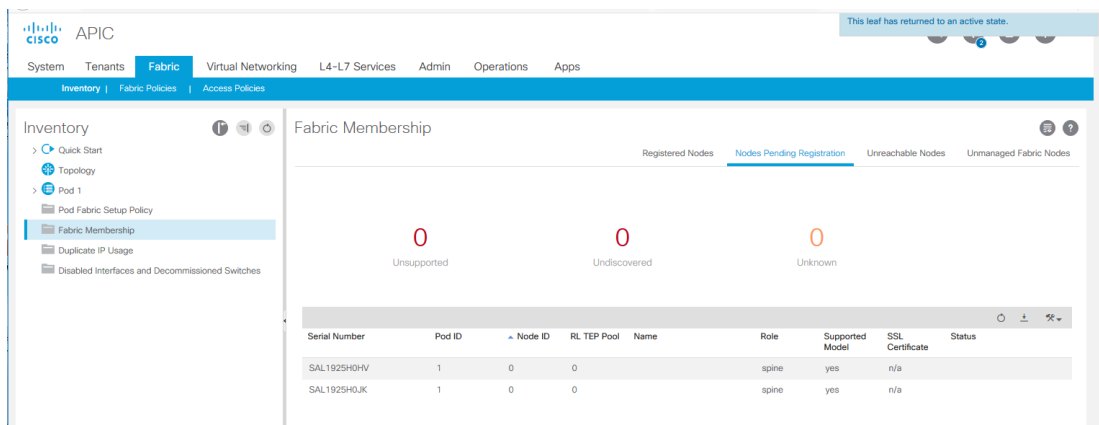
9. Since each Leaf Switch connect to every Spine switch, once the first Spine completes the discovery process, you should see all remaining Leaf switch pending registration. Go ahead with Registering all remaining nodes and wait for all switches to transition to an Active state.

Figure 18. Fabric Discovery Completed
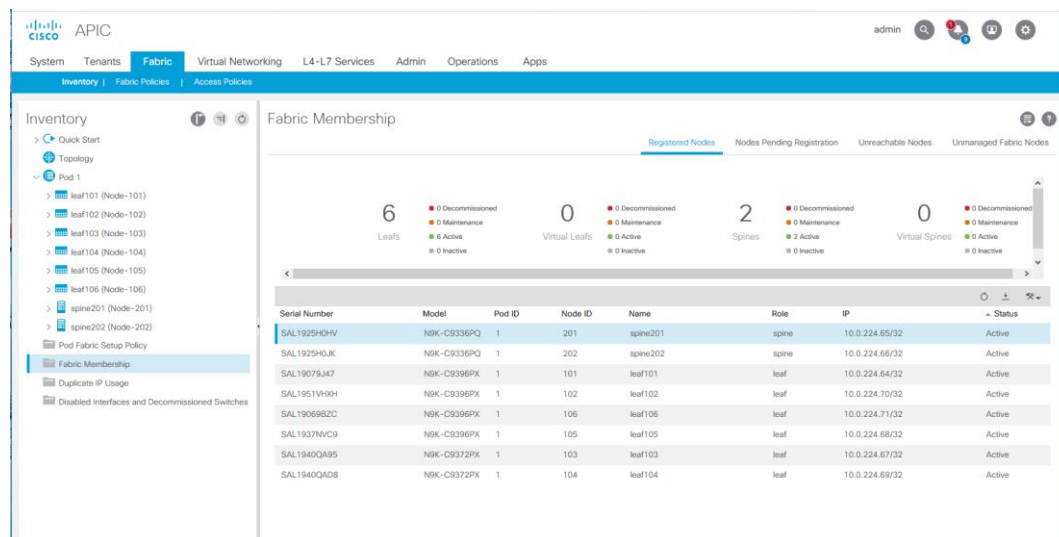
10. With all the switches online & active, our next step is to finish the APIC cluster configuration for the remaining nodes. Navigate to System > Controllers sub menu > Controllers Folder > apic1 > Clusters as Seen by this Node folder.

    From here you will see your single APIC along with other important details such as the *Target Cluster Size* and *Current Cluster Size*. Assuming you configured apic1 with a cluster size of 3, we'll have two more APICs to setup.

Figure 19. Viewing APIC Cluster Information
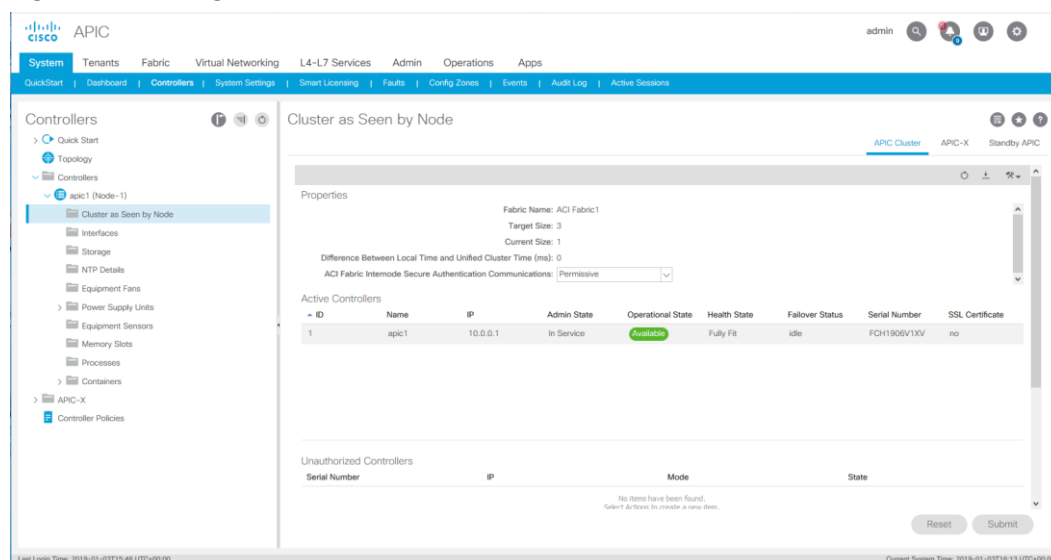
Setup Remainder of APIC Cluster

At this point we would want to now open the KVM console for APIC2 and begin running through the setup Dialogue just as we did for APIC1 previously. When joining additional APICs to an existing cluster it's imperative that you configure the same Fabric Name, Infra VLAN and TEP Pool. The controller ID should be set

to ID 2. You'll notice that you will not be prompted to configure Admin credentials. This is expected as they will be inherited from APIC1 once you join the cluster.

Figure 20. Configuring APIC2
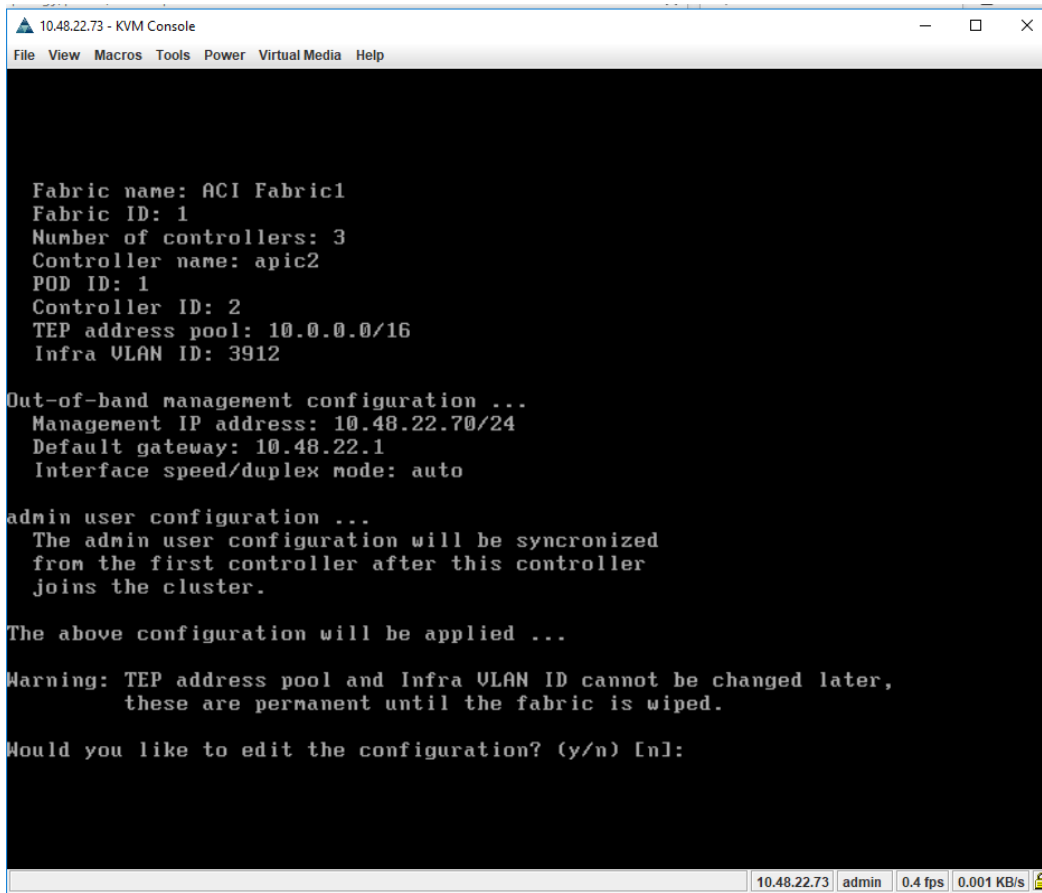
Allow APIC2 to fully boot and bring its service online. You can confirm everything was successfully configure as soon as you see the entry for APIC2 in the Active Controllers view. During this time, it will also begin syncing with APIC1's config. Allow 4-5 mins for this process to complete. During this time you may see the State of the APICs transition back & forth between *Fully Fit* and *Data Layer Synchronization in Progress.* Continue through the same process for APIC3, ensuring you assign the correct controller ID.

Figure 21. APIC Controllers Synchronizing
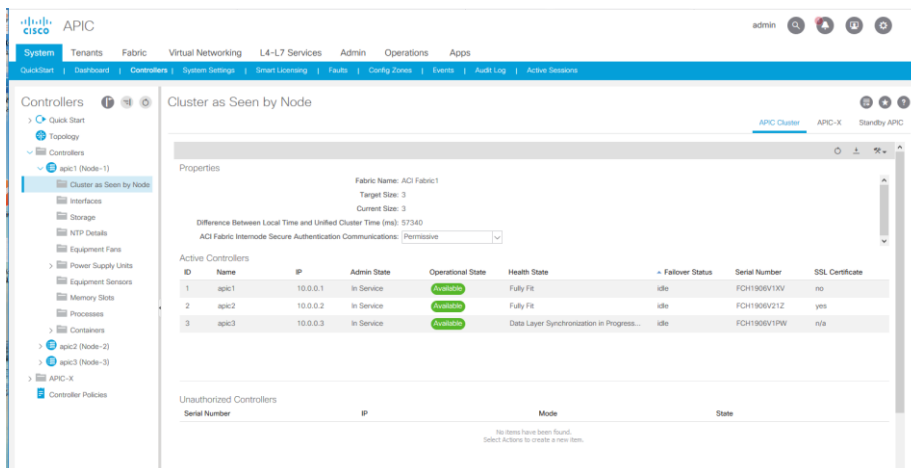
Figure 22. APIC Cluster Discovery Completed (Fully Fit)
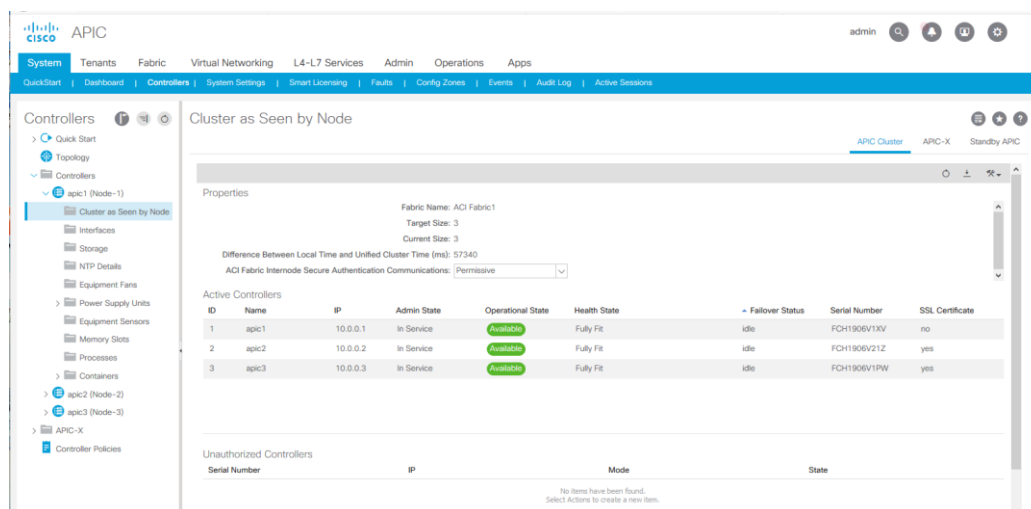
This concludes the entire fabric discovery process. All your switches & controllers will now be in sync and under a single pane of management. Your ACI fabric can be managed from any APIC IP. All APICs are active and maintain a consistent operational view of your fabric.

Setting Up Day 1 Fabric Policies

With our fabric up & running, we can now continue configuring some basic policies based on our best-practice recommendations. These policies will include:

- Out-of-Band Node Management IPs

- NTP

- System Settings (Connectivity Preferences, Route Reflectors etc)

- DNS

- SNMP

- Syslog

- Firmware Policies

- Exports (Configuration & Techsupport)

Out-of-Band Management IPs

In order to be able to connect to your switches directly and to enable some external services such as NTP, you're going to need to configure Management IPs for your switches. During the Initial fabric setup, you would have assigned Management IPs to your APICs, now we'll do the same for your Spine & Leaf nodes.

1. Navigate to Tenants > mgmt > Node Management Addresses > Static Node Management Addresses.

2. Right click on the Static Node Management Addresses folder and select *Create Static Node Management Addresses.* A window will appear and you will start by defining a range of Nodes. If the address block you intend to assign to your nodes is sequential, it simplifies the config. You will enter the range of nodes that corresponds to the sequential range of addresses you plan on assigning. For example, if you had only 2 sequential addresses, you could create the first block for your first two nodes. Additional nodes can be added indecently if your address blocks are non-sequential. Next put a check on the *Out-Of-Band Addresses* box. For the Management EPG, leave it as default. For the OOB IPv4 Address, enter the

starting IP & subnet mask. Ex. 10.48.22.77/24. This will increment this address by 1 for each Node in the range. Lastly enter the Gateway IP and click submit. If you need to add additional blocks for a different IP or node range, you can repeat the steps accordingly. It's also recommended create static addresses for each of your APICs. This will be required latter on for external services such as SNMP.

Figure 23. Creating a Static Node Management Addresses block
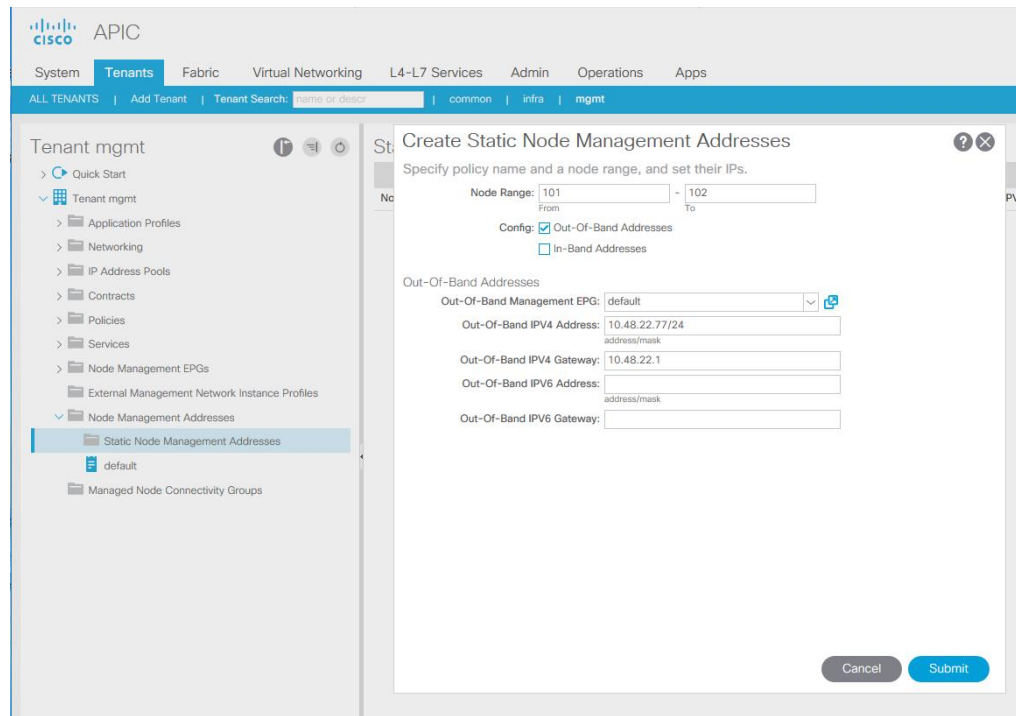
Figure 24. Completed Node Management Addresses
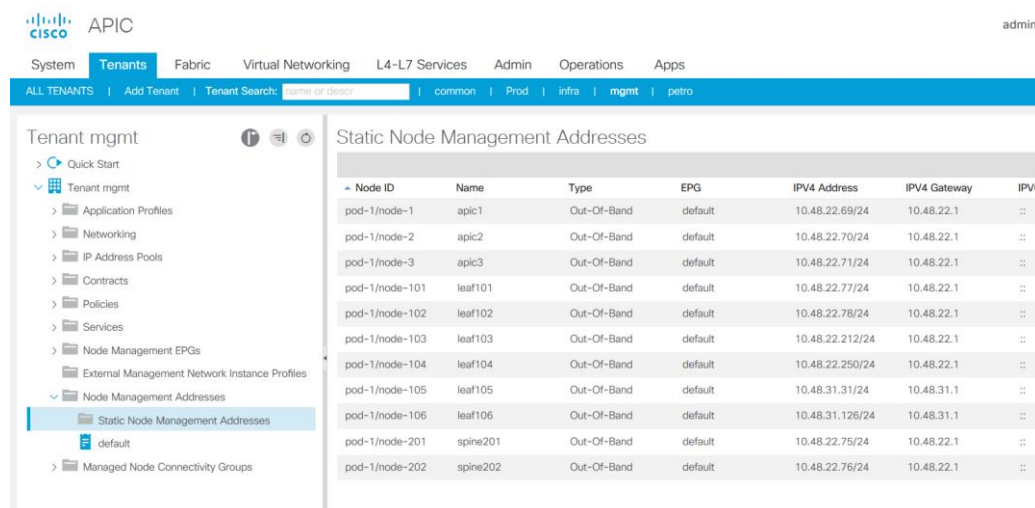
NTP Configuration

Time synchronization plays a critical role in the ACI fabric. From validating certificates, to keeping log files across devices consistent it's strongly encouraged to sync your ACI fabric to 1 or more reliable time sources. Setting up NTP is simple.

Note: Simply creating an NTP policy does not apply it to your fabric. You will also need to assign this policy to a "Pod Policy" which will be covered later in this chapter.

1. Navigate to Fabric > Quickstart, and click on the "Create an NTP Policy Link"

Figure 25. Creating NTP Policy
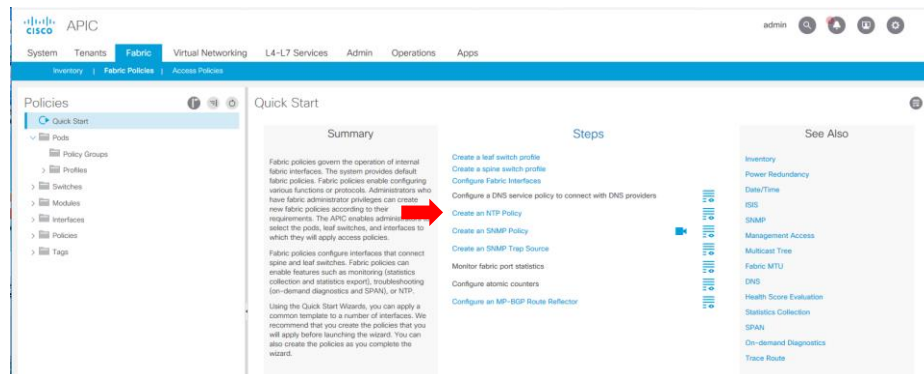
2. A new window will pop up and ask for various information. Provide a name for your policy and set the State to *Enabled.* If you'd like your Leaf switches to serve time requests to downstream endpoints, you can enable the *Server* option. For downstream endpoints they can use the Management Interface IP of nodes as an NTP Server source.   When done, click Next to define the NTP Sources.

    Note: Using a Bridge Domain SVI (Subnet IP) as an NTP Source for downstream clients is not recommended. When a Leaf switch is enabled as NTP server, it will respond on any Interface. Issues can arise when attempting to use the SVI address of a Leaf, rather than the management IP.

Figure 26. Creating Date Time Policy
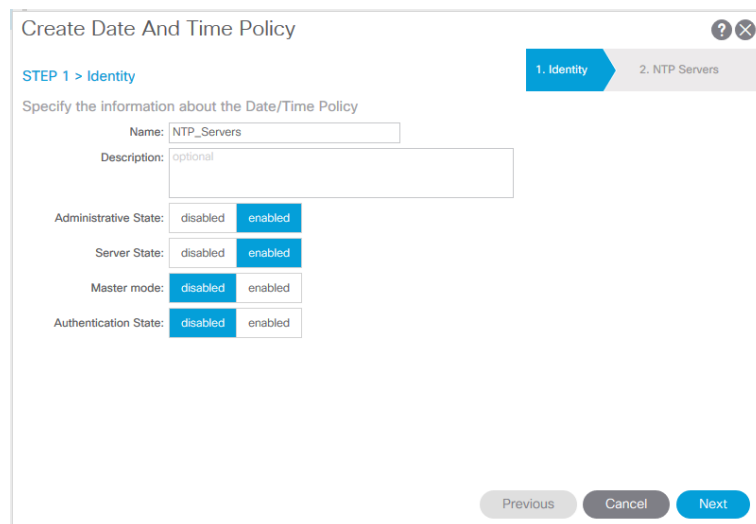
3. Next you will define one or more NTP providers. This will be the upstream device the nodes will poll for Time synchronization. You can optionally assign one of the sources as *Preferred*, which will force the switches to always to attempt time sync with this source first, then re-try against alternate sources. We recommend at least 2 different NTP sources to ensure availability. Leave all the default options and click OK.

Figure 27. Configuring NTP Providers

4. The last step for Time configuration is to set the Display Format and/or Time Zone for your fabric. Navigate to Fabric > Fabric Policies > Pod > Date and Time > default.

5. Configure the Display time to be either local time or UTC and assign the appropriate Time Zone where your APICs are located.

Figure 28. Configuring Time Zone
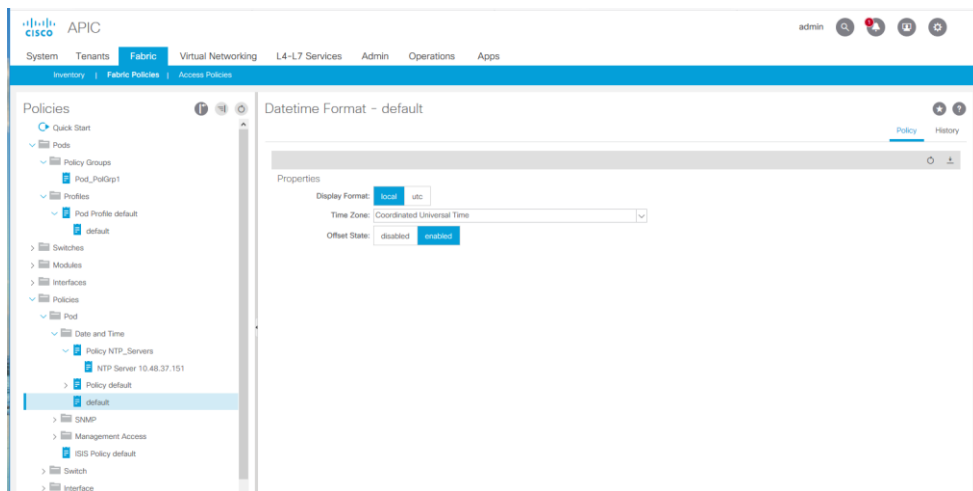
System Settings

The next section we're going to cover include a bunch of different unrelated fabric settings that are grouped in the same UI panel. Navigate to System > System Settings sub menu. Below includes a list of settings that are recommended to be changed with a brief explanation. Understand that these suggestions are suited for most customers but can be modified pending specific needs. Some of the options below may have you wondering - Why they aren't enabled by default? Some of the settings below may have been made configurable after the fact, so respective of not changing default behaviors, sometimes the optimal settings need to be changed manually.

■ APIC Connectivity Preference: Change to *ooband*. This option select which Management EPG profile should be used by default if both oob & inband management is configured. Out-of-band is far simpler to understand and fine for most use cases.

■ Global AES Passphrase Encryption Settings: Assign Passphrase and Enable (two steps). This is used to export secure fields when generating a Config Export task for your configuration. Without this enabled, all secure fields including passwords and certificates are omitted from your Config Export, requiring you to have to manually re-apply them upon config import. This makes your import/export tasks fast & secure.

- Fabric Wide Settings: Enable Force Subnet Check and Enforce Domain Validation.

  - Enabling the subnet check applies to Gen2 or later switches. This feature limits the local learning of endpoint MAC & IPs to only those belonging to a Bridge Domain defined subnet. This feature Is explained in greater detail in the ACI Endpoint Learning Whitepaper.

  - Enforcing Domain Validation restricts EPG VLAN usage by ensuring that the respective Domain & VLAN Pool are bound to the EPG. This prevents accidental or malicious programming of EPGs to use VLAN IDs they may not be permitted to use. Tenant Admins (responsible for Application level policies) typically have different permissions than the Infrastructure Admin (responsible for networking & external connectivity) in ACI – and with this separation of roles you can have one user role responsible for allowing certain VLAN Ranges per Domains using RBAC & Security Domains. Then this limits your Tenant admins access to only specific Domains and in-turn the VLANs they're permitted to assign to EPGs via Static Path Bindings etc.

  - BGP Route Reflector: The ACI fabric route reflectors use multiprotocol BGP (MP-BGP) to distribute external routes within the fabric. To enable route reflectors in the ACI fabric, the fabric selects the spine switches that will be the route reflectors and provide the autonomous system (AS) number. Once route reflectors are enabled in the ACI fabric you can configure connectivity to external routers. Assign an Autonomous System Number (ASN) to your fabric and configure up to eight Spines as Route Reflectors.   In a multipod environment it would be recommended to spread them across Pods.

Note: Simply creating a Route Reflector policy does not apply it to your fabric. You will also need to assign this policy to a "Pod Policy" which will be covered later in this chapter.

Figure 29. System Systems – Fabric Wide Settings
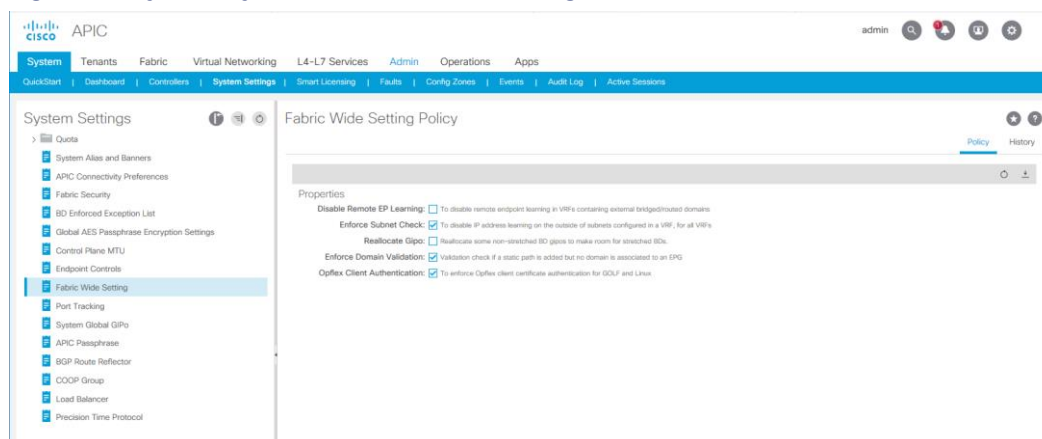
Creating & Applying a Pod Policy

Now that you've configured various fabric policies, some of which require to be assigned to your nodes via Fabric Pod Policy. You can use the default policy or create a new one. The Pod Policy group is collection of policies previously created and applied to one or more Pods in your fabric. This granularity allows you to apply different policies to different Pods if required. In our use case we're going to create a single Pod Policy Group name "PolGrp1"

1. Native to Fabric > Fabric Policies sub menu > Pods > Policy Groups folder

2. Right click on the Policy Groups folder and select *Create Pod Policy Group*

3. For the Policies we previously created, you'll want to select them here. If you configured the default polices, you can leave the default policy name selected, otherwise you'll want to select the appropriate policy name.

Cisco Systems. All printed copies and duplicate soft copies are considered uncontrolled and the original online version should be referred to for the latest version.

Page 25 of 215

Note: In our example we configured one policy for NTP with a named policy "NTP_Servers" whereas all other policies are using the "default" policy. This is to show the two valid ways to configure various Pod Policies.

Figure 30. Creating Pod Policy Group

4. Navigate to Fabric > Fabric Policies sub menu > Pods > Profiles > Pod Profile > default. With the *default* Pod Selector selected in the navigation pane, change the Fabric Policy Group to the one created in the previous step.

Figure 31. Assigning Pod Policy Group to Pod Profile
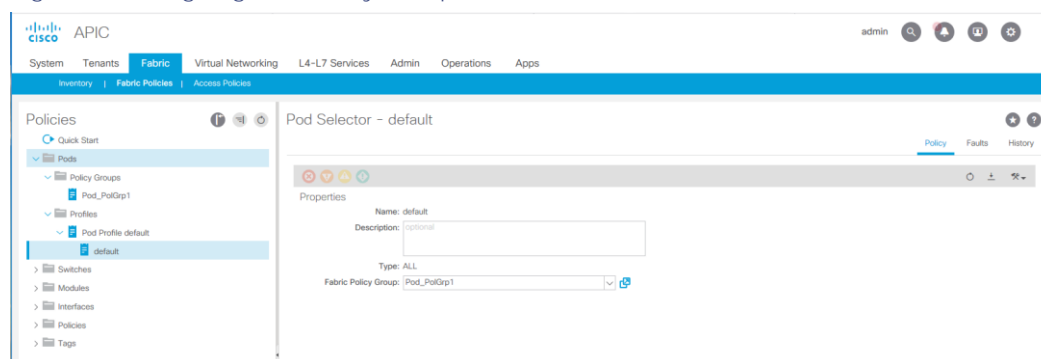
DNS Configuration

To allow your APIC to resolve hostname to external resources, such as Virtual Machine Managers (VMMs), Remote Locations etc, you'll need to configure the DNS profile.

1. Navigate to Fabric > Fabric Policies sub menu > Policies > Global > DNS Profiles > default

2. In the work pane, click the + sign to add one or more DNS providers. You can assign a Preferred provider if you wish.

3. In the DNS Domains, you can click the + sign to assign a domain suffix that should be used by your APIC when resolving DNS hostnames.

Figure 32. Configuring a DNS Profile
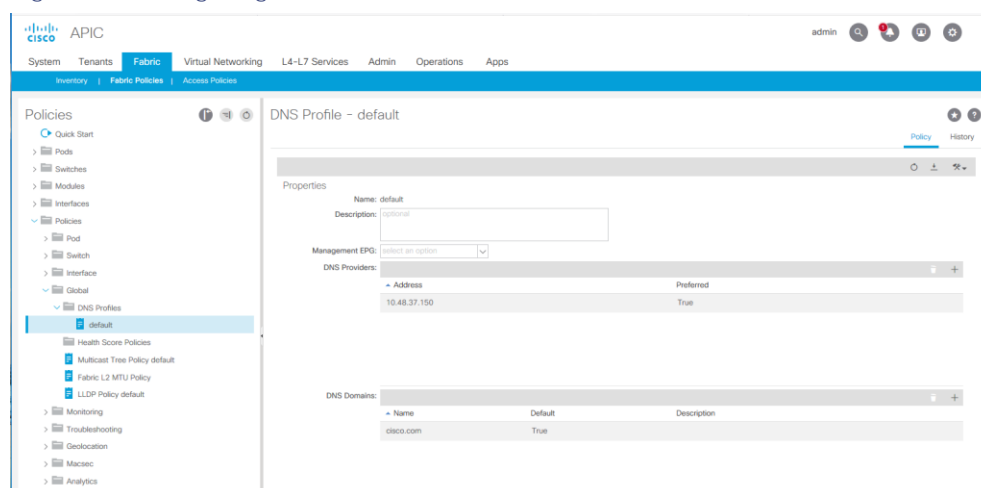
Securing Management Access

By default, the APIC will allow a few default protocols such as HTTPS & SSH to access the out-of-band (oob) management interfaces. When we start adding additional services such as SNMP and other monitoring services, we will need to explicitly add all "allowed" protocols to the Management Contracts/Filters including both HTTPS & SSH. Telnet is disabled by default and Cisco recommends using only SSH for Management CLI access as it's a more secure protocol. It's a best practice to configure your Management Contracts on Day 1 to enforce a complete whitelist model against access to your management interfaces. The next procedure will walk you through configuring

1. Navigate to Tenants > mgmt > Contracts > Filters

2. Right-click on the Filters folder and select Create Filter.

3. In the pop-up, provide a name. Ex. allow-ssh

4. Under Entries, click the + sign and add a filter entry with the following values (all other columns can be left blank):

    a. Name: allow-ssh

    b. EtherType: IP

    c. IP Protocol: TCP

    d. Source Port/Range: 22-22

    e. Dest. Port/Range: unspecified

5. Click Submit.   Next step will be creating the Contract & attach the newly created Filter.

Figure 33. Creating SSH Filter for Management Access

Figure 33. Creating SSH Filter for Management Access

6. Repeat Steps 1 – 5 to allow Web Access using HTTPS (TCP 443)

Figure 34. Creating HTTPS Filter for Management Access

7. Navigate to Tenants > mgmt > Contracts > Out-Of-Band Contracts

8. Right Click on the Out-of-Band Contracts folder and select *Create Out-Of-Band Contracts*

9. Provide a name for the contract. Ex. a*llow-management*

10. Click the + sign to add a subject. Provide a name for the Contract Subject. Ex. *allow-management*

11. Click the + sign, select each filter then click *Update* to add both filter chains previously created. Ex. *allow-ssh & allow-https*

Figure 35. Creating Contract Subject

Figure 35. Creating Contract Subject

12. Click Ok then click Submit. With the security rules created, we'll create a Management Instance Profile and attach these rules.

13. Navigate to Tenants > mgmt > External Management Network Instance Profile

14. Right-click on the folder and select *Create External Management Network Instance Profile* or select an existing Profile.

15. Provide a name for the Profile. Ex. 'default' if not already created

16. Under Consumed Out-of-Band Contracts click the + sign and add your previously created allow-management contract then click *Update*

17. Under the Subnets, enter the target Subnet you wish to restrict accessing the APIC then click *Update.* In our example we are using a 0.0.0.0/0 entry to allow unrestricted access.

18. Click Submit.

Figure 36.Adding Contracts to External Management Network Profile
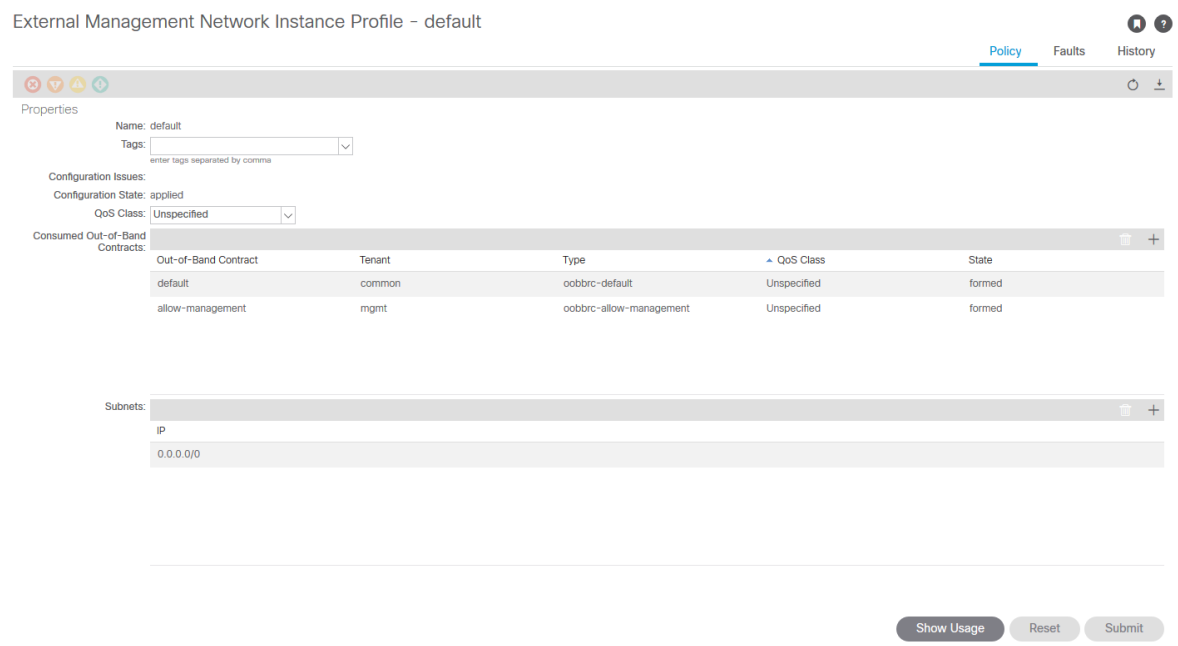
SNMP Configuration

SNMP monitoring is an important aspect of ensuring your ACI administrators are quickly made aware of any potential problems. It's just one of the many monitoring aspects you can take advantage of and widely adopted standard. Assuming you have a SNMP monitoring system already setup, this section will guide you through setting up the ACI portion of SNMP monitoring. Configuring SNMP requires that you have already configured your Static Management Node Addresses and have In-band or Out-of-Band network connections to your switches. SNMP is configured through two different Scopes; Global & VRF Context.   We will start with the Global scope configuration as the VRF Context requires your User tenant to be created.

Global Scope configuration will allow you to monitor the physical status of the fabric including Interfaces, Interface States, Interface Stats and environmental information.

1.   Navigate to Fabric > Fabric Policies sub menu > Policies > Pod > SNMP folder

2.   Right-Click on the SNMP folder and select *Create SNMP Policy*.

3.   Give your policy a name. Ex. SNMP_Pol

4.   Set the admin state to *Enabled*

5.   Optional items include the Contact & Location details.

6.   Under Community Policies click the + sign.

Creating Out-of-Band Filters

To allow SNMP via your Management Interfaces on your switches, you'll need to create and apply the respective Filter & Contracts.

1.   Navigate to Tenants > mgmt > Contracts > Filters

2.   Right-click on the Filters folder and select Create Filter.

3. In the pop-up, provide a name. Ex. allow-snmp

4. Under Entries, click the + sign and add a filter entry with the following values (all other columns can be left blank):

    a. Name: allow-snmp

    b. EtherType: IP

    c. IP Protocol UDP

    d. Source Port/Range: 161 – 162

    e. Dest. Port/Range: 161 – 162

5. Click Submit.   Next step will be creating the Contract & attach the newly created Filter.

Figure 37. Creating SNMP Filter

6. Navigate to Tenants > mgmt > Contracts > Out-Of-Band Contracts > allow-management

7. Click on the Contract Subject previously created named *allow-management*

8. Click the + sign to add a filter chain and select the previously created SNMP filter. Ex. *allow-snmp*

9. Click Ok then click Submit. With the security rules created, we'll create a Management Instance Profile and attach these rules.

Cisco Systems. All printed copies and duplicate soft copies are considered uncontrolled and the original online version should be referred to for the latest version.

Page 31 of 215

Figure 38. Add SNMP Filter to Management Contract
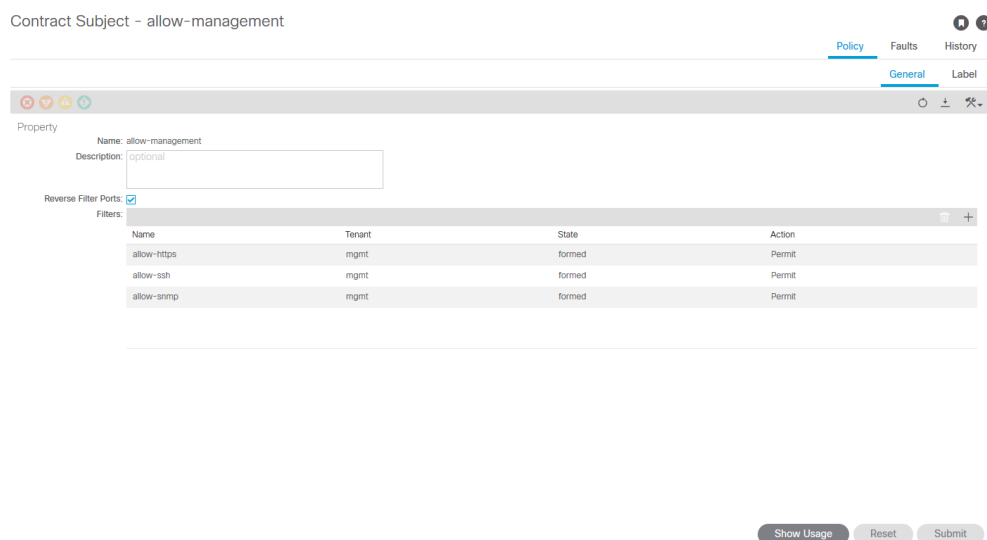
Syslog Configuration

Another useful tool for monitor the fabric is the popular Syslog policy to aggregating faults and alerts. Syslog configuration is comprised of first defining one or more Syslog Destination targets, then defining Syslog policies within various locations within the UI to accommodate Fabric, Access Policy and Tenant level syslog messages. Enable all these syslog sources will ensure the greatest amount of details are captured but will increase the amount of data & storage requirements depending on the logging level you set.

Create Syslog Remote Location

1. Navigate to Admin > External Data Collectors > Monitoring Destinations > Syslog

2. From the Actions Menu select Create Syslog Monitoring Destination Group

3. Provide a name for the Syslog Group. Ex. syslog_servers

4. Leave all other options default and click Next

5. Under Create Remote Destinations, click the "+" icon

    a. Enter hostname or IP address.

    b. If necessary, modify any additional details as you wish such as severity & port number

    c. Set the Management EPG as default (Out-of-band)

    d. Click OK.

6. If necessary, add additional Remove Destinations

7. Click Finish

Create Fabric Level Syslog Source

The fabric Syslog policy will export alerts for monitoring details including physical ports, switch components (fans, memory, PSUs etc) and linecards

1. Navigate to Fabric > Fabric Policies sub menu > Policies > Monitoring > Common Policy > Callhome/Smart Callhome/SNMP/Syslog/TACACs.

2. From the Actions Menu select *Create Syslog Source*

    a. Provide a name for the source. Ex fabric_common_syslog

    b. Leave the severity as warnings unless desired to increase logging details

    c. Check any additional Log types such as Audit Logs (optional)

    d. Set the Dest Group to the Syslog Destination Group previously created.

    e. Click Submit

Figure 39. Creating Fabric Common Syslog Source
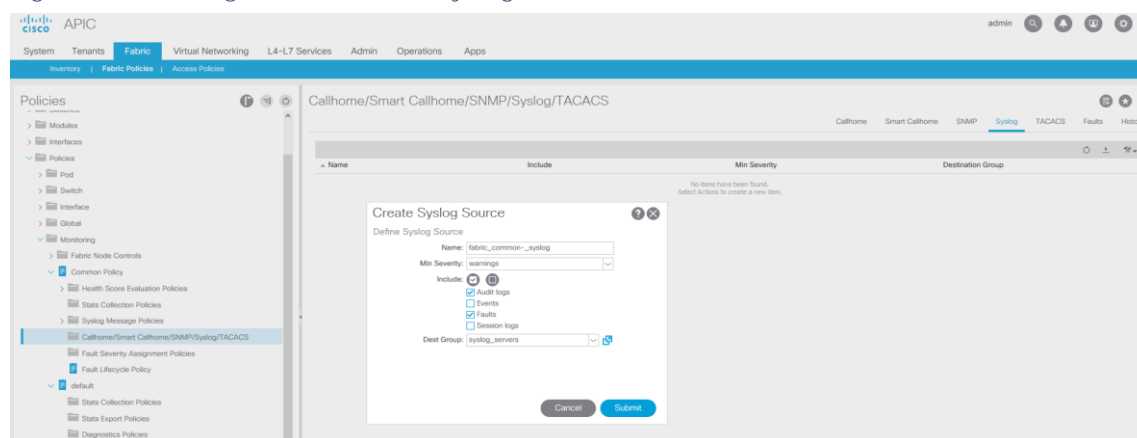
3. Navigate to Fabric > Fabric Policies sub menu > Policies > Monitoring > default > Callhome/Smart Callhome/SNMP/Syslog/TACACs.

4. In the work pane, set the Source Type to Syslog

5. Click the "+" icon to add a Syslog Source

    a. Provide a name for the source. Ex fabric_default_syslog

    b. Leave the severity as warnings unless desired to increase logging details

    c. Check any additional Log types such as Audit Logs (optional)

    d. Set the Dest Group to the Syslog Destination Group previously created.

    e. Click Submit

Figure 40. Creating Fabric default Syslog Source
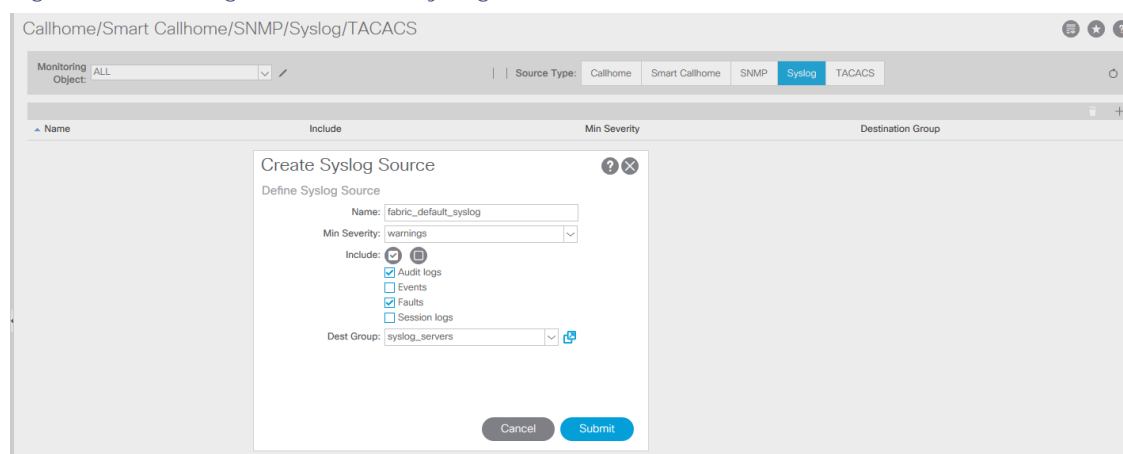
Creating Access Level Syslog Policy

The Acess Syslog policy will export alerts for monitoring details including VLAN Pools, Domains, Interface Policy Groups, and Interface & Switch Selectors Policies.

1. Navigate to Fabric > Access Policies sub menu > Policies > Monitoring > default > Callhome/Smart Callhome/SNMP/Syslog/TACACs.

2. In the work pane, set the Source Type to *Syslog*

3. Click the " +" icon to add a Syslog Source

   a. Provide a name for the source. Ex access_default_syslog

   b. Leave the severity as warnings unless desired to increase logging details

   c. Check any additional Log types such as Audit Logs (optional)

   d. Set the Dest Group to the Syslog Destination Group previously created.

   e. Click Submit

Figure 41. Creating Access Level Syslog Policy
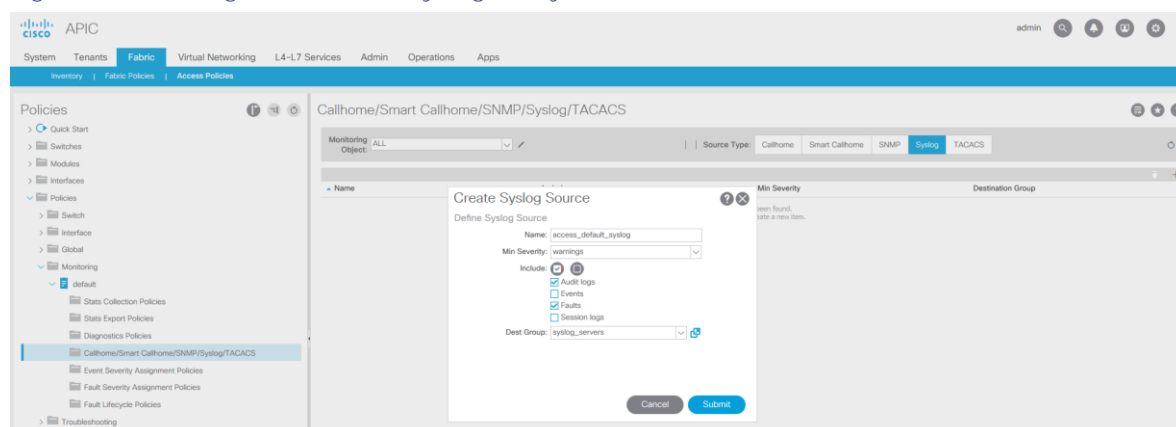
Creating Tenant Level Syslog Policies

Tenant level syslogging will include all tenant related polices include Application Profiles, EPGs, Bridge domains, VRFs, external networking etc. To simplify the syslog configuration across multiple tenants you can

leverage Common Tenant syslog configuration and share that across other tenants. This would provide a consistent level of logging for all tenants. Alternately if you wanted to have varied levels of logging/severities for different Tenants, you could create the respective Syslog policy within each tenant. For our purposes we will deploy a single consistent syslog policy using the Common Tenant.

1. Navigate to Tenants > common > Policies > Mentoring > > default > Callhome/Smart Callhome/SNMP/Syslog/TACACs.

2. In the work pane, set the Source Type to *Syslog*

3. Click the "+" icon to add a Syslog Source

    a. Provide a name for the source. Ex tenant_default_syslog

    b. Leave the severity as warnings unless desired to increase logging details

    c. Check any additional Log types such as Audit Logs (optional)

    d. Set the Dest Group to the Syslog Destination Group previously created.

    e. Click Submit

Figure 42. Creating a Common tenant Syslog Policy
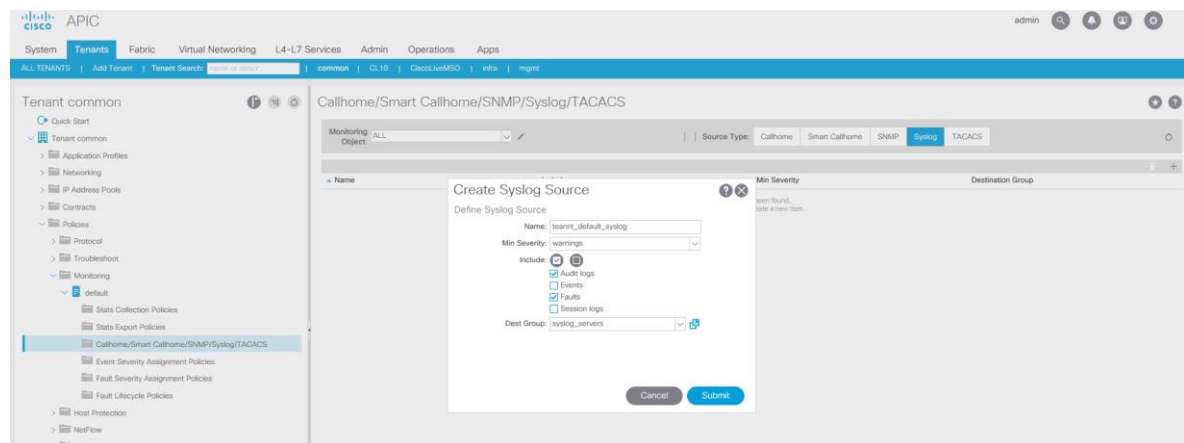
4. Navigate to Tenants > *Your_Tenant* > Policy tab

5. Set the *Monitoring Policy* drop down box to be the default policy from the common tenant.

Figure 43. Assigning default common Tenant syslog Policy to a Tenant
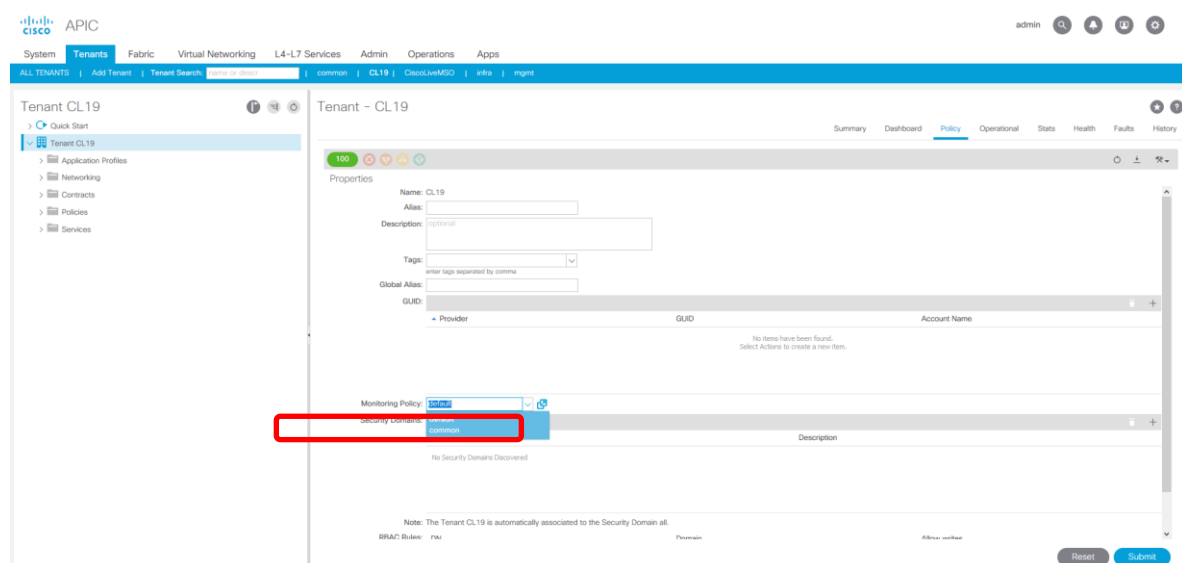
Upgrade & Firmware Policies

Firmware upgrades may not be thought as a Day 1 task but setting up the appropriate policies is. By default, when you install APIC from scratch, it will be deployed only with running firmware. The APICs firmware repository will be empty at this point. Our first task is going to be to load the "current" firmware images for the version you plan on running on your fabric today. In our example we deployed the APIC with 4.0(1h), and therefore we'll need to download the APIC & Switch images for 4.0(1h) / 14.0(1h) respectively to the APIC.

Adding Software Images to the APIC

*It's assumed by this point you've already downloaded the appropriate APIC & Switch images from Cisco.com and stored them locally.

1. Navigate to Admin > Firmware > Images tab

2. From the Actions drop down, select *Add Firmware to APIC*

3. From here you can select a Local Upload via your browser from your workstation, or you can use a Remote location. Remote locations can be HTTP or SCP sources. Select the appropriate option.

4. Enter the appropriate details and click submit. If the details are valid, the download status should progress.

5. Repeat the process for the corresponding second image (switch or APIC).

6. Once completed, you should both images in the repository.

Figure 44. Uploading APIC FW to APIC via HTTP

Figure 45. Uploading Switch FW to APIC via SCP

Figure 46. Viewing Software Images Within APIC Repository
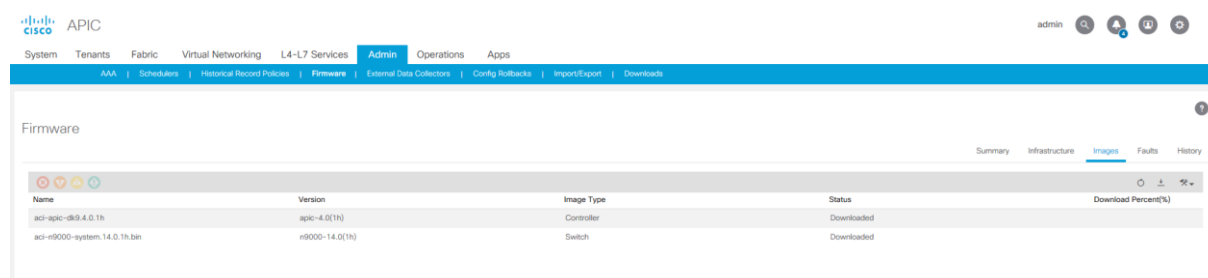
About ACI Upgrades

When performing ACI fabric upgrades, there are typically two categories of upgrades. One task for the controllers and one or more subsequent tasks for the switch nodes. The controller upgrade task completely automated. Once initiated, each controller node will upgrade serially. This ensures that the next controller begins its turn only once the previous controller has come back online and sync'd back with the cluster. Any standby controllers will also be upgrade at this time. Once all controllers have been upgraded, the switches will be next. A common practice is to group nodes by odd or even nodes IDs. Typically switch nodes are grouped together in redundant pairs by function such as Border Leafs, Compute Leafs, Service Leafs etc. Assuming all functional node pairs contain at least one odd & one even node, the impact will be minimal. This does assume that most if not all workloads are dual connected to Leaf nodes with an Odd & Even ID. The odd/even group upgrades are simply a guideline. You may wish to further divide your switch upgrade groups based on physical pods, remote leafs, and Virtual Leafs (vLeafs). For each additional group comes with less risk in the event any issues manifest with switches running the new software version, but this will have an impact on the overall upgrade window duration. A standard upgrade for a Leaf switch is typically 15-30mins, but this can often vary based on the underlying software activities required. Depending on whether the upgrade is a major or minor version change, additional components may be upgraded during this time, for example Erasable Programmable Logic Devices (EPLDs), SSD firmware and other switch components. Switches upgraded within the same Upgrade Group will be upgrade simultaneously, except for VPC Peer switches. VPC Peer Switches will always be upgraded serially regardless of the upgrade group membership to prevent unexpected network outages. When considering on upgrading, always review the Release Notes & Firmware Management Guide for the new version to confirm the upgrade path, compatibility and known caveats. An Upgrade Matrix tool is available to help you determine the correct upgrade path:
https://www.cisco.com/c/dam/en/us/td/docs/Website/datacenter/apicmatrix/index.html

We strongly recommends that you connect, configure & test console access to all controllers and switches prior to attempting any upgrade. In the event a controller or switch upgrade encounter issues, this may be the only method to access the device to troubleshooting and/or recover.

Upgrading APIC

Assuming you've uploaded the respective Software images to the APIC, you can now initiate the upgrade process.

1. Navigate to Admin > Firmware > Infrastructure tab > Controllers

2. From the Actions Menu, select Schedule Controller Upgrade.

3. Set the target firmware version

4. Set the Upgrade start time

5. [Optional] Ignore Compatibility Check. This will allow you to perform an upgrade that might not follow a supported upgrade path, which could potentially cause disruptions. For pre-Prod/Lab this is safe to enable.

6.   Click Submit.

Allow the controller process to complete which could take anywhere between 30mins > 1hr+ depending on the number of controllers in the cluster.

Upgrading Switch Nodes

Once all controllers have been successfully upgraded, and the cluster shows "fully fit", you can proceed with upgrading the switches. The procedure below assumes you intend on performing the upgrade immediately. Alternately, you can schedule the upgrade to start during a scheduled date/time.

1.   Navigate to Admin > Firmware > Infrastructure tab > Nodes

2.   From the Actions menu, select Schedule Node Upgrade

3.   Select the group type (Switch or vPod). Always perform the Switch upgrade prior to upgrading vPods

4.   Provide a name for the Upgrade group. Ex. " *Odd_Nodes*"

5.   Set the target firmware

6.   Check Enable Graceful Maintenance

7.   From the Node Selection, choose Manual

8.   Check the Switch nodes that correspond to your group. Ex. 101, 103, 201

9.   Click Submit

Figure 47. Creating Switch Upgrade Tasks for Odd nodes

Based on your Start Time option, the upgrade will immediately begin. You can remain on this screen to monitor progress. During the process the switches will download the appropriate software image from the APIC, perform the software upgrade, then reload with the new image. During this time, you may see the

Cisco Systems. All printed copies and duplicate soft copies are considered uncontrolled and the original online version should be referred to for the latest version.

Page 39 of 215

Switches disappear from the Fabric Inventory. Barring no issues, each switch will reappear once the upgrade has been completed. Be patient as some switches may take slightly longer than others.

Export Policies

Export policies encompass a wide range of options which include Configuration Backup and troubleshooting logs. Some of these policies can and should be setup during your deployment phase. A few common polices to configure include a Weekly (or even daily) configuration backup task, and a Core Export policy. Both of these options will come in handy if/when needed. The next section will walk you through setting up a few of these recommended policies.

Setting up a Remote Location

1. Navigate to Admin > Import/Export sub-tab > Export Policies > Remote Locations folder

2. From the Actions menu, select *Create Remote Location*

3. Enter a friendly name for the Remote Location. Ex. *lab_ftp*

4. Provide the Hostname / IP for the remove device

5. Select the respective protocol: ftp / scp / sftp

6. Enter the Remote path. Ie. */home/files*

7. Provide your username & password

8. Click Submit

Now with the Remote Location created, you will be able to use/reference it with subsequent export policies. Note: The remote location path & credentials will not be validated until an Export police attempts to send data to it.

Core Export Policy

In the event a service or Data Management Engine (DME) fail, it's helpful to provide the core dump file to Cisco TAC for analysis. Configuring a Core Export policies ensures a copy of this is maintained in the event of a device failure. Along with the Core Export you can optionally include a techsupport bundle when the failure occurs – this prevents any log roll over if an issue occurs and goes unnoticed for an extended period of time.

1. Navigate to Admin > Import/Export > Export Policies > Core > default

2. In the work pane, change the collection type to be *Core and TechSupport.*

3. Uncheck the Export to Controller check box

4. Set the Export Destination to the Remote Location previously created.

Daily Configuration Export Policy

Since the configuration file for ACI is relatively small, it's a good practice to configure a Daily Configuration export task.

1. Navigate to Admin > Import/Export > Export Policies > Configuration

2. From the Actions Menu select Create Configuration Export Policy

3. Provide a name for your policy. Ex. DailyConfigExport

4. Choose the desired format JSON/XML.

5. From the Scheduler drop down box, select Create Trigger Scheduler

   a) Enter a name for the Scheduler. Ex. daily_12am

   b) Under the Schedule Window click the "+" icon to add

   c) Change the Window Type to Recurring

   d) Enter a Window Name. Ex. daily_12am

   e) Click Ok

   f) Click Submit

6. Change the Export Destination to the one created Previously.

7. Click Submit

Note: The Global AES Encryption Setting should be already enabled. If not, please refer to earlier in this guide to assign an AES Passphrase and enabling Encryption. Configuring a Configuration Export Policy without this enabled will remove all Secure Fields & Passwords from your Exported file, making the Import process more involved requiring having to manually re-assign the passwords for VMM domains, Remote Locations and other policies.

Smart Licensing

Smart Licensing for Cisco Product provides customers with a single pane of glass to manage all their software subscriptions & licenses for purchased & entitled products. Smart Licensing within ACI is currently not enforced. Though Cisco highly encourages the use of Smart Licensing, failing to register your fabric with the Cisco Smart Software Manager (CSSM) has no impact on functionality. Registering your fabric with CSSM requires that the APIC has reachability to tools.cisco.com and DNS properly configured. There is a great Technote that fully explains the Registration & licensing process for ACI, so it will not be covered further in this guide. Please see the following link for details on configuration Smart Licensing for ACI:
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/3-x/smart_licensing/b_Smart_Licensing.html

Cisco Systems. All printed copies and duplicate soft copies are considered uncontrolled and the original online version should be referred to for the latest version.

Page 41 of 215

Fabric Configuration

This section describes the fabric configuration.

Introduction to Access Policies

How do you configure a VLAN on a port in an ACI fabric? In this section we will explain how these specific policies are configured and how you can achieve this goal. To begin with it's important to understand that we do not configure a VLAN directly on a port but use policies which will allow us to scale configuration and apply similar behavior to a group of objects such as switches or ports.

Let's focus on a sample use case. We have a server connected to our ACI fabric. This server has 2 NIC's and these adapters are configured in an LACP port-channel. The server is connected on port 1/9 on ACI leafs 101 and 102. On these ports we would like to configure the EPG name "EPG-Web" which will result in the server being able to provide Web services in the network. In order todo so we will use an external VLAN 1501 which will allow us to extend the EPG out of the ACI fabric. For more information regarding Tenant / APP / EPG configuration and extension please consult the "Tenant" chapter.

Switch policies

The first thing we need to do is create a policy in ACI which defines which switches need to be used. For this we use a "Switch profile" in which we define trough a "Switch Selector" which switches are part of it.

Figure 48. Switch policies
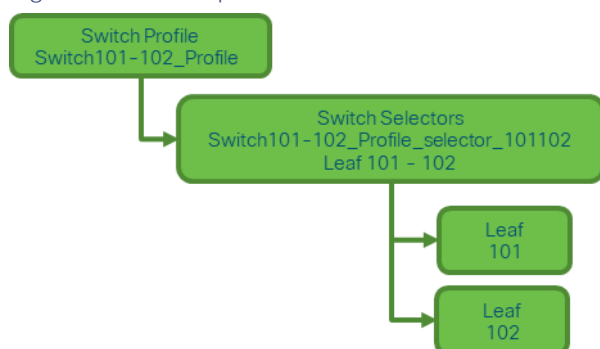
In most ACI deployments we recommend configuring 1 switch profile per switch, and 1 switch profile per vPC domain using a naming scheme which indicates the nodes which are part of the profile.

The wizard deploys a naming scheme which can easily be followed. The name scheme consists of Switch{node-id}_Profile. As an example "Switch101_Profile" will be for a switch profile containing node-101 and Switch101-102_Profile for a switch profile containing switches 101-102 which are part of a vPC domain.

Interface Policies

Once we have defined which switches need the configuration we can define the ports on these switches which require the configuration. We do this by creating an "Interface Profile" which consists of 1 or more "Access Port Selectors".

Figure 49. Interface policies

To form the relationship between this "Interface Profile" and the switches involved we link the "Switch Profile" to the required "Interface Profile.

Interface profiles can be used in many different ways. Similar to switch profiles, a single interface profile can be created per physical switch along with an interface profile for each vPC domain. These policies should have a 1 to 1 mapping to their matching switch profile. Following this logic, the fabric access policies are greatly simplified, and easy for other users to follow.

The default naming schemes employed by the wizard can also be used here. It consists of the switch profile name + ifselector to indicate this profile is used to select interfaces. An example would be Switch101_Profile_ifselector. This Interface profile would be used to configure non vPC interfaces on switch 101 and it would map directly to the Switch101_Profile switch profile.

Figure 50. Switch and interface profiles combined – step1
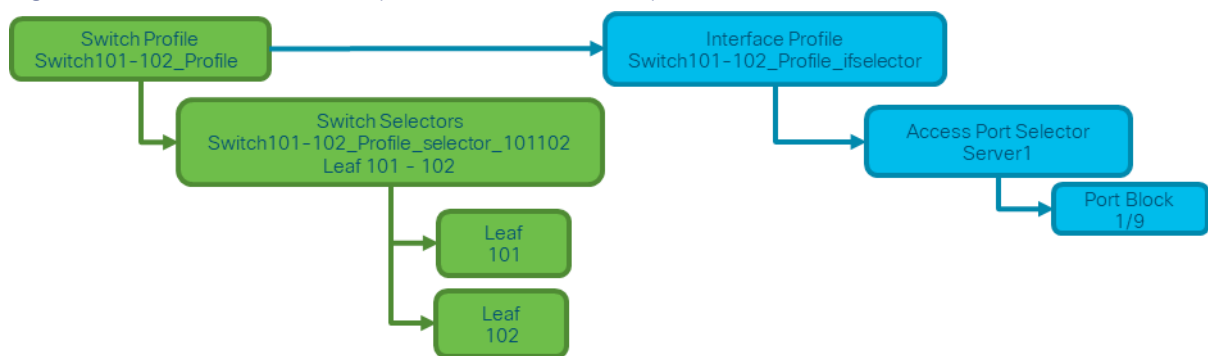
Notice because we have linked an "Interface Profile" with 1 port to the "Switch Profile" which includes 2 leaves, we have now with a simple link associated both switches with port Eth1/9.

At this point in time we have linked ports to switches but we did not define the characteristics of these ports. In order todo so we need to define an "Interface Policy Group" which will define the properties of this port. In our case because we need to create an LACP port-channel, we will create a "VPC Interface Policy Group".

Figure 51. Policy group

To form the relationship between this "VPC Interface Policy Group" and the involved interface we link towards this "Interface Policy Group" from the "Access Port Selector".

Figure 52. Switch and interface profiles combined – step2

vPC

As we are creating an LACP port-channel over 2 switches, we need to define a VPC between Leaf 101 and Leaf 102. In order todo so we create the following objects linked to the leaves.

Figure 53. vPC

VLAN pools

As a next step we will create the VLAN 1501 in our fabric. In order todo so we need to create a "VLAN Pool" with "Encap Blocks"

When considering the size of your vlan pool ranges, keep in mind that most deployments only need a single VLAN pool and one additional pool when using VMM integration. If you plan on bringing VLANs from your legacy network into ACI, define the range of legacy VLANs as your static VLAN pool.

As an example, lets assume VLANs 1-2000 are used in our legacy environment. Create one Static VLAN pool which contains VLANs 1-2000. This will allow you to trunk ACI Bridge Domains & EPGs towards your legacy fabric. If you plan on deploying VMM, a second dynamic pool can be created using a range of free VLAN IDs.

When deploying VLANs on a switch, ACI will encapsulate Spanning-tree BPDUs with a unique VXLAN ID which is based on the pool the VLAN came from. Due to this, it is important to use the same VLAN pool whenever connecting devices which require STP communication with other bridges.

VLAN VXLAN IDs are also used to allow vPC switches to synchronize vPC learned mac and IP addresses. Due to this, the simplest design for VLAN pools is to use a single pool for static deployments and creating a second one for dynamic deployments.

Figure 54. VLAN Pool

Domains

The next step we do is to create a domain. A domain defines the 'scope' of a VLAN pool, i.e. where that pool will be applied. A domain could be physical, virtual or external (either bridged or routed). In our example I will

use a physical domain as I need to connect a bare metal server into the fabric. I will link this domain to my "VLAN Pool" giving access to the required vlan(s).

Figure 55. Physical Domains
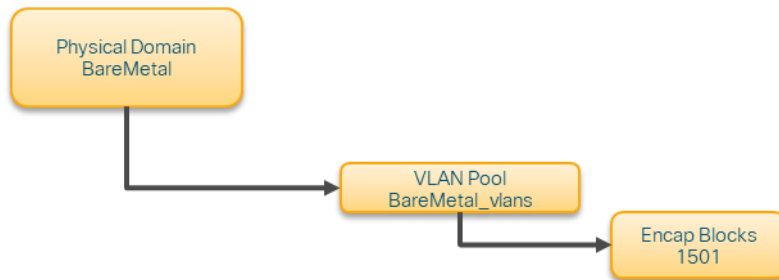
For most deployments, a single physical domain is sufficient for static path deployments and a single routed domain to allow the creation of L3Outs. Both of these can map to the same static VLAN pool. If the fabric is deployed in a multi-tenancy environments or more granular control is required to restrict which users can deploy specific EPGs & VLANs on a port, a more strategic deployment needs to be considered. Domains also provide the ability to restrict user access with Security Domains using Roles Based Access Control (RBAC).

Attachable Access Entity Profile

We now have built two blocks of configuration, on one side we have the switch and interface configurations and on the other side we have the Domain/VLAN(s). We will use an object called "Attachable Access Entity Profile" or AEP to glue these two blocks together.

A "Policy Group" is linked towards an AEP in a "1:n" relationship which means that an AEP's goal is to group Interfaces on Switches together which share similar policy requirements. This means that you can further on in the fabric refer to one AEP which is representing a group of interfaces on specific switches.

Figure 56. Attachable access entity profile
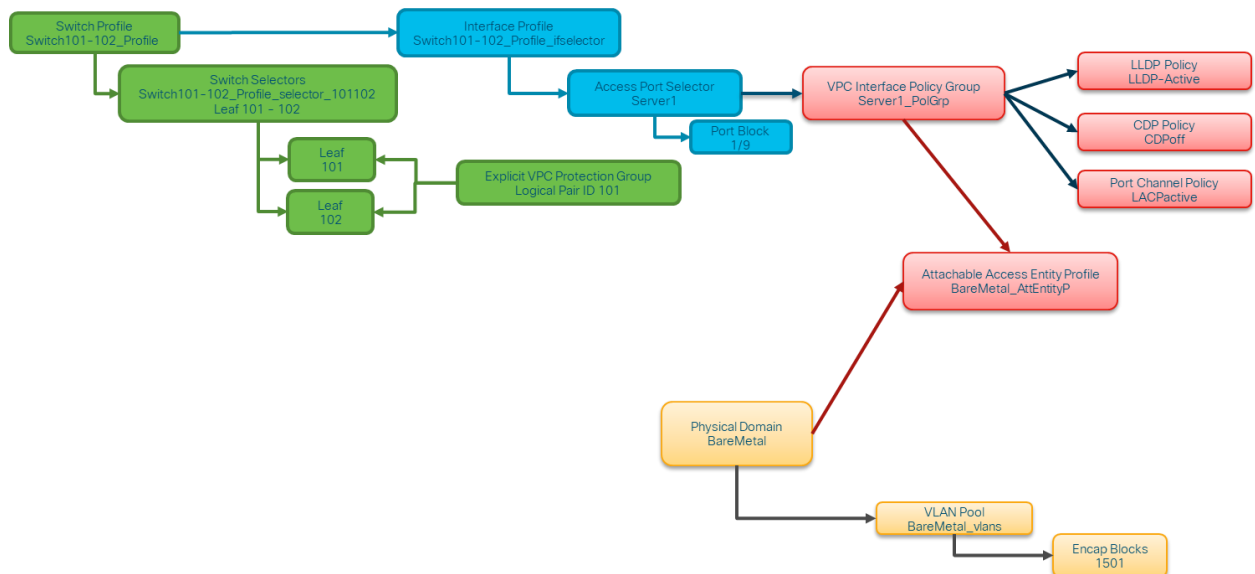
In most deployments, a single AEP should be used for static paths and one additional AEP per VMM domain.

Cisco Systems. All printed copies and duplicate soft copies are considered uncontrolled and the original online version should be referred to for the latest version.

Page 45 of 215

The most important consideration is that VLANs can be deployed on interfaces through the AEP. This can be done by mapping EPGs to an AEP directly or by configuring a VMM domain for pre-provision. Both these configurations make the associated interface a Trunk port ('switchport mode trunk').

Due to this, it is important to create a separate AEP for L3Out when using routed ports or routed sub-interfaces. If SVIs are used in the L3Out, it is not necessary to create an additional AEP.

Tenant, APP and EPG

ACI uses a different means of defining connectivity using a policy-based approach.

The lowest object we have available is called an "Endpoint Group" or EPG. We use the construct of an EPG to define a group of VMs or servers with similar policy requirements. To group these EPG's together we use a logical construct called "Application Profiles" which are part of a specific tenant.

Figure 57. Tenant, APP and EPG
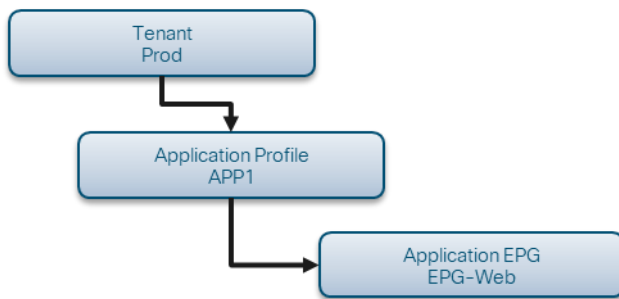
The next step is to link the EPG to the domain, hence, making the link between the logical object representing our workload (the EPG) and the physical switches / interfaces.

Figure 58. EPG to Domain link

Cisco Systems. All printed copies and duplicate soft copies are considered uncontrolled and the original online version should be referred to for the latest version.

Page 46 of 215
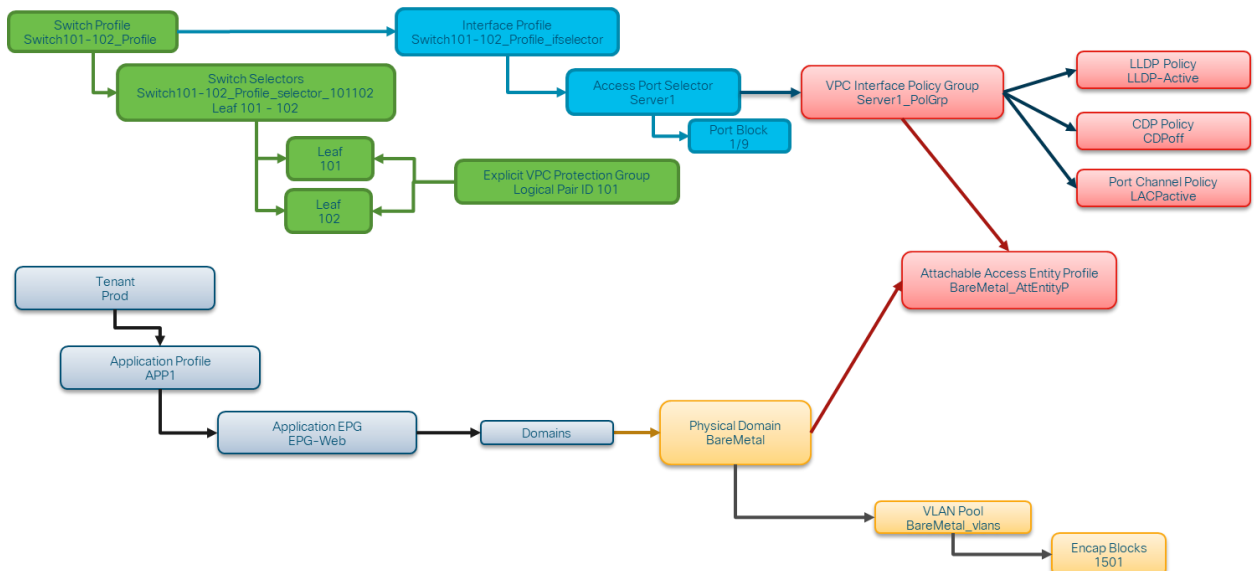
Static Bindings

The last step because we are using a "Physical" domain is that we need to tell the EPG where exactly to program which VLAN out of the "VAN Pool". This will allow the EPG to be stretched externally out of the fabric and it will allow us to connect the bare metal server into the EPG.

Figure 59. Static Bindings

The referenced "Port Encap" off course needs to be resolvable against the "VLAN Pool".

All Together

Figure 60. Bare metal ACI connectivity
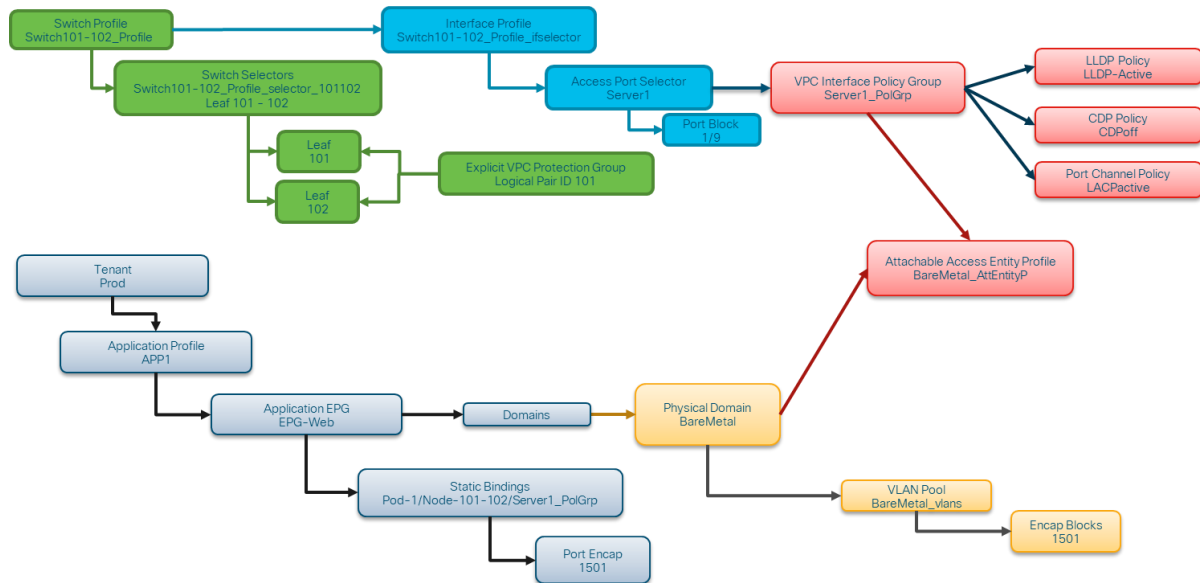
Connecting One More Server

With all the previous policies created, what would it mean to connect one more server on port Eth1/10 on leaf switches 101 and 102 with a port-channel?

Looking at the previous image "Bare metal ACI connectivity", it means we have to create:

■ An extra "Access Port Selector" + "Port Block"

■ An extra "VPC Interface Policy Group"

■ An extra "Static Binding" with "Port Encap"

Notice that due to the fact that we are using an LACP port channel, we need to use a dedicated "Interface Policy Group", being a "VPC Interface Policy Group", because we need to have a single VPC per port-channel / "Access Port Selector".

In the case we would have been using individual links, we could re-use the "Interface Policy Group" we previously created also for the extra server.

The resulting policies would look like the following image.

Figure 61. Connecting server1 into our setup
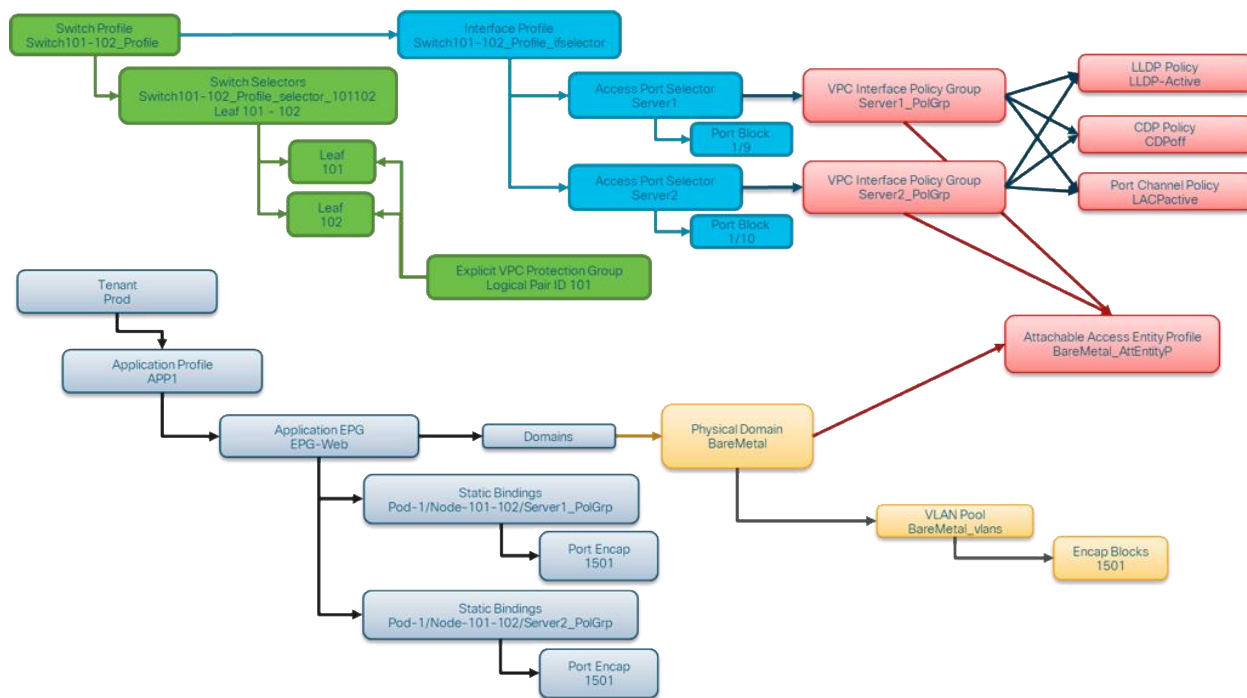
Blade Chassis Connectivity with VMM

In this section we will connect a UCS Fabric Interconnect running a VMware workload into the ACI fabric. As we are starting with a new clean fabric, this means we must create.

■   Access policies to connect FI-A

■   Access policies to connect FI-B

■   VMM Domain

Also keeps in mind that UCS FI connectivity in combination with VMware has the following characteristics:

■   From UCS FI towards the ACI leaf switches we will run a port-channel. We will create 1 for each FI

■   From the ESXi server running on the blade towards the FI's we cannot create a port-channel as we are connecting to 2 individual switches, this means we cannot use a load-balancing algorithm that uses IP-hashing. Make sure to use a load balancing algorithm such as "Virtual Port ID" on VMware side which is switch dependent and which will hence support this topology.

See also the below diagram:

Figure 62. UCS FI blade connectivity
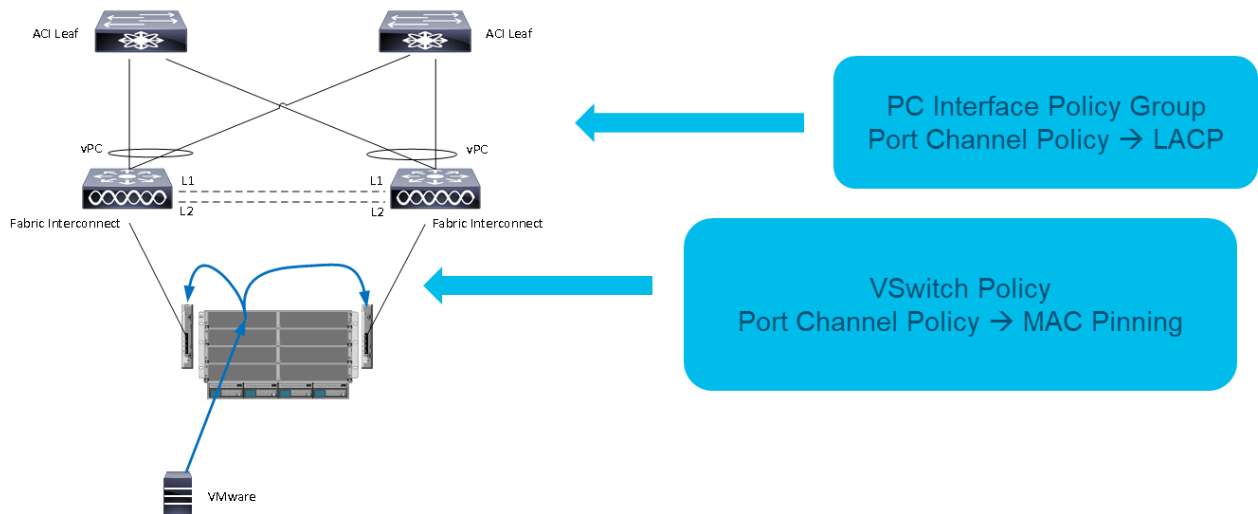
Our connectivity diagram for this UCS Fabric Interconnect looks like the following:

Figure 63. UCS FI blade connectivity in this document
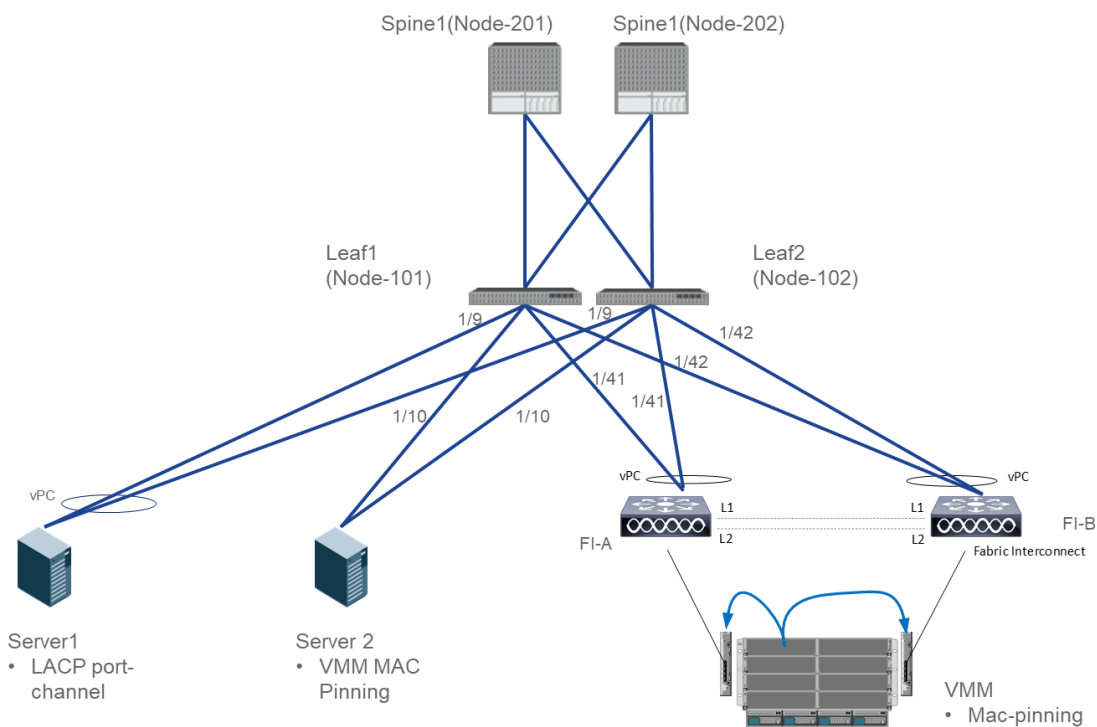
Notice that the FI's are connected symmetric to the ACI fabric, meaning, FI-A is connected to Eth1/41 on both leaf101 and leaf102. FI-B is connected to port Eth1/42 on leaf101 and leaf102. Using symmetric connectivity will greatly simplify the configuration.

Connecting the First UCS Fabric Interconnect (FI-A)

The first step we will do is launch the wizard to create the access policies involved.

1. You can do this by navigating to Fabric > Access Policies > Quick Start

Figure 64. Connecting FI-A: Step 1
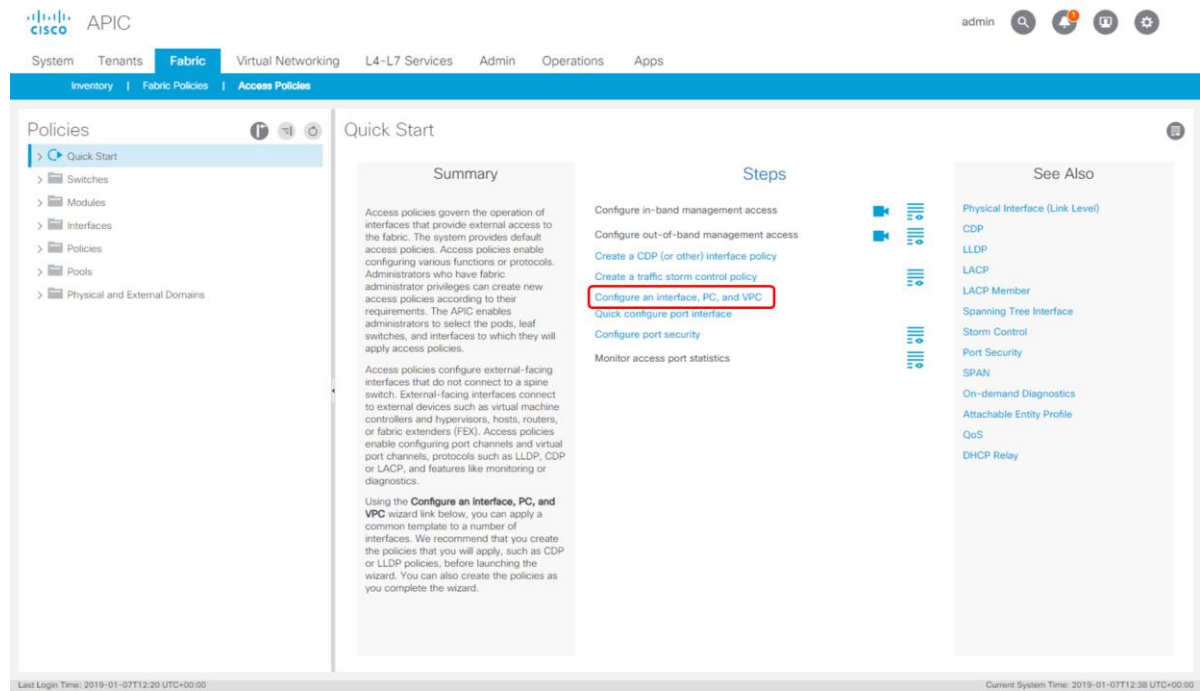
2. This will bring us to the following screen where we will first create a VPC pair on leaf switches 101 and 102.

Figure 65. Connecting FI-A: Step 2
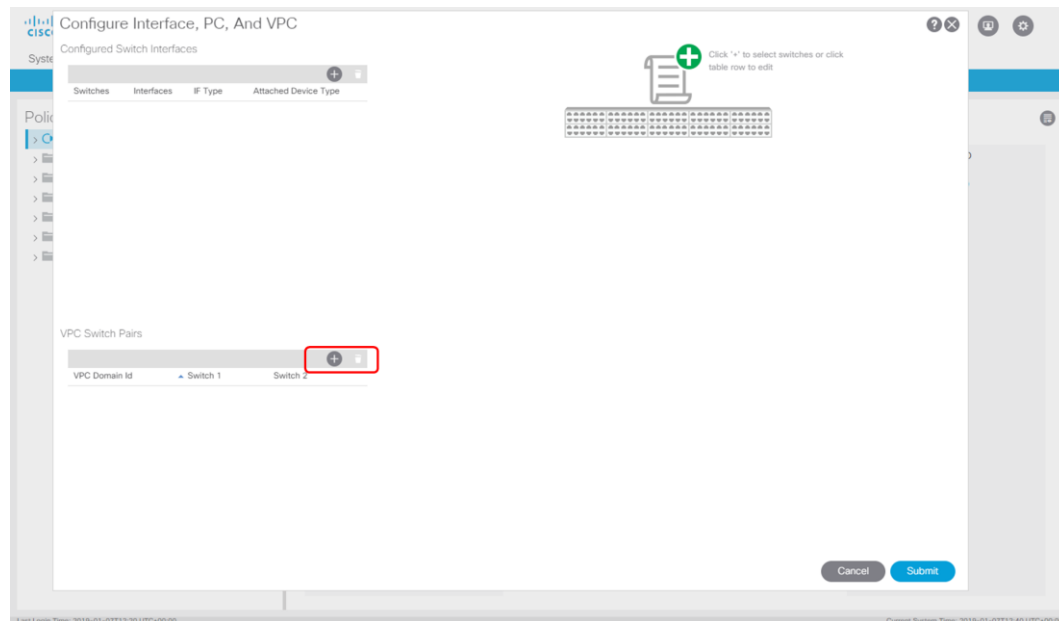
3. Fill in the fields as per the below screen

Figure 66. Connecting FI-A: Step 3
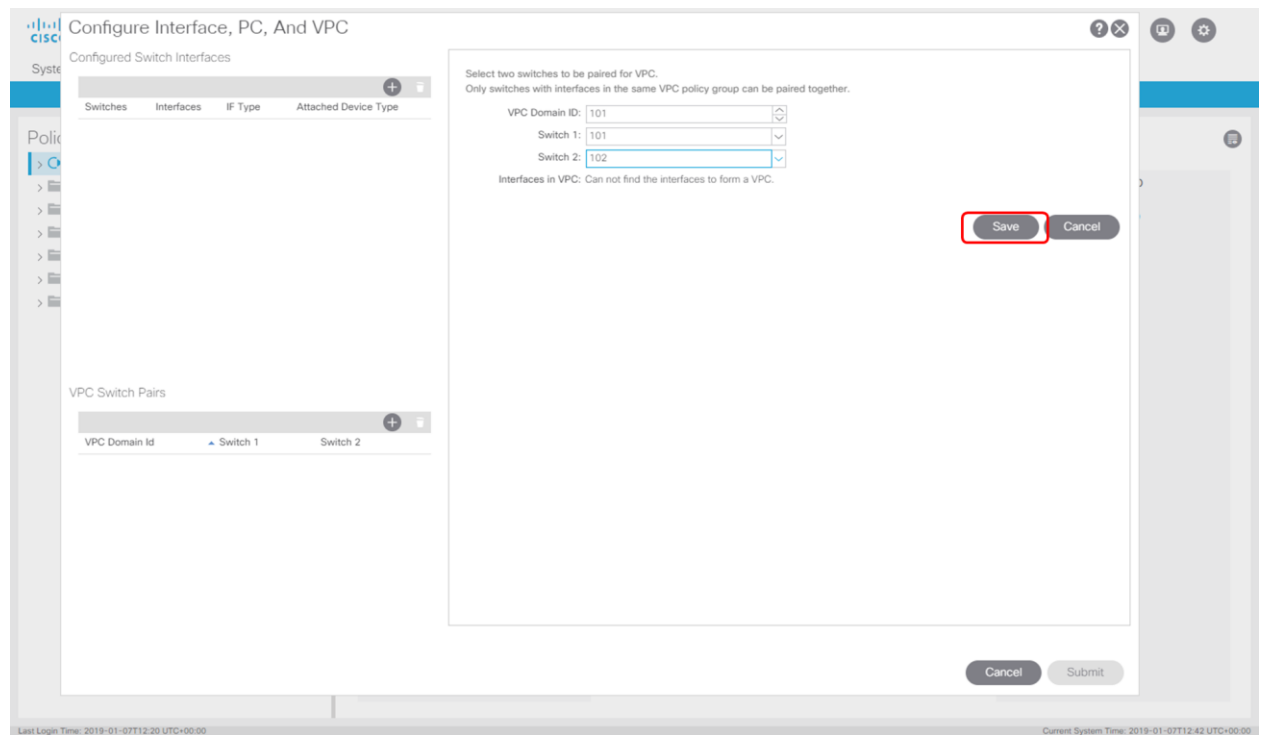
4. Notice we must fill in a unique ID for the "VPC Domain ID". In our case we have chosen 101 referring to the first leaf in our VPC port-channel. When done make sure to click "Save".

Figure 67. Connecting FI-A: Step 4
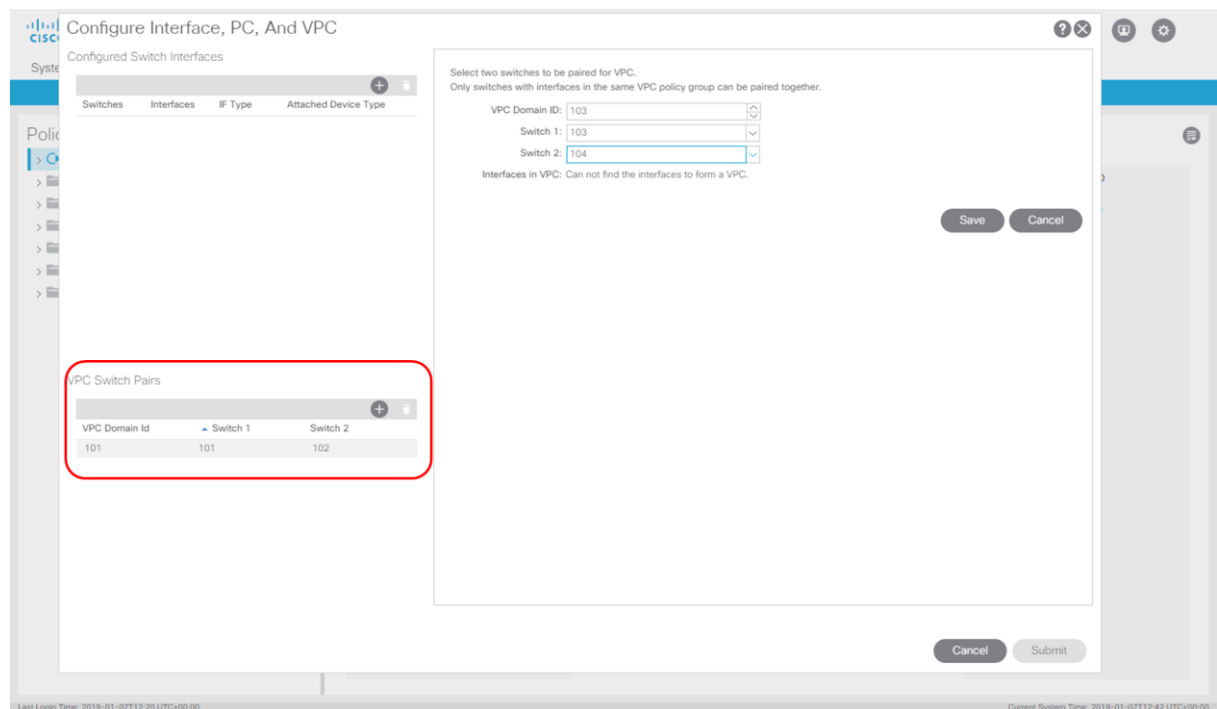
After saving you can see the VPC appearing on the left in the wizard showing that you have correctly created this.

5. You can verify in the "Access Policies" that this policy has been correctly created:

Figure 68. Connecting FI-A: Step 5
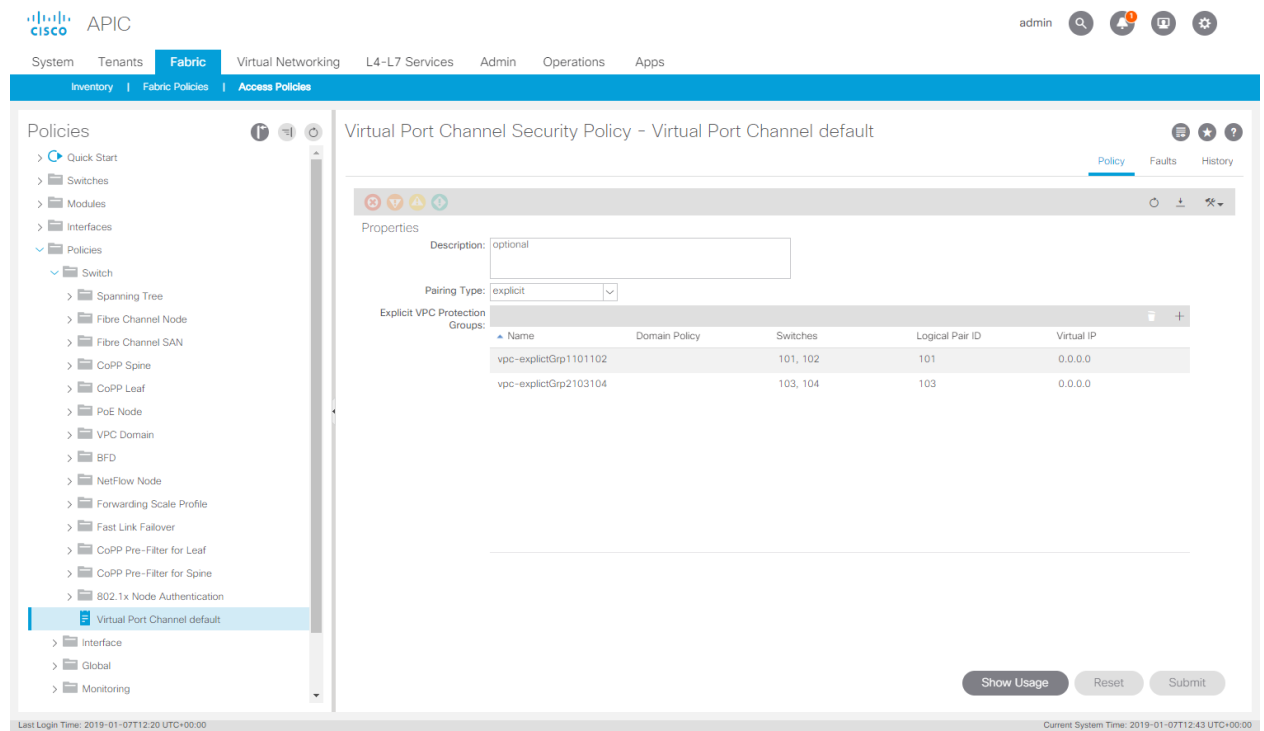
6. Next we will start the wizard again and on the left select the VPC pair where we need to create the UCS FI-A connectivity.

Figure 69. Connecting FI-A: Step 6
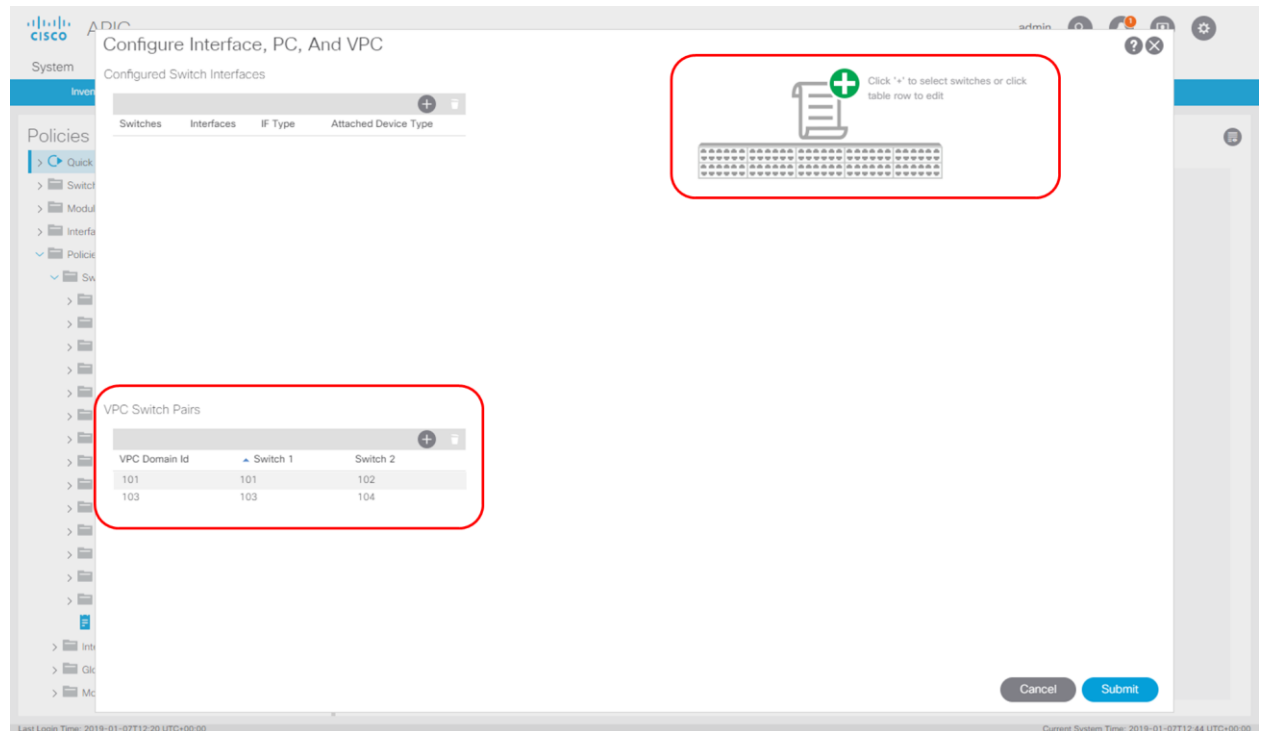
7. After doing this we click on the + sign on the top right to start the wizard.

Figure 70. Connecting FI-A: Step 7
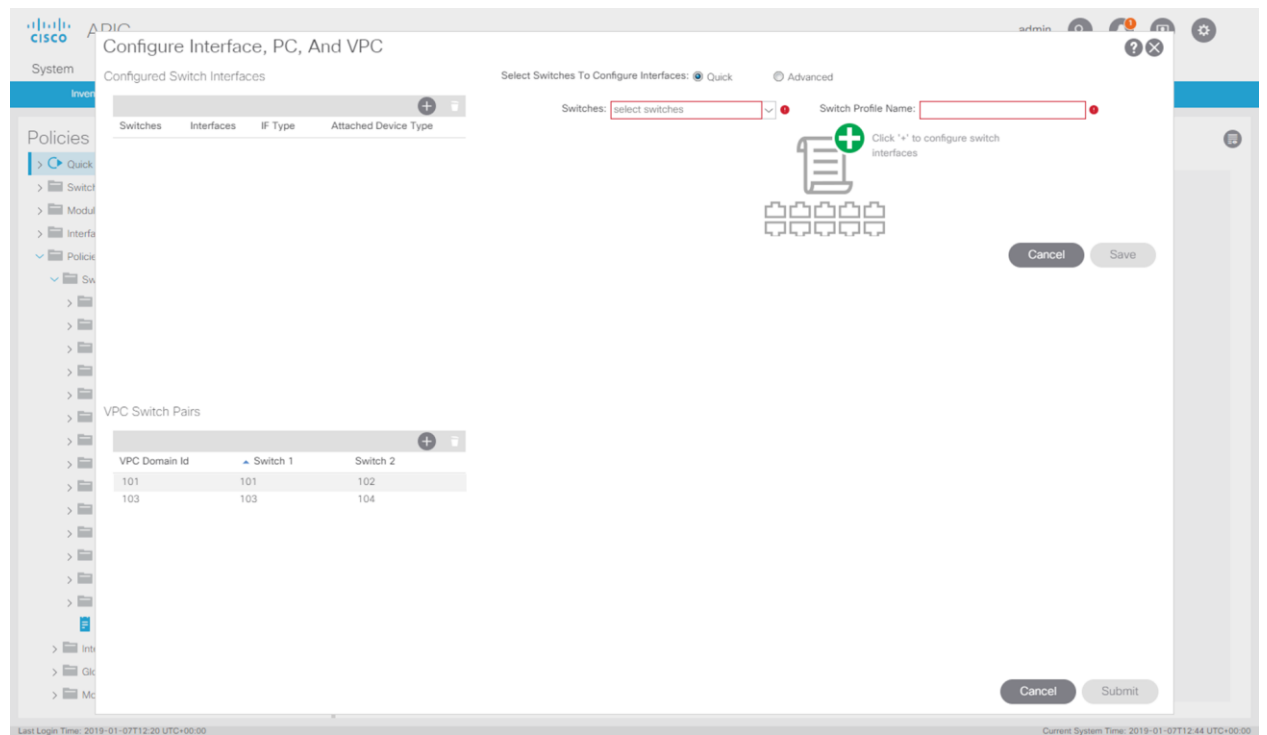
8. The first step we need to do is create a "Switch Policy".

Figure 71. Connecting FI-A: Step 8
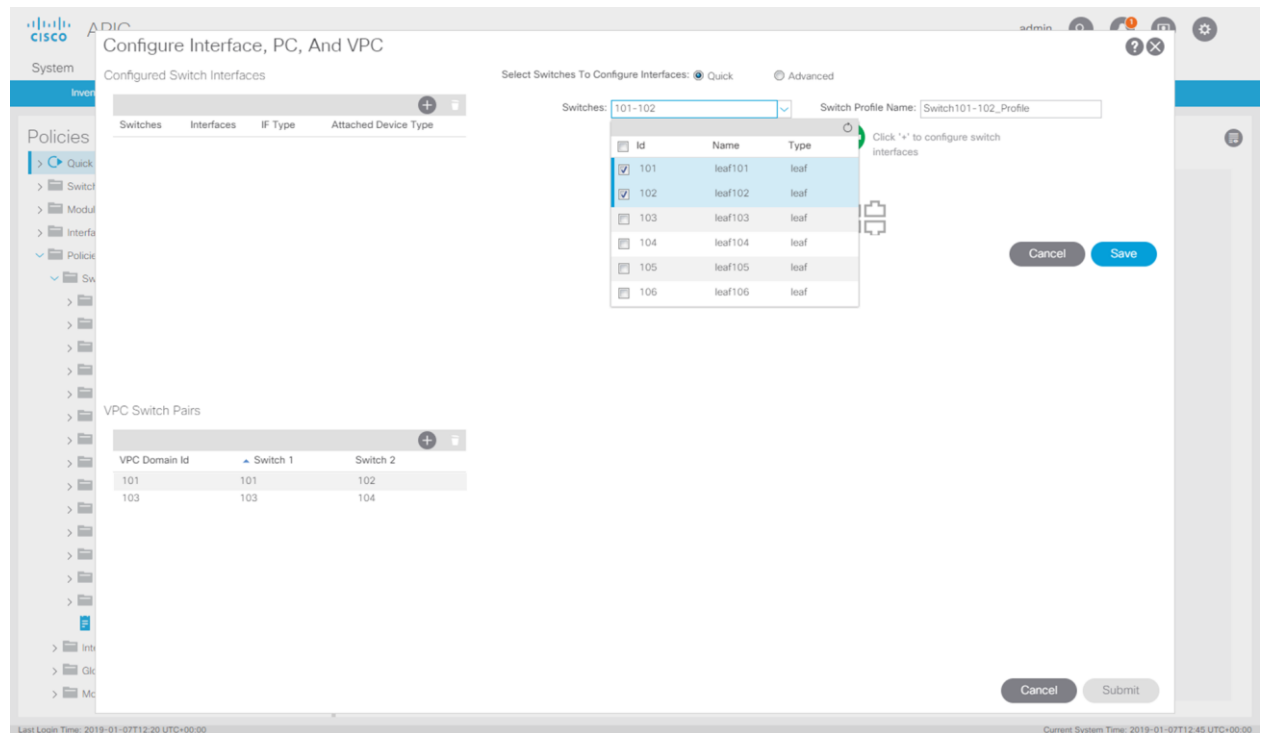
9. After filling in the required information, click "Save".

Figure 72. Connecting FI-A: Step 9
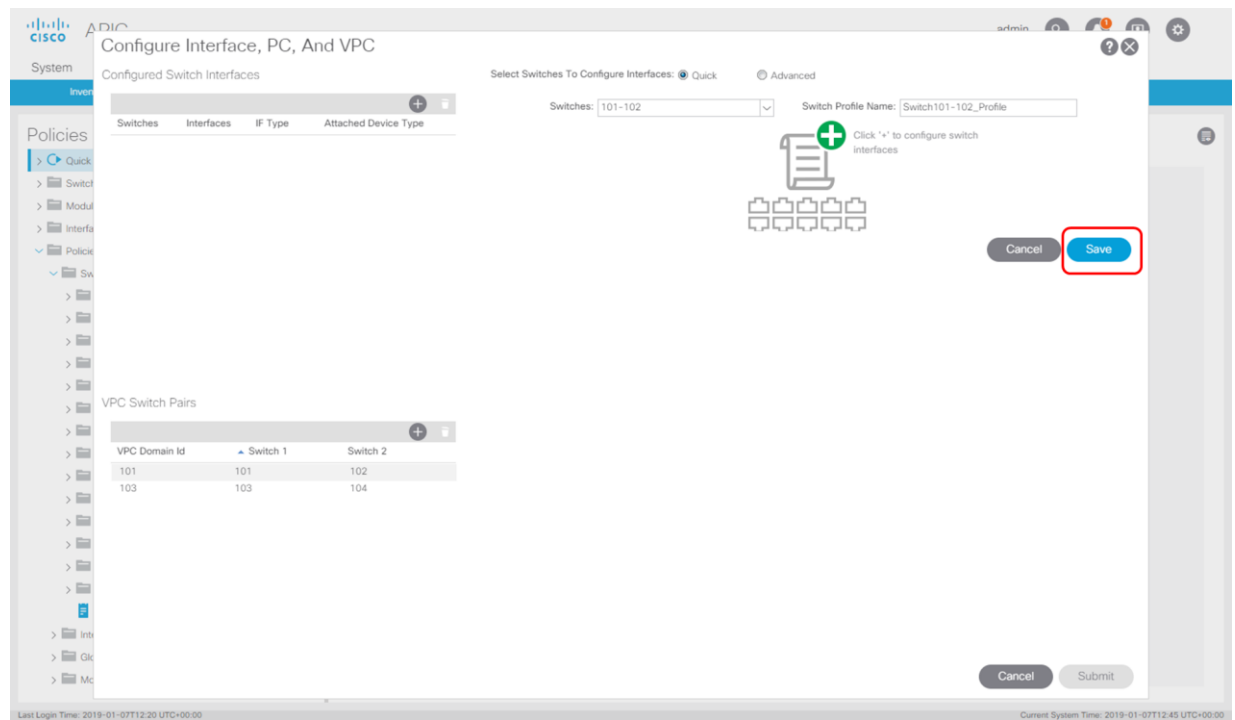
10. This will launch the "Interface Policy" creation wizard.

Figure 73. Connecting FI-A: Step 10
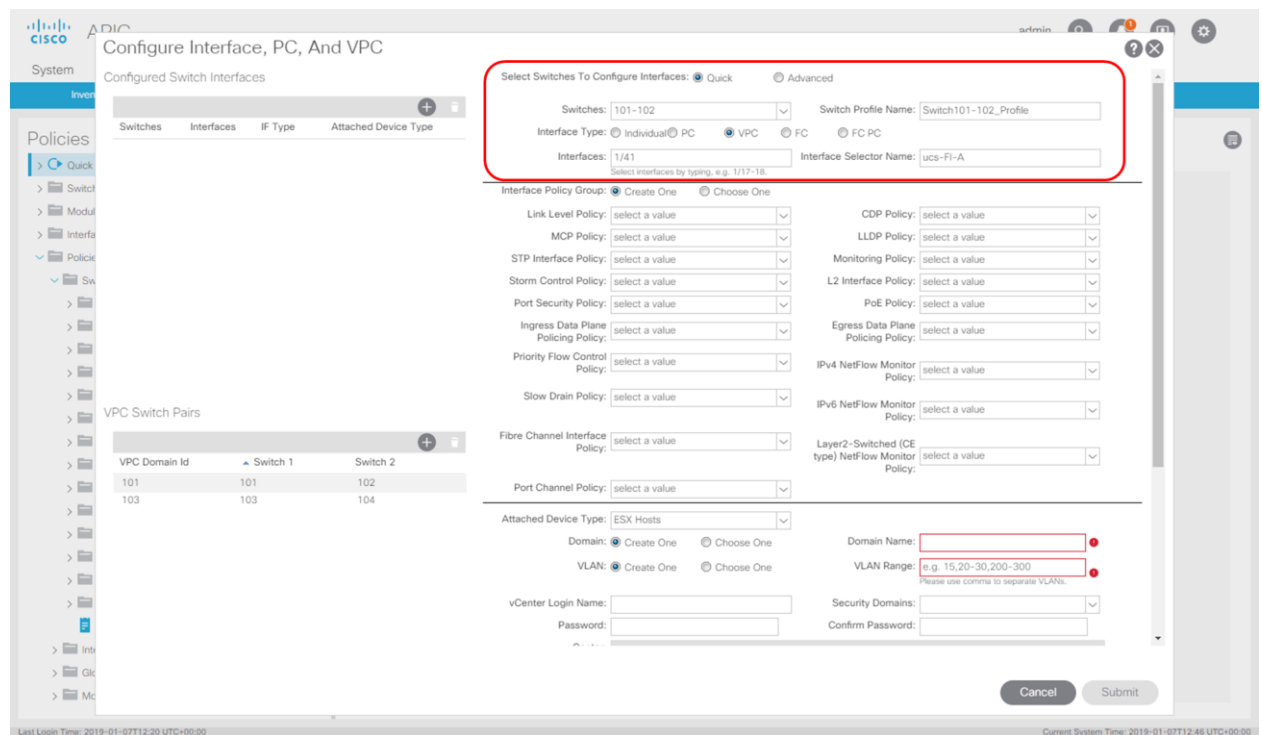
We will first fill in the interface where FI-A will connect, and we will also give it the name "ucs-FI-A".

11. After this we will create a CDP policy.

    Give the "CDP Interface Policy" a name, select "Admin State" to "Disabled" and select "Submit".

Figure 74. Connecting FI-A: Step 11
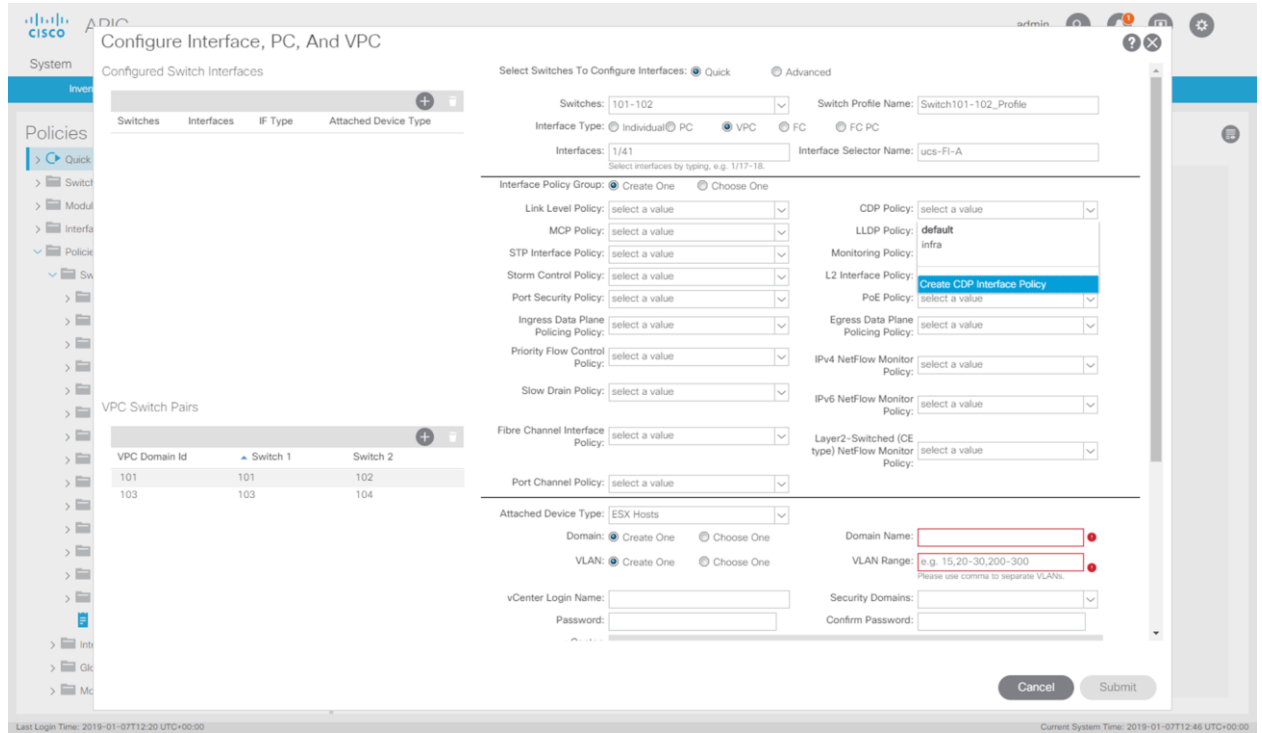
12. We will now create the LLDP policy.

Figure 75. Connecting FI-A: Step 12
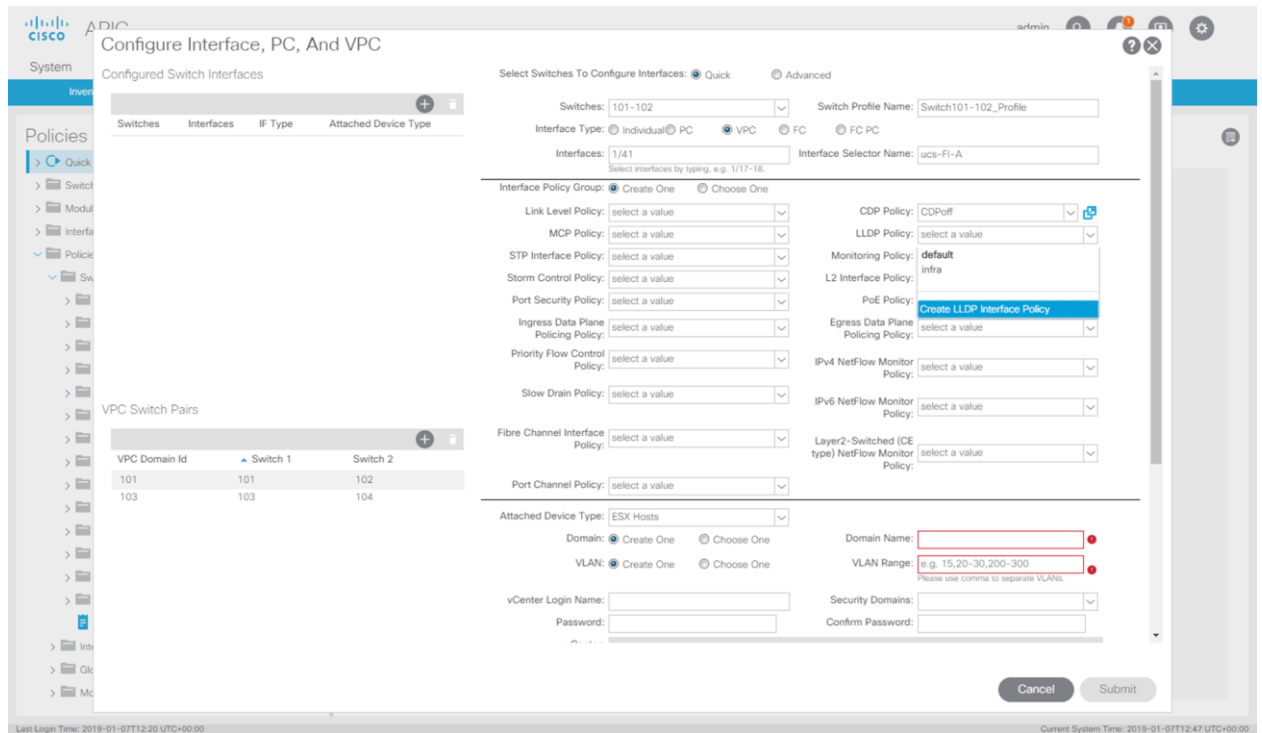
13. Give the policy a name and click "Submit".

Figure 76. Connecting FI-A: Step 13

14. In the next step we will create a "Port Channel Policy".

Figure 76. Connecting FI-A: Step 14

15. Give this policy a name, select the correct "Mode" and click "Submit".

Figure 77. Connecting FI-A: Step 15

16. Now click on the "+" sign to add a new vCenter. In order to do so first verify in your vCenter client what the name of your Datacenter is.

    We will first select "Associated Device Type" as "ESX Hosts" and then fill in a "Domain Name" and "VLAN Range". We will then configure the "vCenter Login Name" and "Password".

Figure 77. Connecting FI-A: Step 16

Figure 78. Connecting FI-A

17. After having filled in the required information click "OK" to continue.

Figure 78. Connecting FI-A: Step 17

18. Now as a last step we will fill in the required vSwitch configuration.

Figure 79. Connecting FI-A: Step 18

Notice we are using "Mac Pinning-Physical-NIC-load" which is a switch-independent protocol due to the use of FI's. We also enable LLDP at vSwitch level.

19. Now click " Save"

Figure 80. Connecting FI-A: Step 19
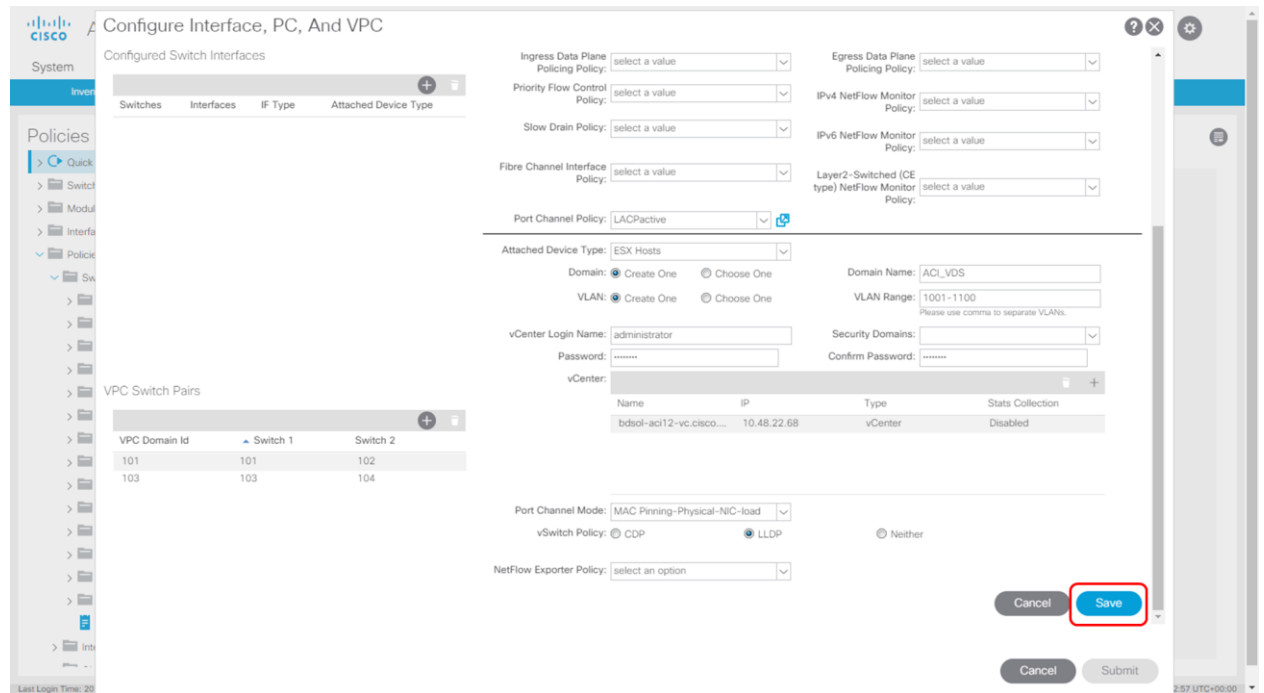
20. Now click " Save"

Figure 81. Connecting FI-A: Step 20
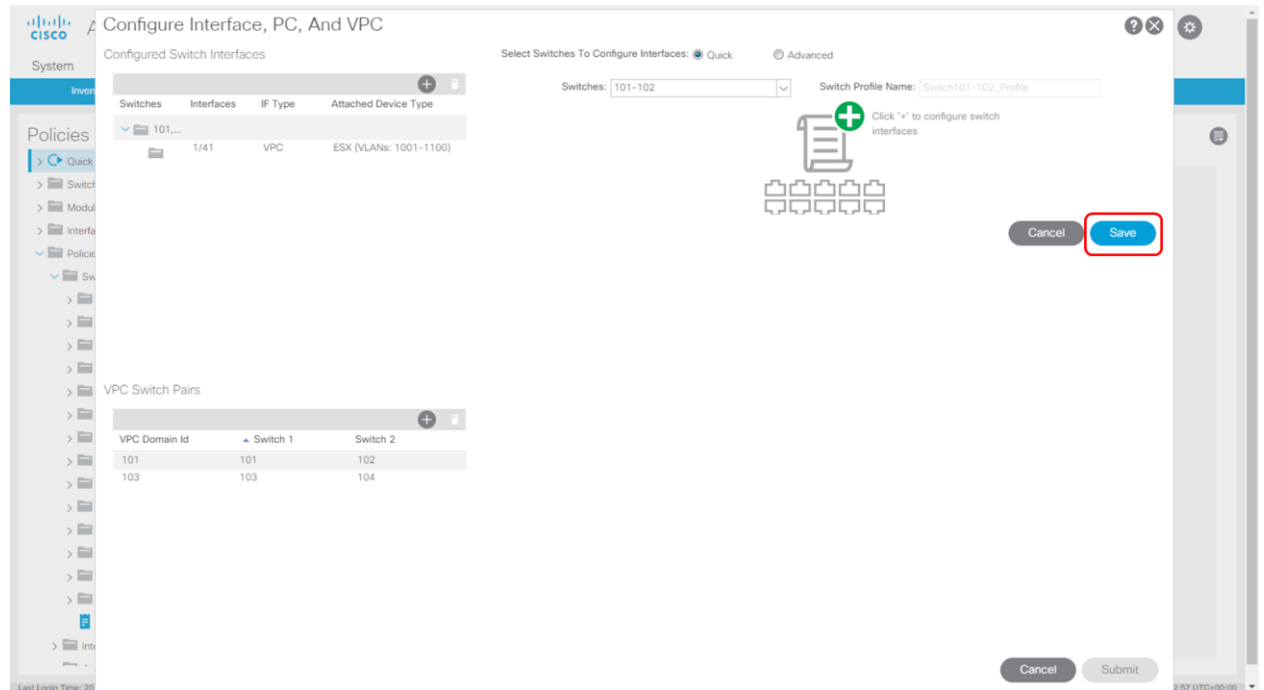
21. And last click "Submit"

Figure 82. Connecting FI-A: Step 21
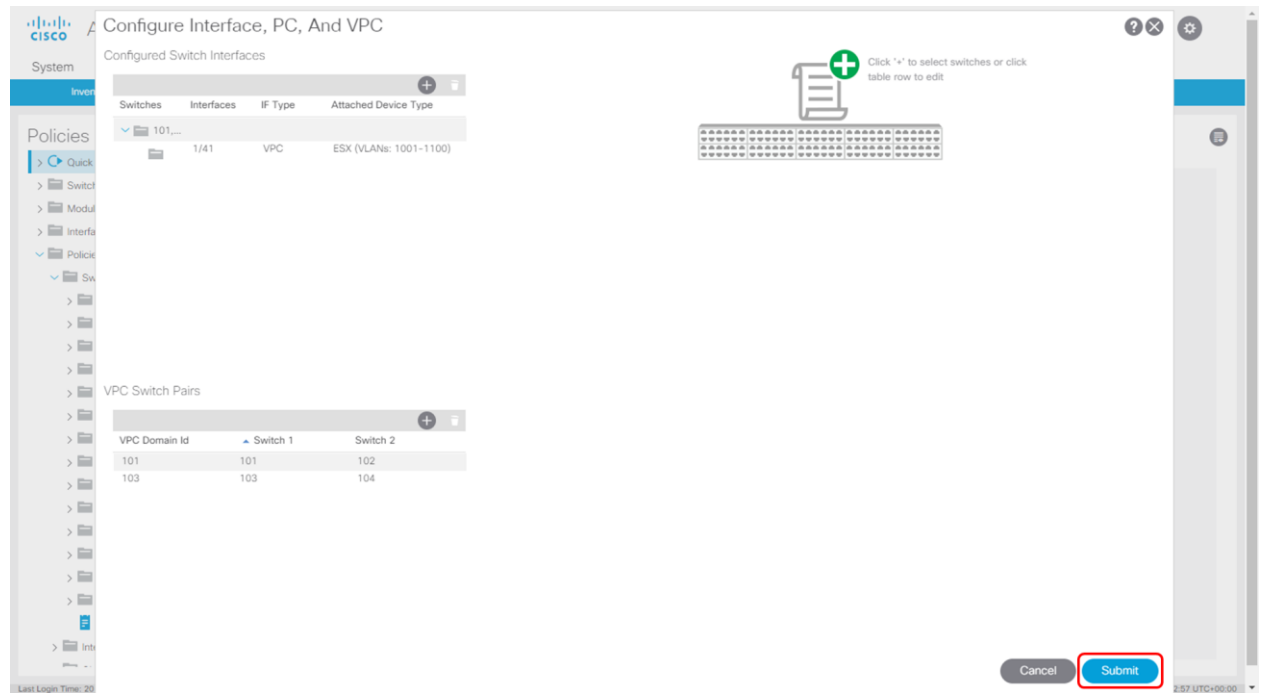
22. You will now notice if you relaunch the wizard all the configuration has been created:

Figure 83. Connecting FI-A: Step 22
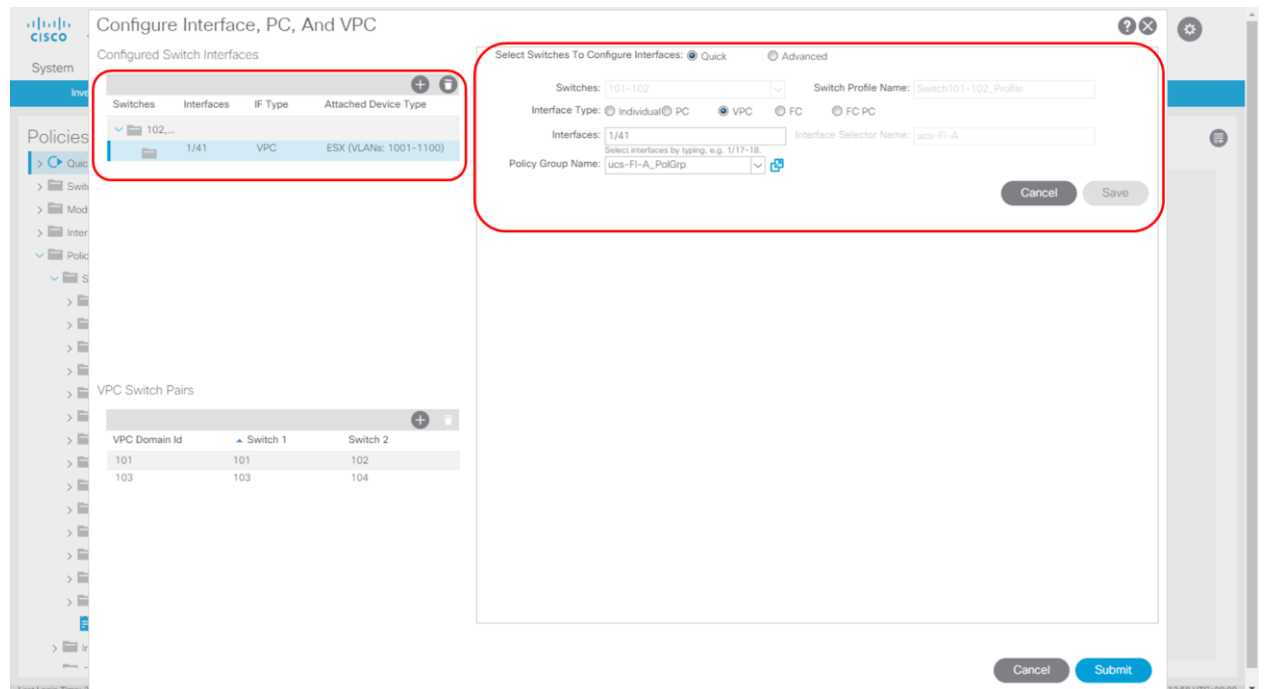
Make sure to select the create "Interface Profile" on the left which will provide the information you can see on the right.

23. In order to match the above configuration make sure to to add in UCM the following config:

■ VLAN 1001-1100 in the ACI connected vNIC's trough a vNIC template (make sure to use the vnic templates referenced for the NIC's in the service profile)

■ Enable LLDP in the Network Control Policy used in your vNIC template

Figure 84. Connecting FI-A: Step 23
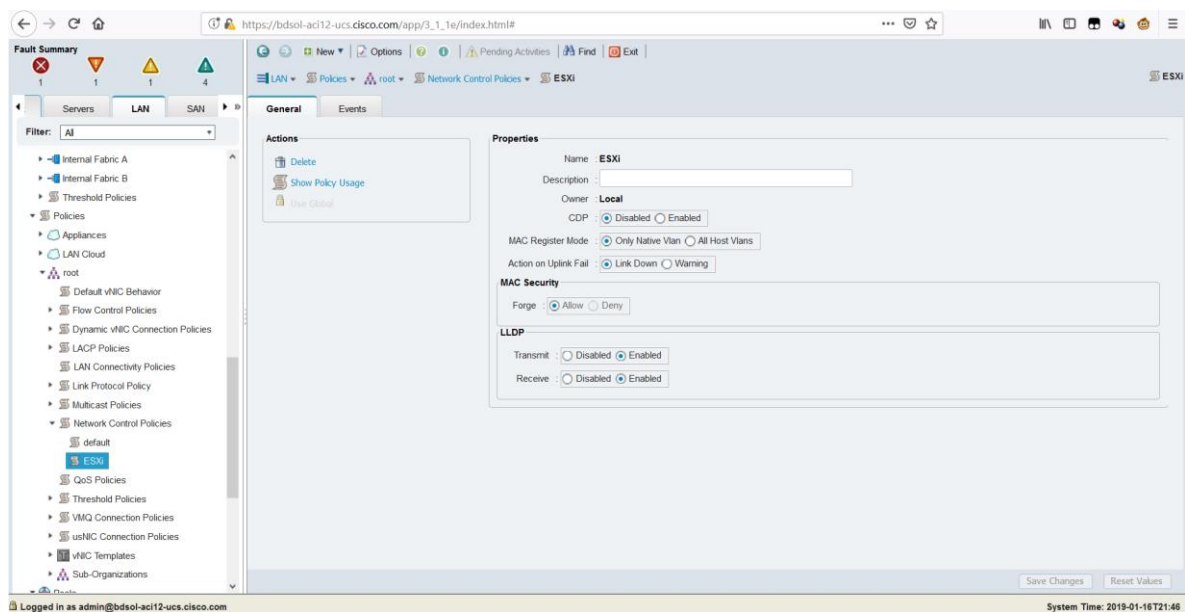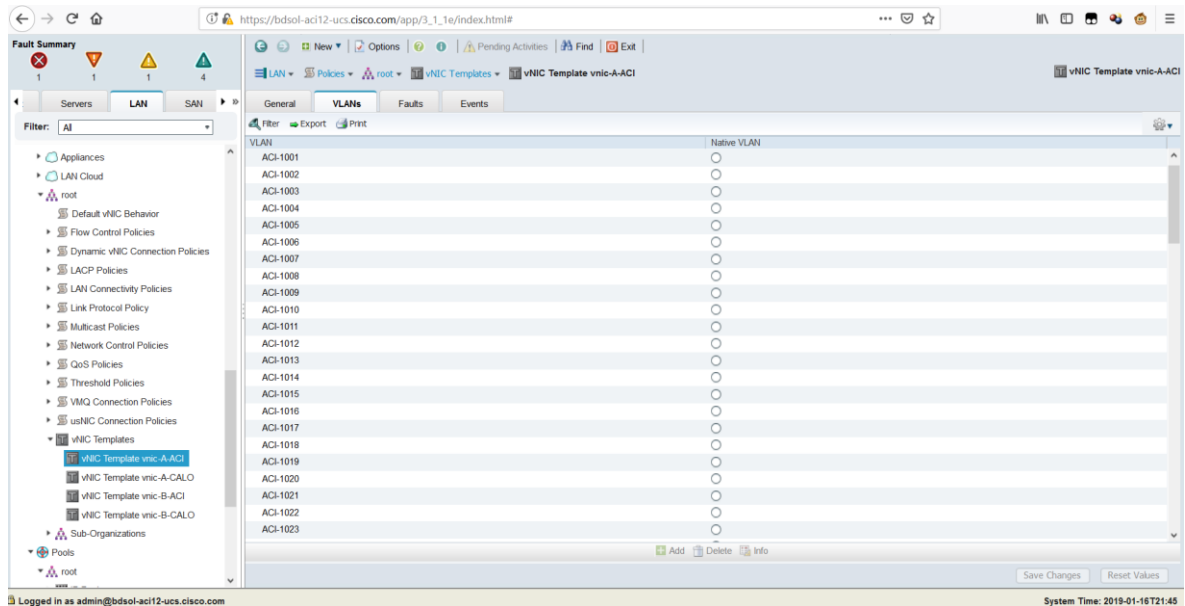
Connecting the Cecond UCS Fabric Interconnect (FI-B)

Now we will connect the second fabric interconnect. Because we already created some policies, you will now notice we will reuse a lot of them.

1. In order todo so we will start the wizard again under Fabric > Access Policies.

Figure 85. Connecting FI-B: Step 1
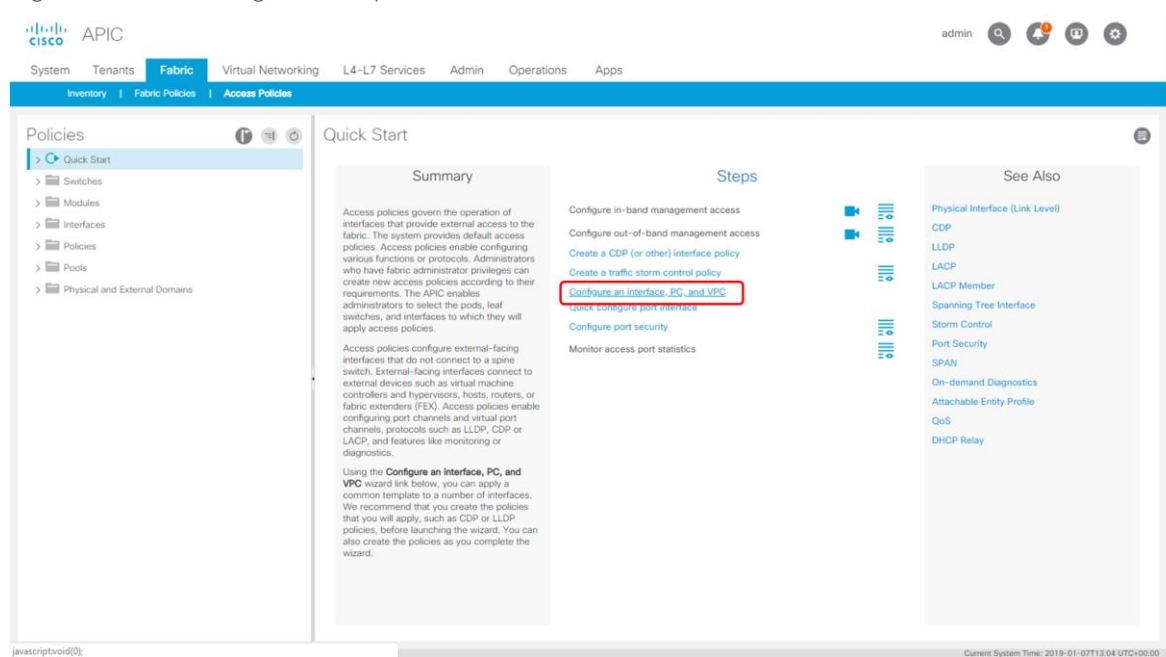
2. After launching the wizard, the following screen appears.

Figure 86. Connecting FI-B: Step 2
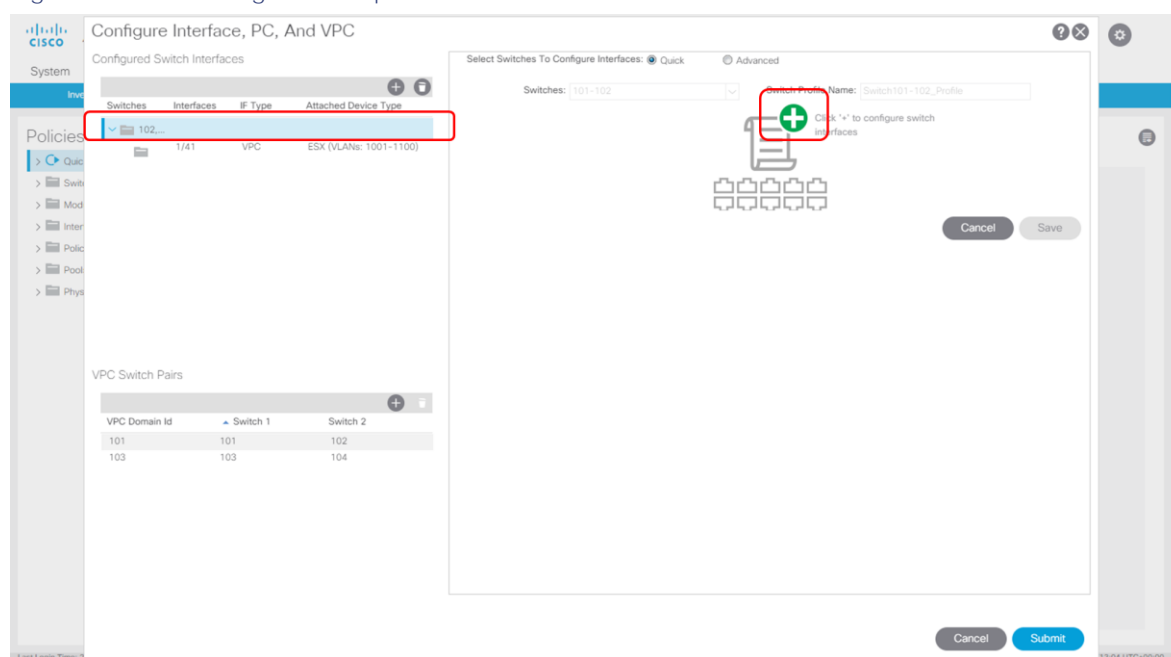
3. On the left select the switch pair (the "Switch Policy") and then on the right click "+". The following screen appears:

Figure 87. Connecting FI-B: Step 3
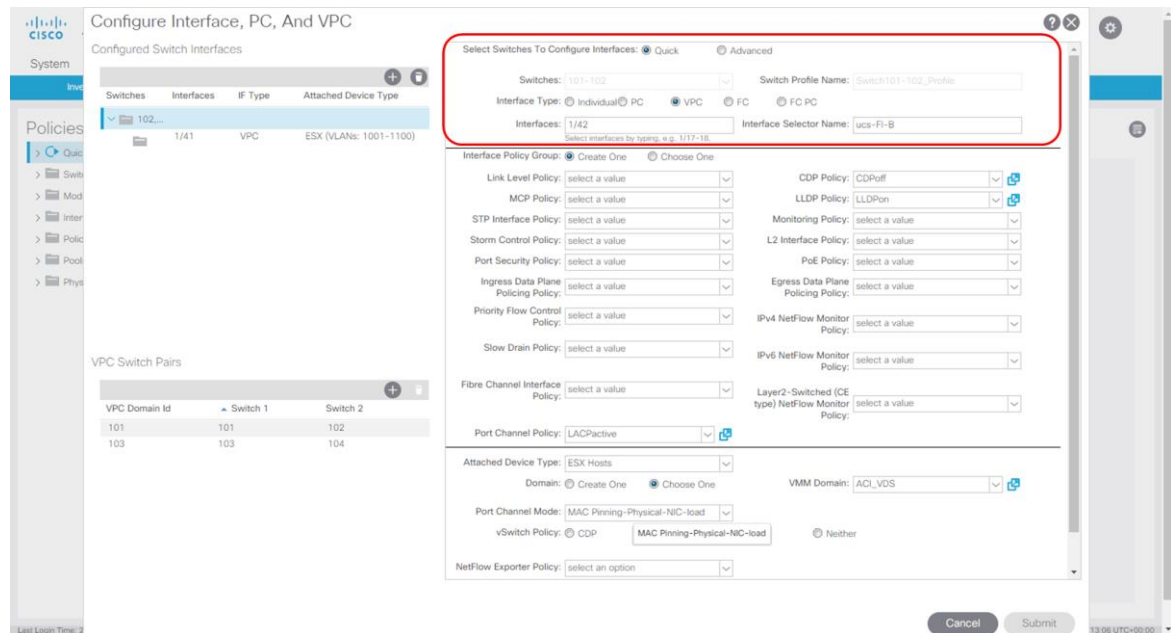
4. Fill in the correct interface under "Interfaces" and give the "Interface Policy" and new, in our case we will choose ucs-FI-B. In a next step we will define our "Interface Policy Group" and create a new one.

Figure 88. Connecting FI-B: Step 4
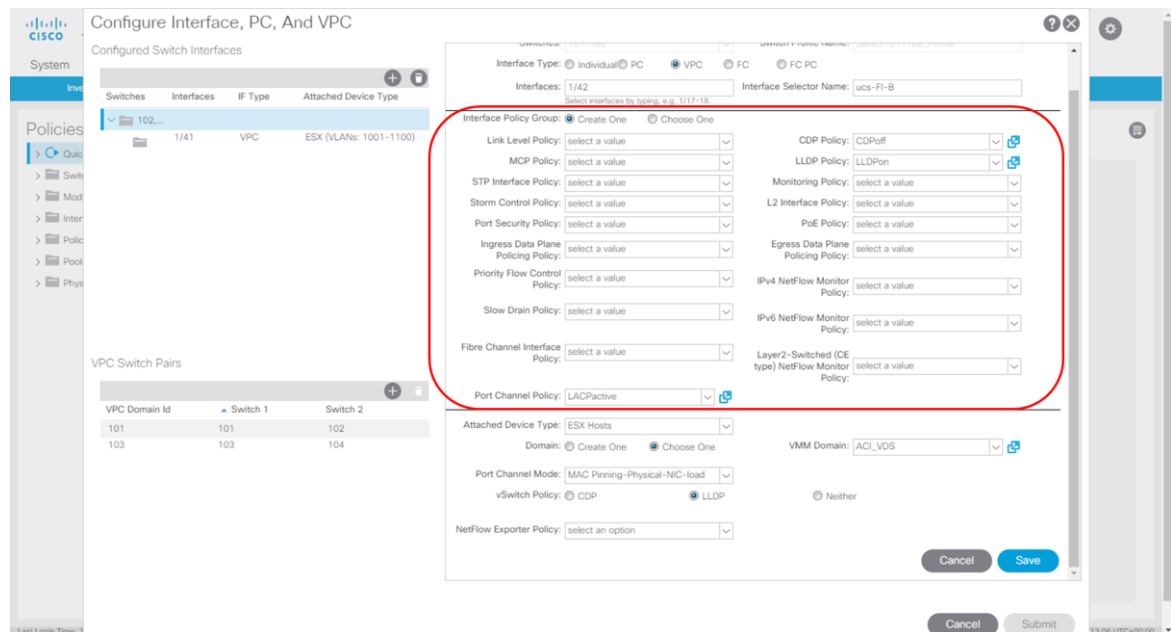
5. Notice we are reusing the earlier created CDPoff, LLDPon and LACPactive policies and we move to the Domain section.

Figure 89. Connecting FI-B: Step 5
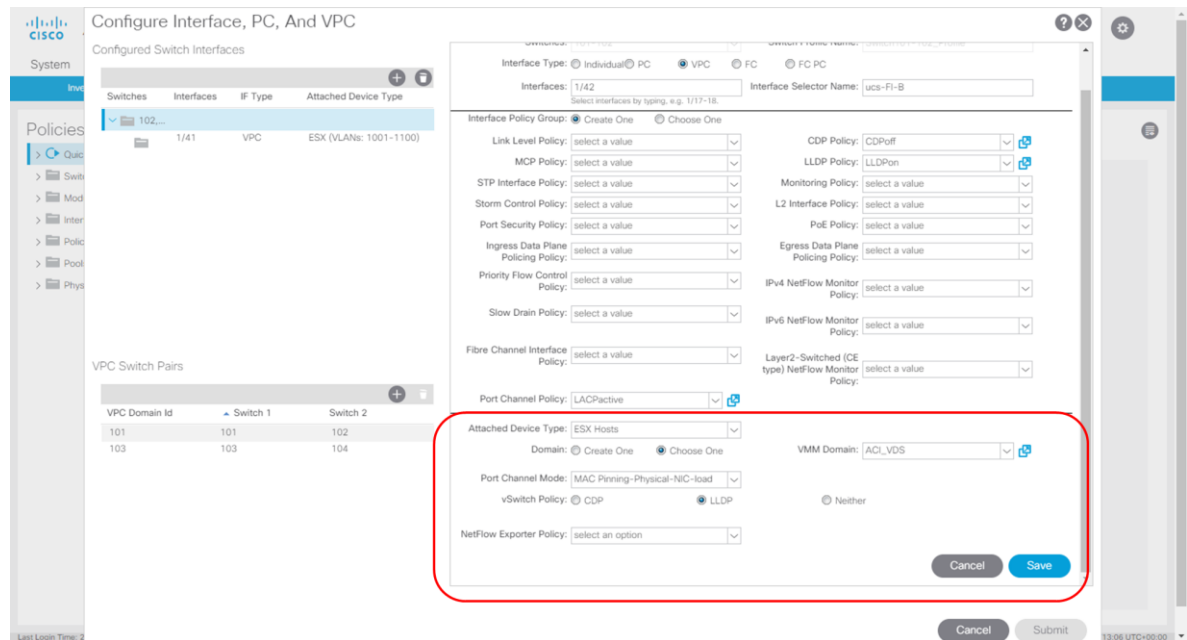
6. In the domain section we choose the existing VMM Domain named ACI_VDS and choose again "MAC Pinning-Physical-NIC-load" and LLDP.

    After doing this we select "Save".

Figure 90. Connecting FI-B: Step 6
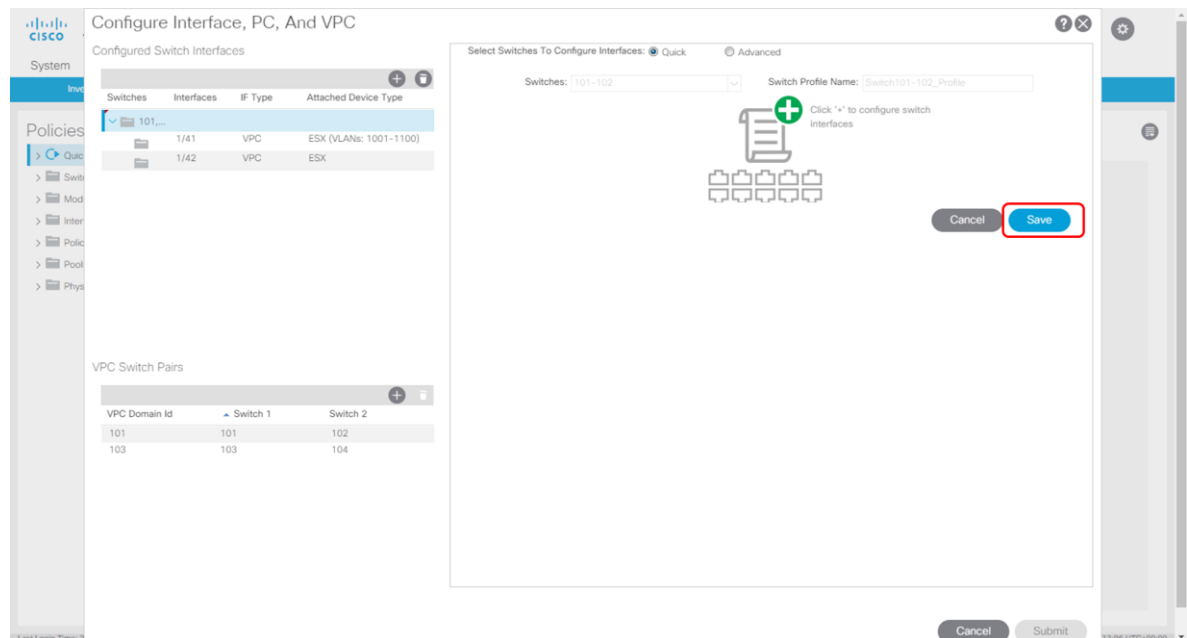
Then click "Save" again.

7. And finally click "Submit".

Figure 91. Connecting FI-B: Step 7
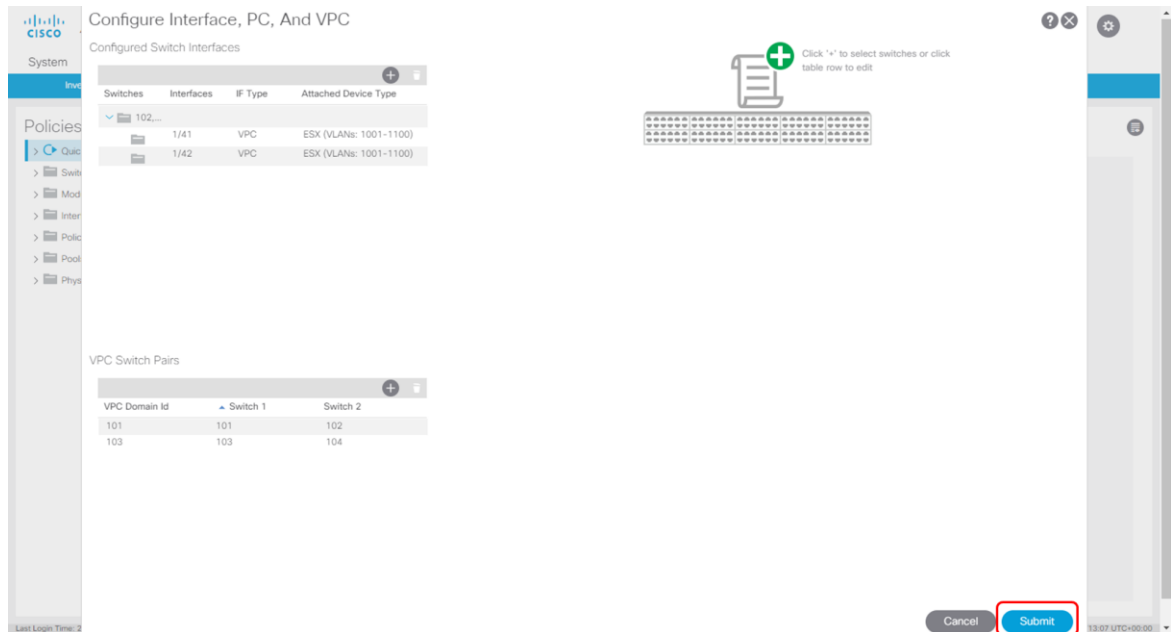
8. Although from the ACI perspective everything has been created, in VMware we still need to add the servers to the created vSwitch

Figure 92. Connecting FI-B: Step 8
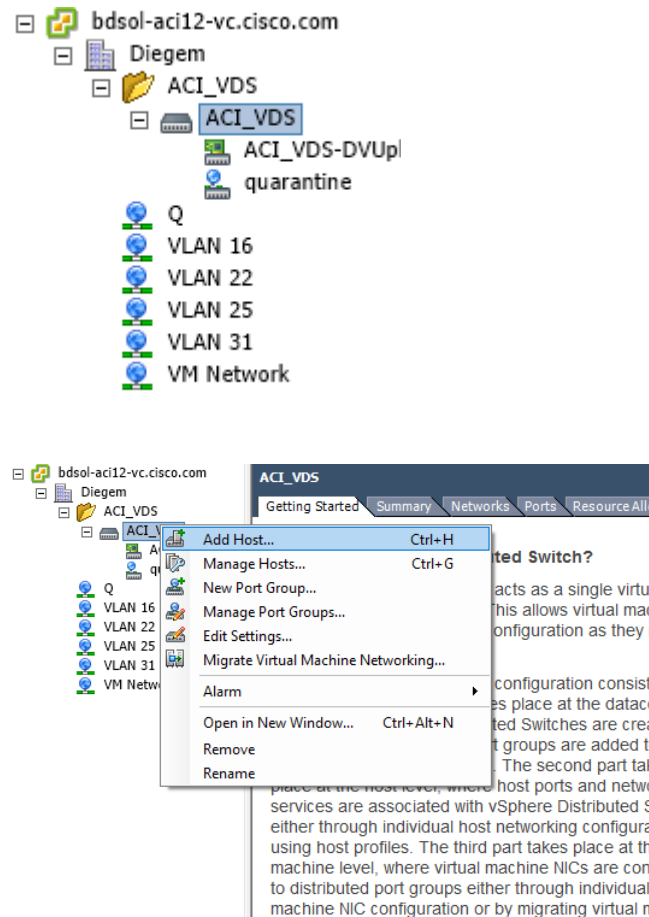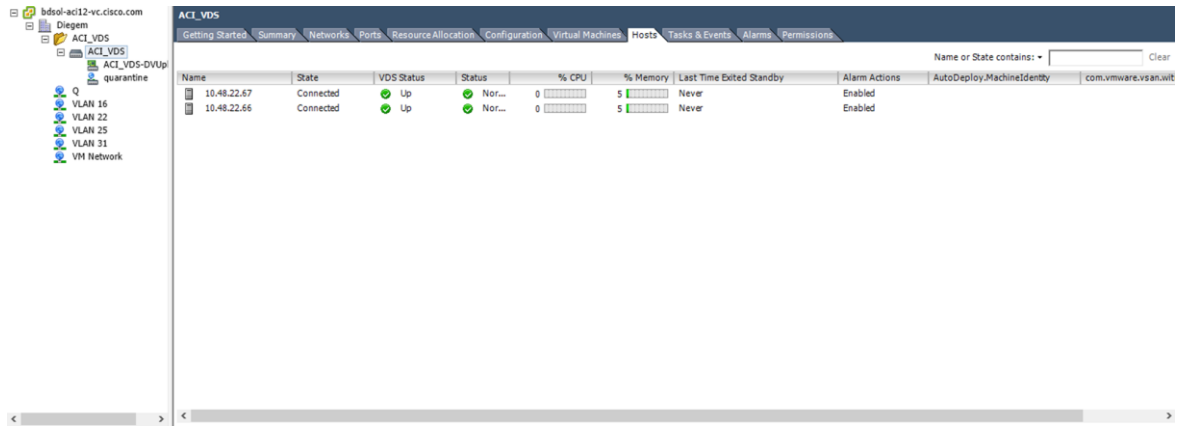
9. You will now notice under Virtual Networking Inventory that the servers have been discovered by ACI and that the connectivity towards the leaves is active.

Figure 93. Connecting FI-B: Step 9
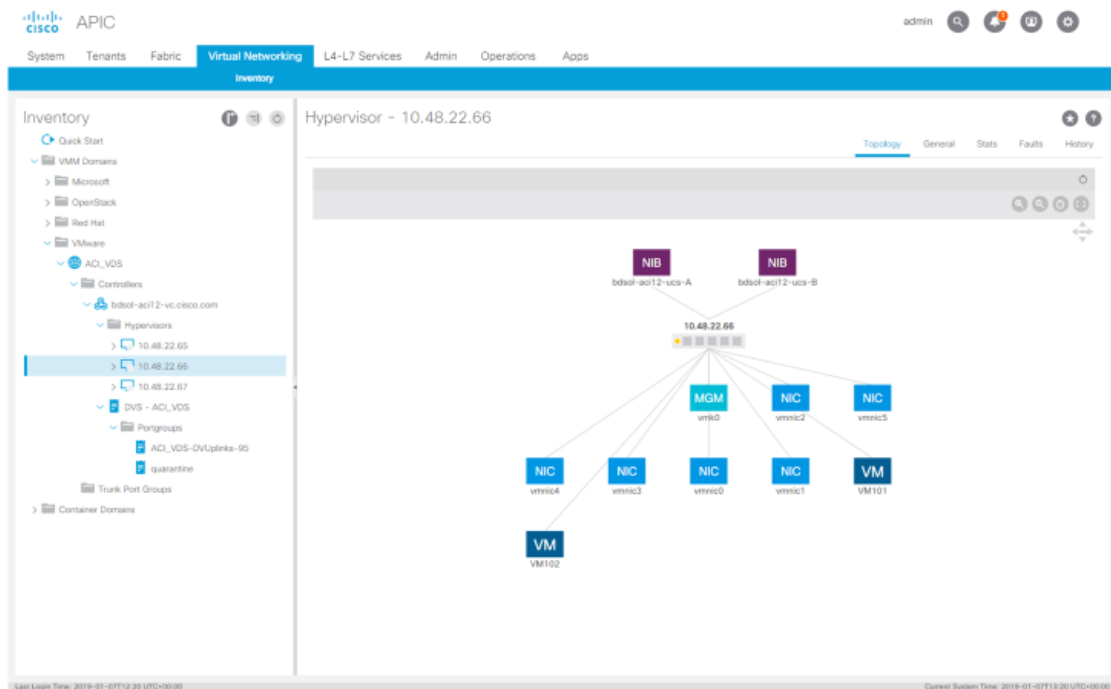
Overview of Created Policies

So, what happened when we have been executing this wizard? The following image shows all policies that have been created.

Cisco Systems. All printed copies and duplicate soft copies are considered uncontrolled
and the original online version should be referred to for the latest version.

Page 66 of 215

Figure 94. Connecting FI: Policy overview
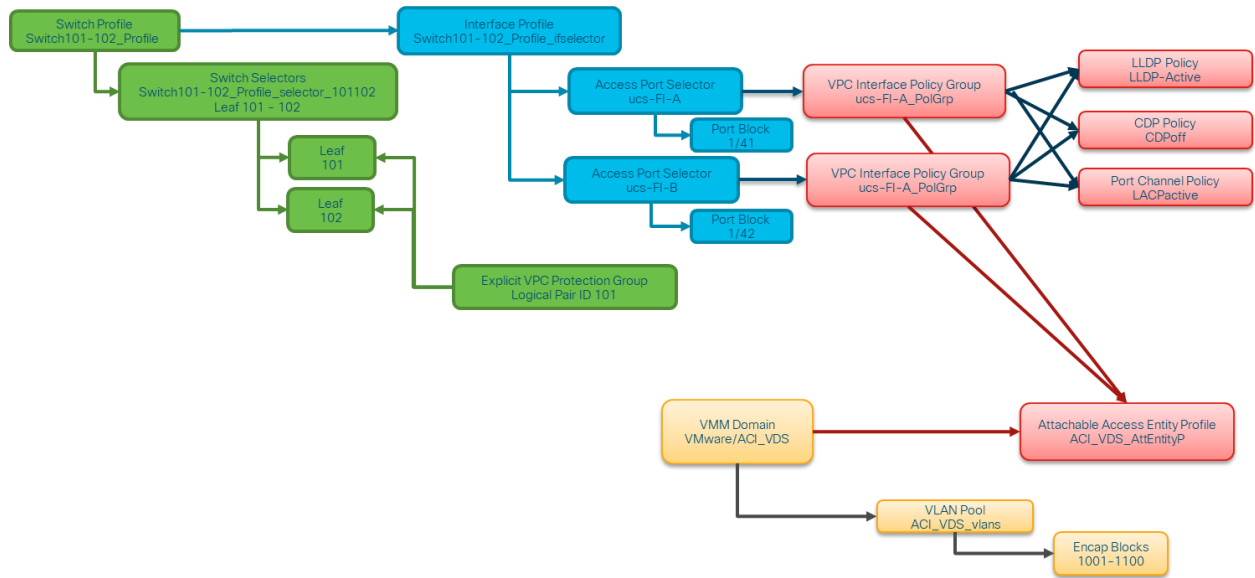
This means we now have a VMM domain that is fully equipped with access policies which can be linked to EPG's which we will later create.

The following is an overview of the created policies in detail.

Figure 95. Connecting FI: Explicit VPC Protection group

Cisco Systems. All printed copies and duplicate soft copies are considered uncontrolled and the original online version should be referred to for the latest version.

Page 67 of 215

Figure 96. Connecting FI: Switch Policy
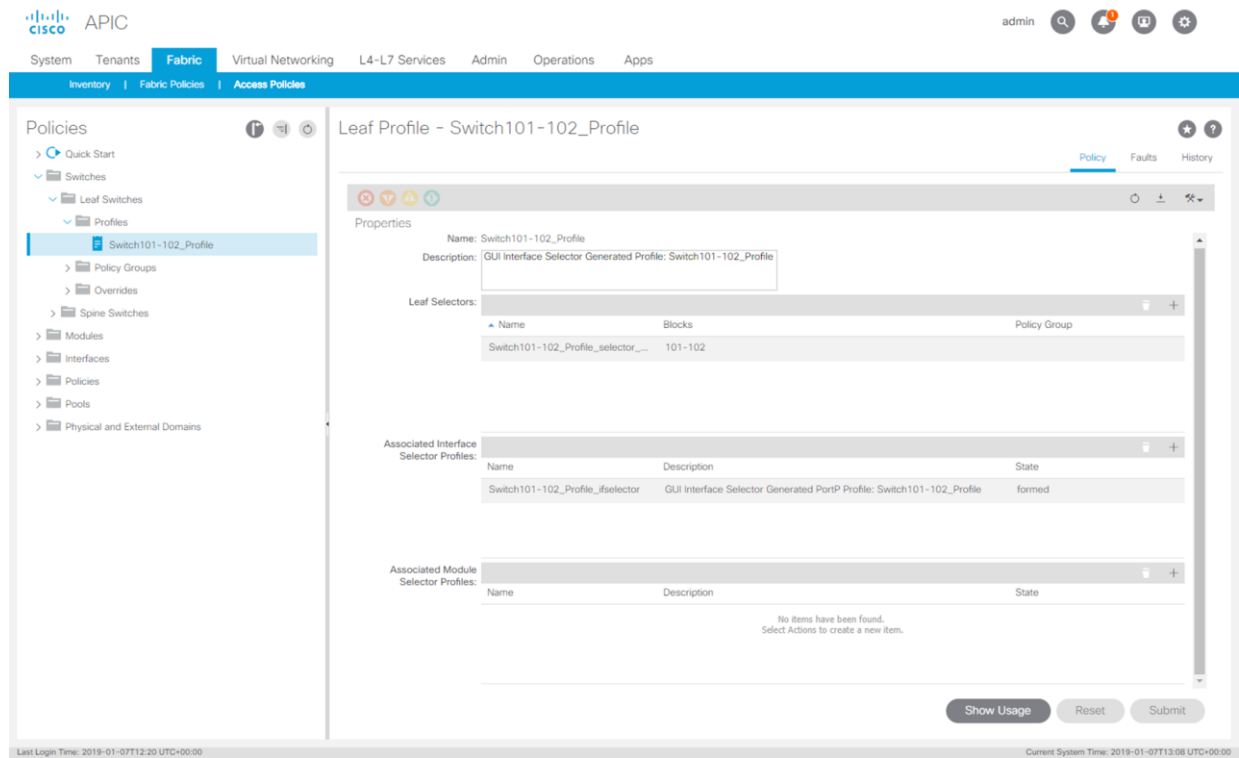
Figure 97. Connecting FI: Interface Profile
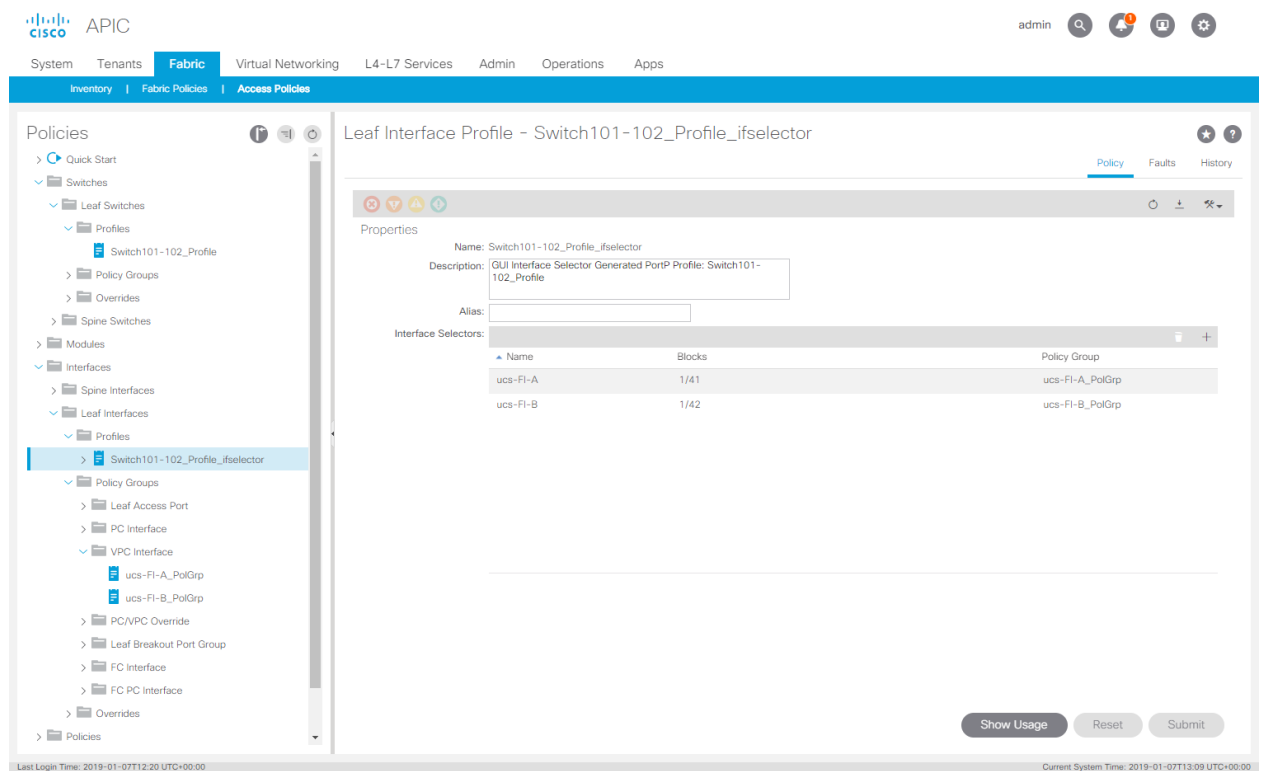
Figure 98. Connecting FI: Access port selectors

Cisco Systems. All printed copies and duplicate soft copies are considered uncontrolled
and the original online version should be referred to for the latest version.

Page 69 of 215

Figure 99. Connecting FI: VPC Interface Policy Group

Cisco Systems. All printed copies and duplicate soft copies are considered uncontrolled
and the original online version should be referred to for the latest version.

Page 70 of 215

Figure 100. Connecting FI: CDP Interface Policy

Figure 101. Connecting FI: LLDP Interface Policy

Figure 102. Connecting FI: Port Channel Policy

Figure 103. Connecting FI: Attachable Access Entity Profile

Figure 104. Connecting FI: VLAN Pool

Figure 105. Connecting FI: VMM Domain

Bare Metal Connectivity with Existing VMM Domain

The following sections describe bare metal connectivity with VMM domains.

Connecting the Bare Metal VMM Server

In this section we will connect server2 to the fabric and prepare it for VMM workload. The connectivity is as per the below diagram.

Figure 106. Bare metal with VMM: connectivity

Cisco Systems. All printed copies and duplicate soft copies are considered uncontrolled and the original online version should be referred to for the latest version.

Page 75 of 215

1. We will first launch the Access Policy wizard to configure this connectivity by going to Fabric > Access Policies > Quick Start

Figure 107. Bare metal with VMM: Step 1
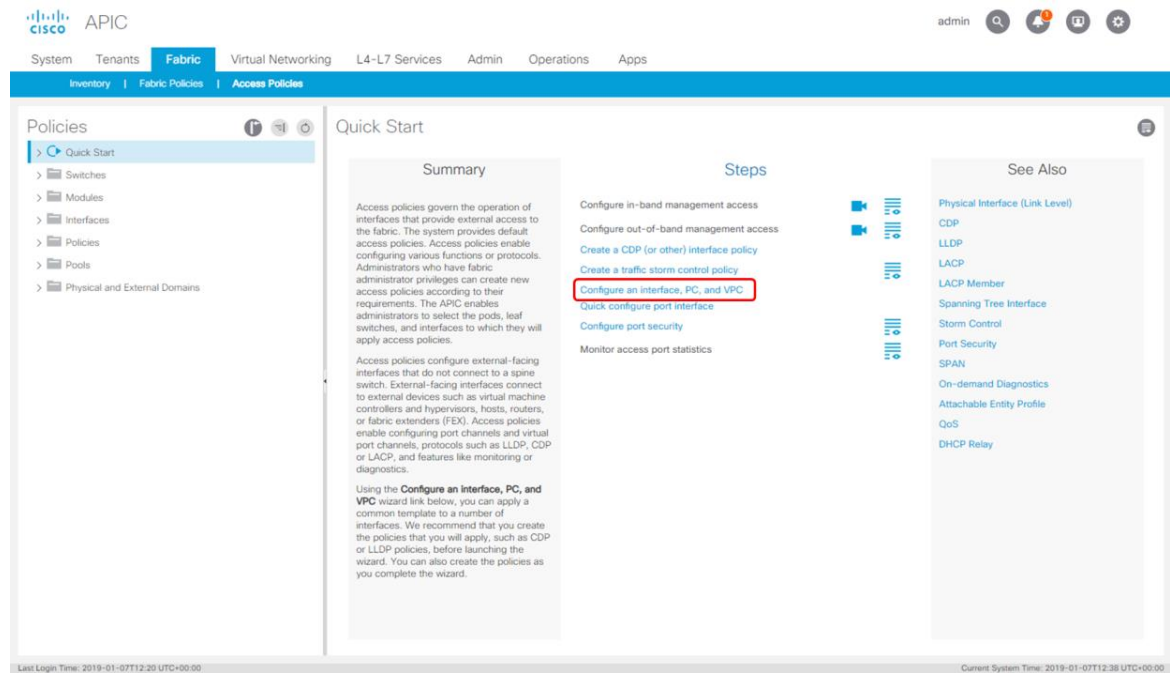
2. The following screen will appear. Select the 101-102 "Switch Policy" on the left and click on the right " + " to provide the access port details.

Figure 108. Bare metal with VMM: Step 2
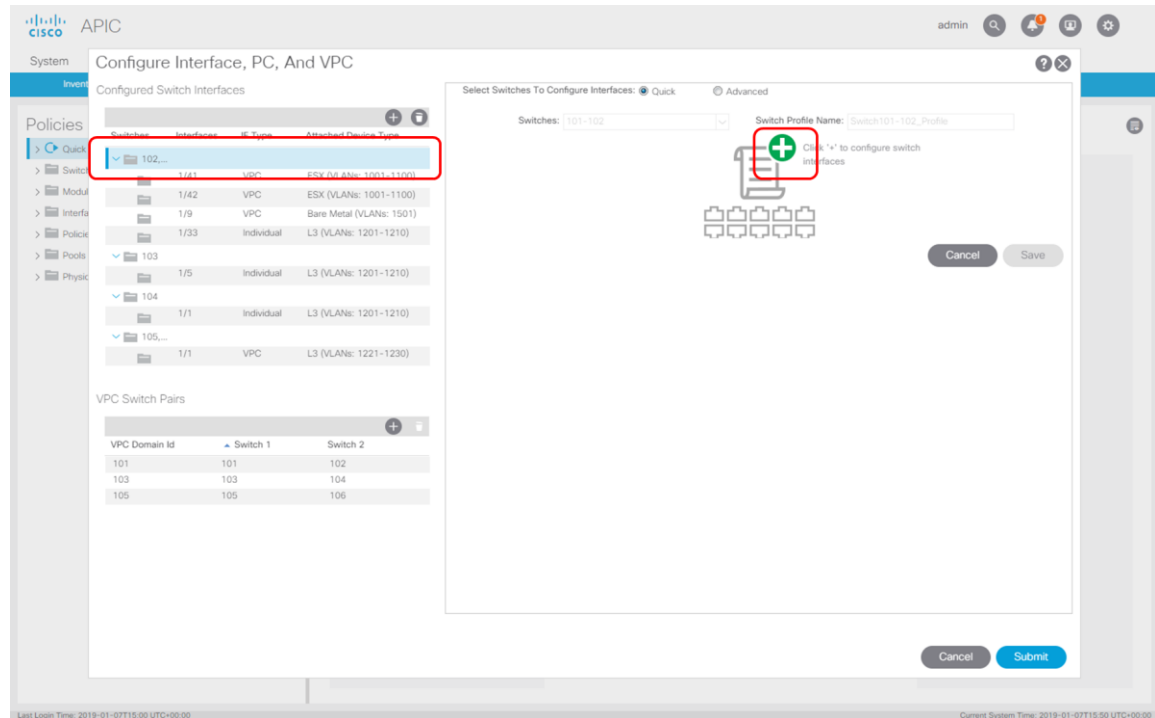
3. As per our diagram we will be using "MAC pinning" hence we need to use Individual links and make a normal "Access Policy".

Figure 109. Bare metal with VMM: Step 3
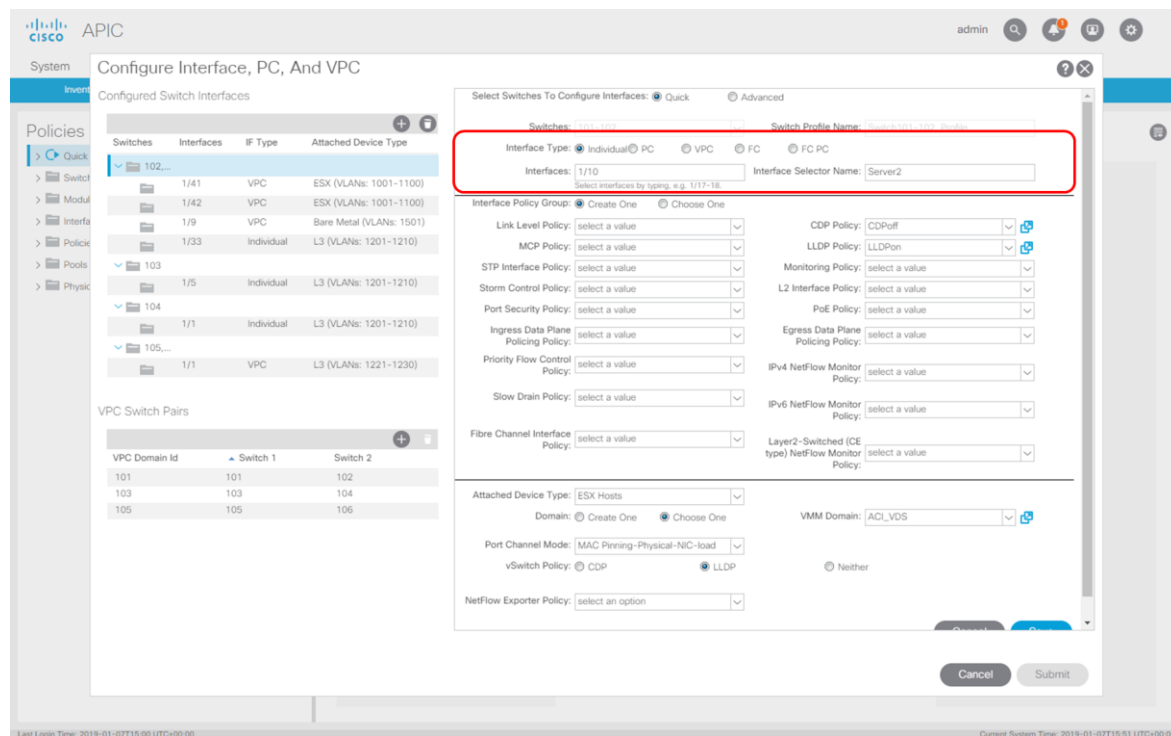
4. As in previous section we select CDPoff and LLDPon policy and leave all other policies to the defaults.

Figure 110. Bare metal with VMM: Step 4
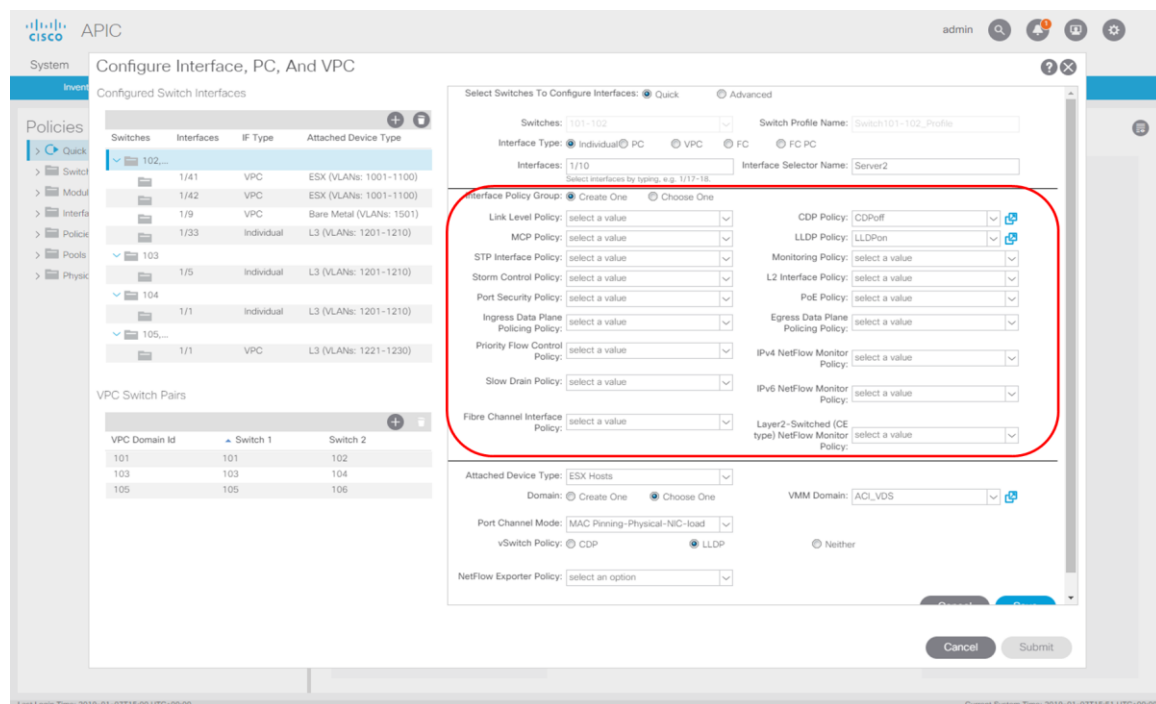
5. In the domain section we will use an existing "VMM Domain" and as a consequence of this we do not have to fill in the VMM VLAN Pool as this is automatically pickedup.

Notice we have to configure the "Port Channel Mode" and as our already previously connected Fabric Interconnects blades have been configure with "MAC Pinning-Physical-NIC-load" we will use a similar load balancing algorithm here. We will also enable LLDP to enable dynamic VMM learning. After having configured this we click "Save"

Figure 111. Bare metal with VMM: Step 5
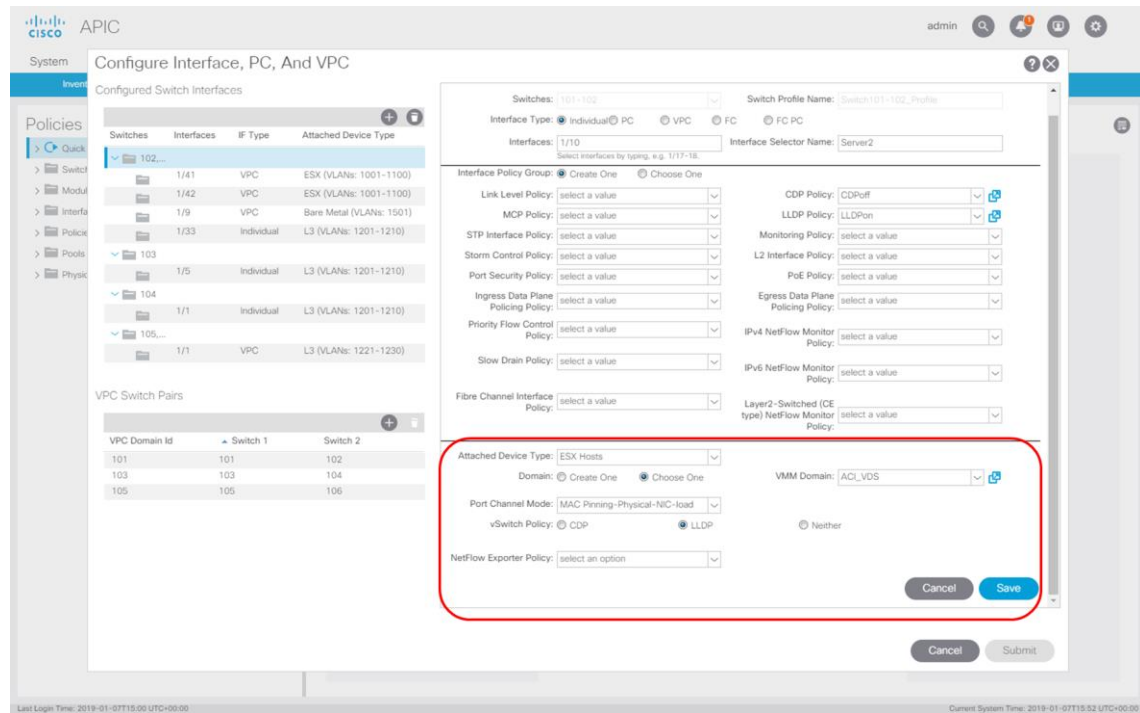
6. Again "Save"

Figure 112. Bare metal with VMM: Step 6
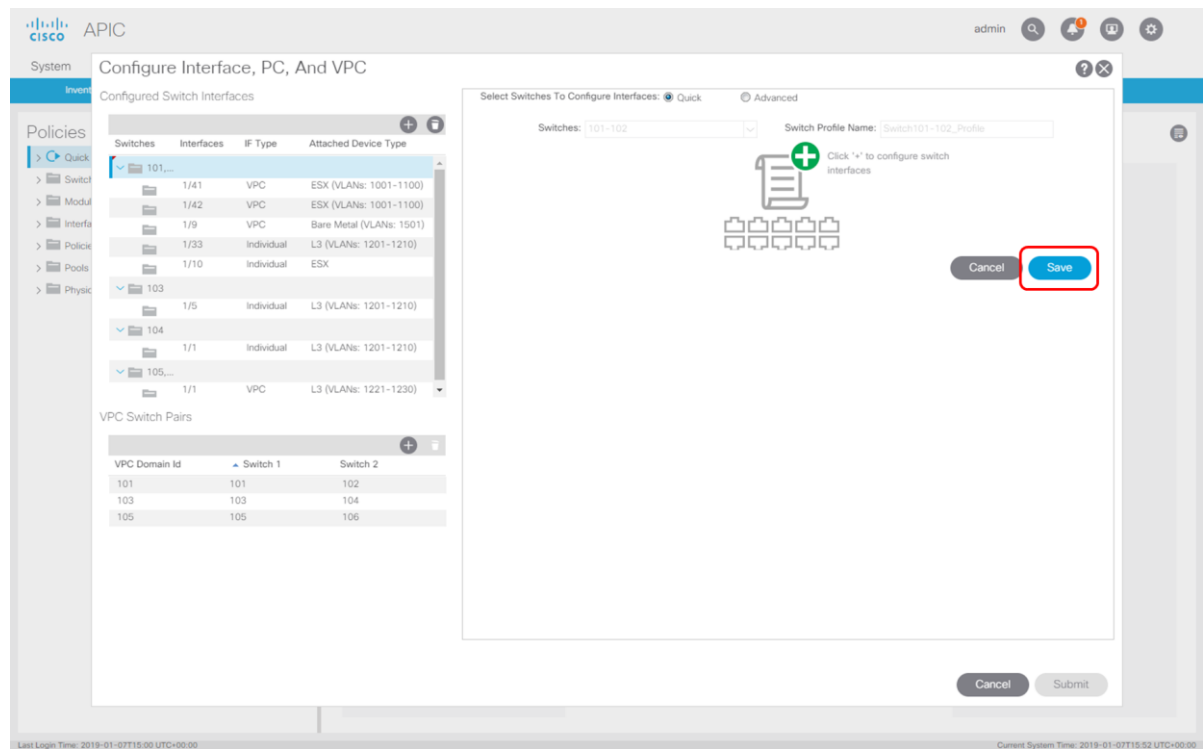
7. And finally we click "Submit"

Figure 113. Bare metal with VMM: Step 7
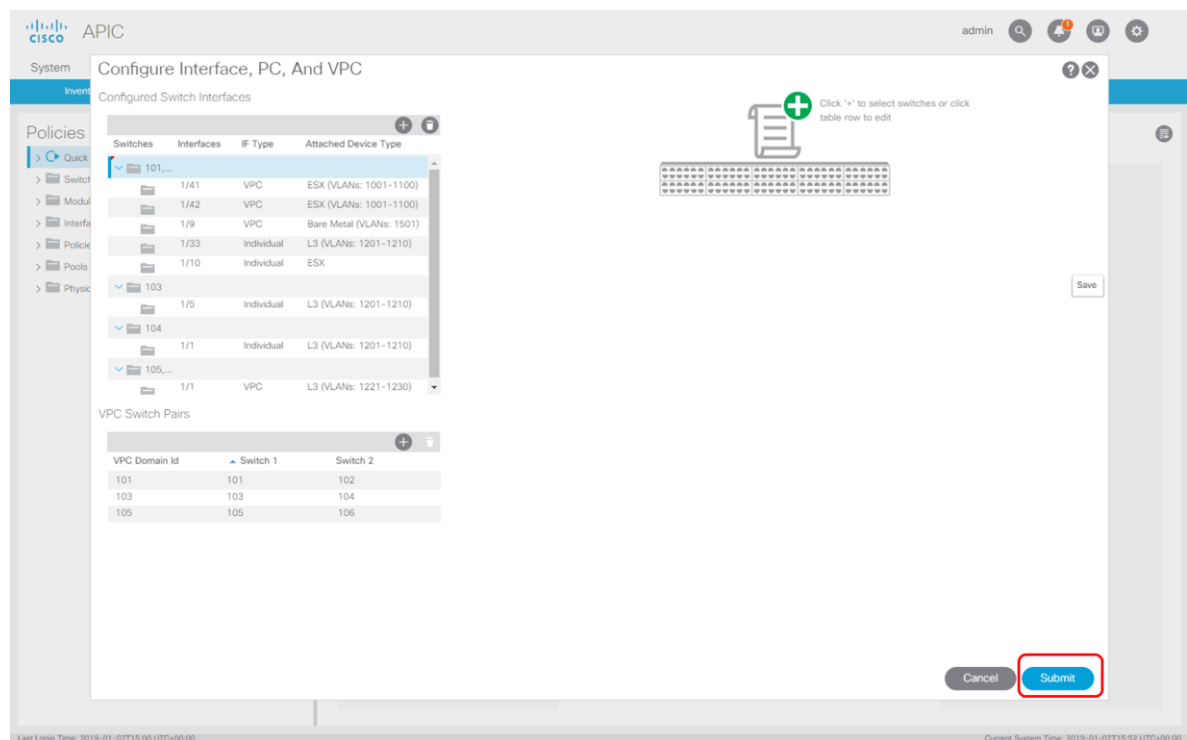
Overview of created policies

So, what happened when we have been executing this wizard? The following image shows all policies that have been created or linked to (already existing).

Figure 114. Bare metal with VMM: Piolicy overview
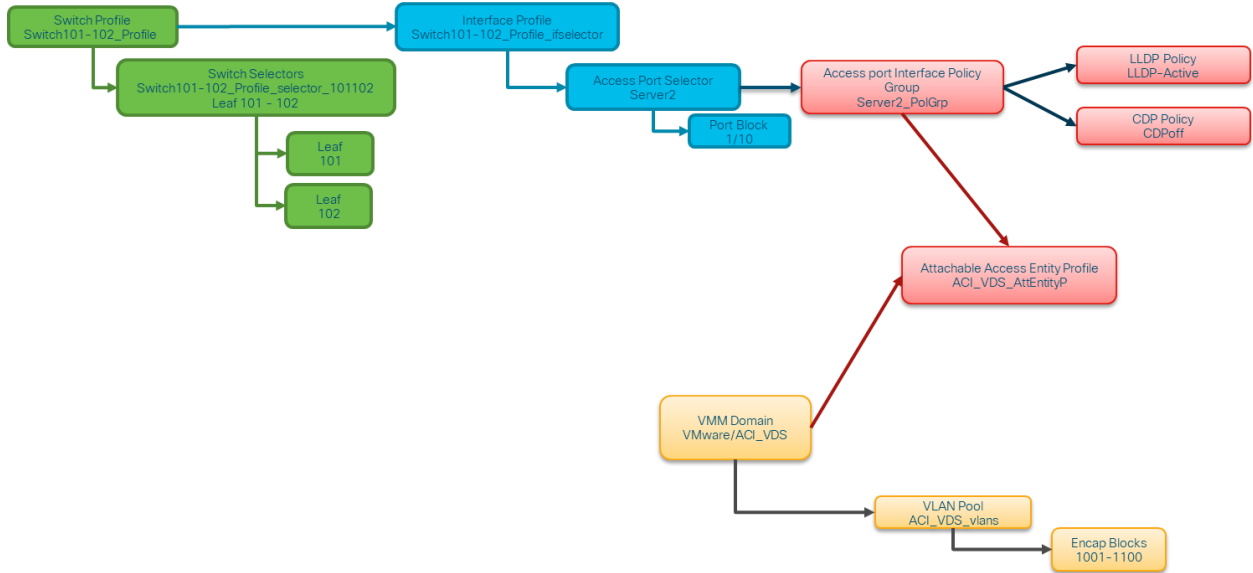
This means we now have all policies in place to connect our server "server2" to the existing VMM domain and later on we can use the "VMM Domain" to run virtual workload on our server.

The following is an overview of the created or re-used policies in detail.

Figure 115. Bare metal with VMM: Switch Policy

Figure 116. Bare metal with VMM: Interface Profile

Figure 117. Bare metal with VMM: Access port selectors

Figure 118. Bare metal with VMM: Access Port Policy Group

Cisco Systems. All printed copies and duplicate soft copies are considered uncontrolled
and the original online version should be referred to for the latest version.

Page 82 of 215

Figure 119. Bare metal with VMM: CDP Interface Policy
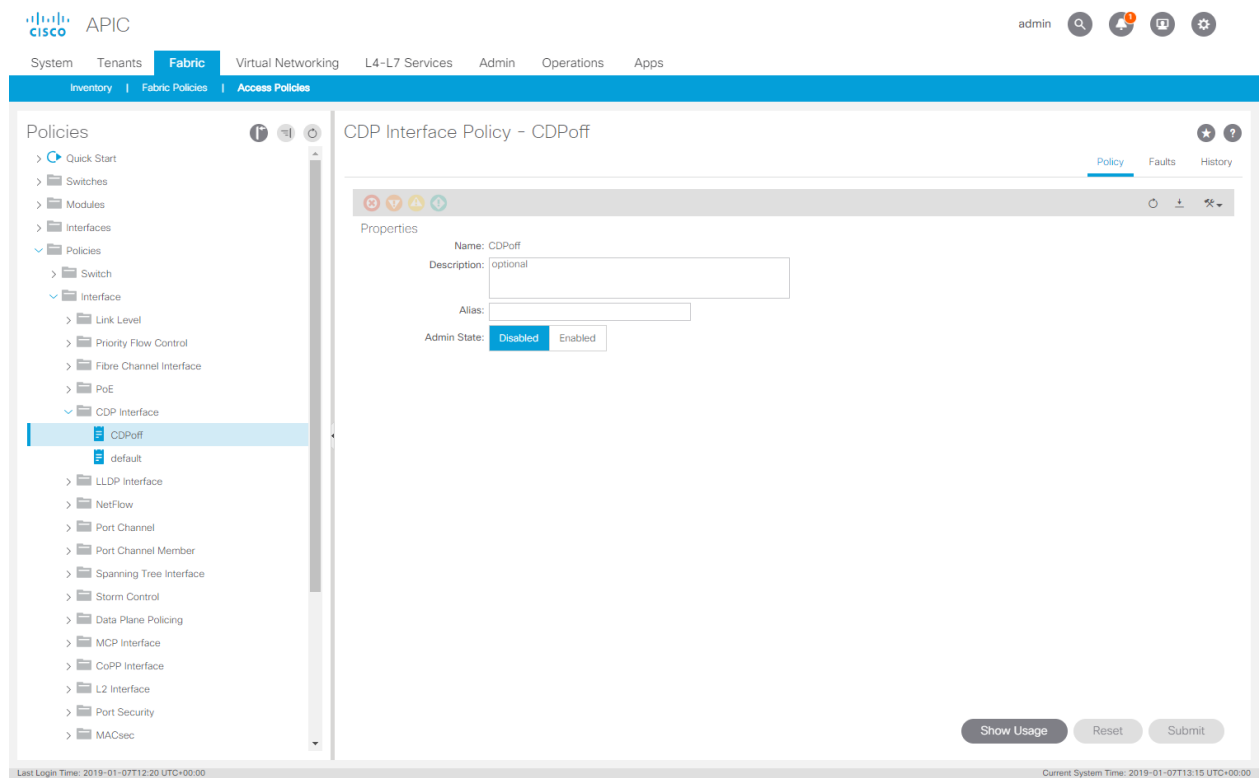
Figure 120. Bare metal with VMM: LLDP Interface Policy
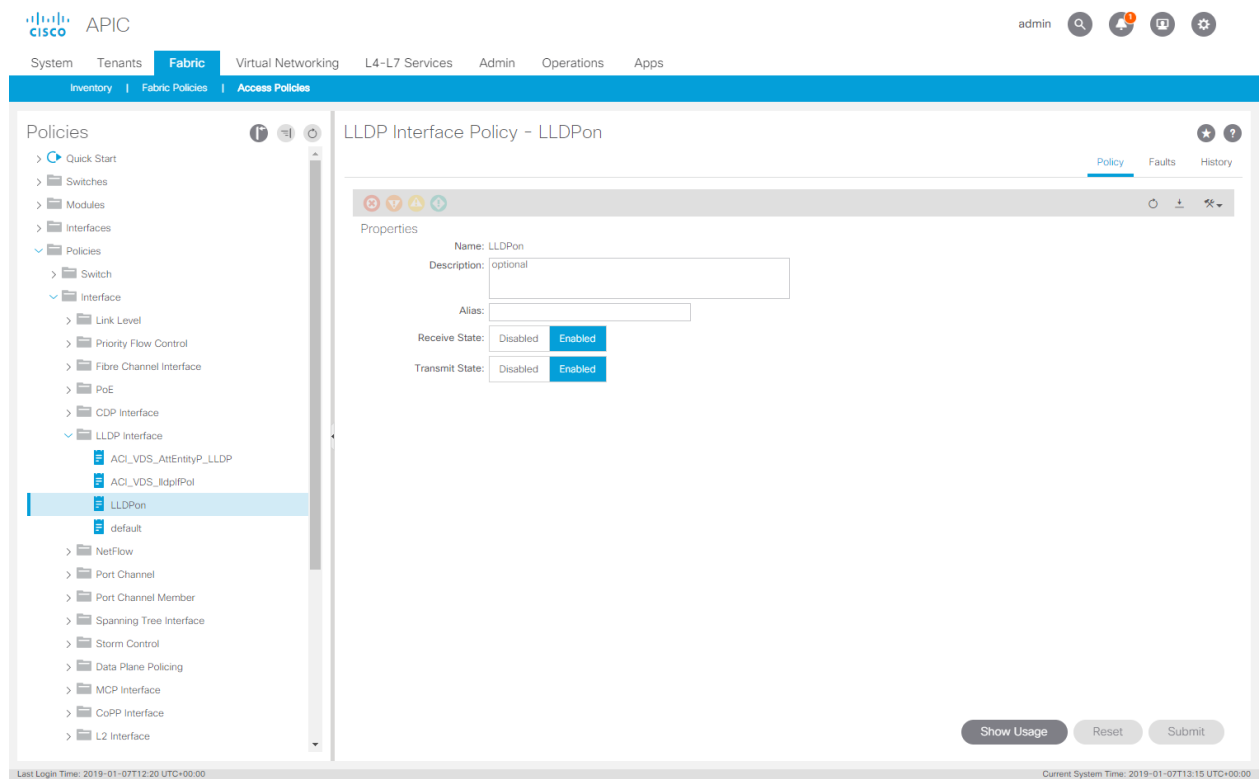
Figure 121. Bare metal with VMM: Attachable Access Entity Profile

Figure 122. Bare metal with VMM: VLAN Pool

Figure 123. Bare metal with VMM: VMM Domain
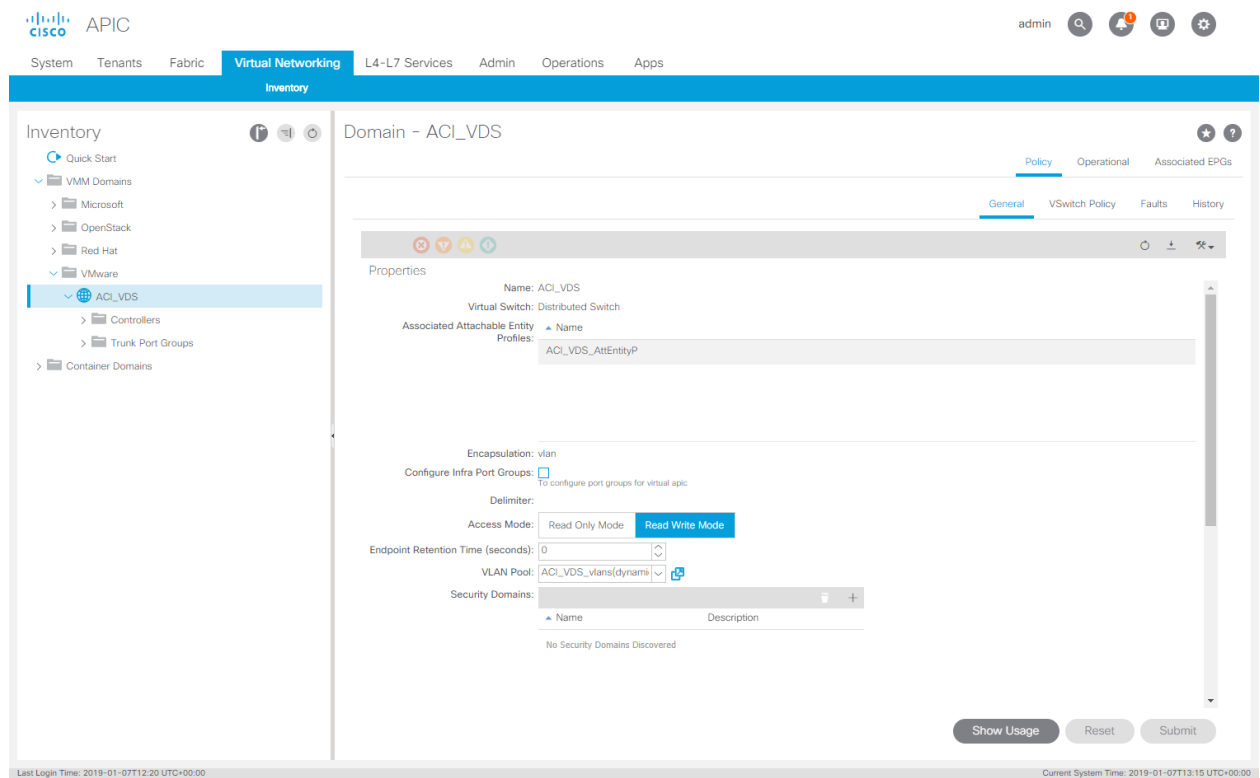
Bare Metal Connectivity with Physical Domain

The following sections describe bare metal connectivity with physical domain.

Connecting the Bare Metal

In this section we will connect server1 to the fabric and prepare it for bare metal workload. The connectivity is as per the below diagram.

Figure 124. Bare metal with physical domain: connectivity

1. We will first launch the Access Policy wizard to configure this connectivity by going to Fabric > Access Policies > Quick Start

Figure 125. Bare metal with physical domain: Step 1
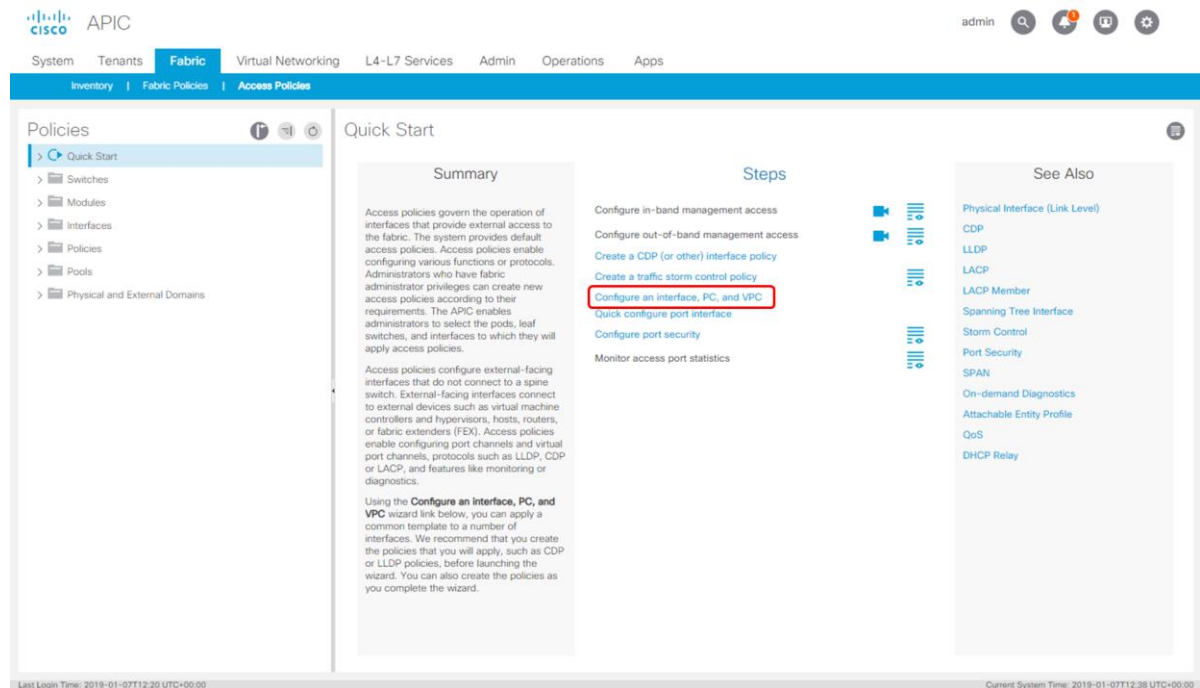
2. The following screen will appear. Select the correct "Switch Policy" on the left and click "+" on the right.

Figure 126. Bare metal with physical domain: Step 2
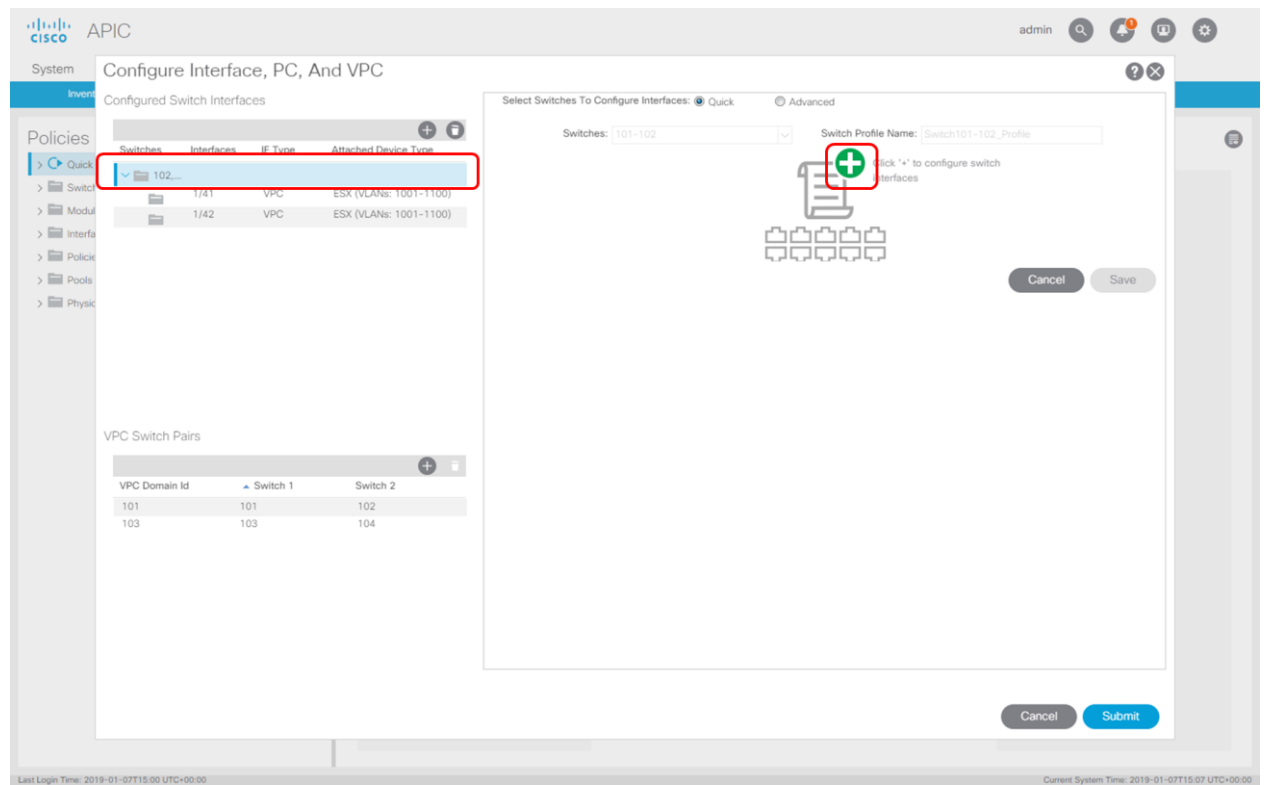
3. Select "VPC" Interface Type and fill in the "Interfaces" and "Interface Selector Name".

Figure 127. Bare metal with physical domain: Step 3
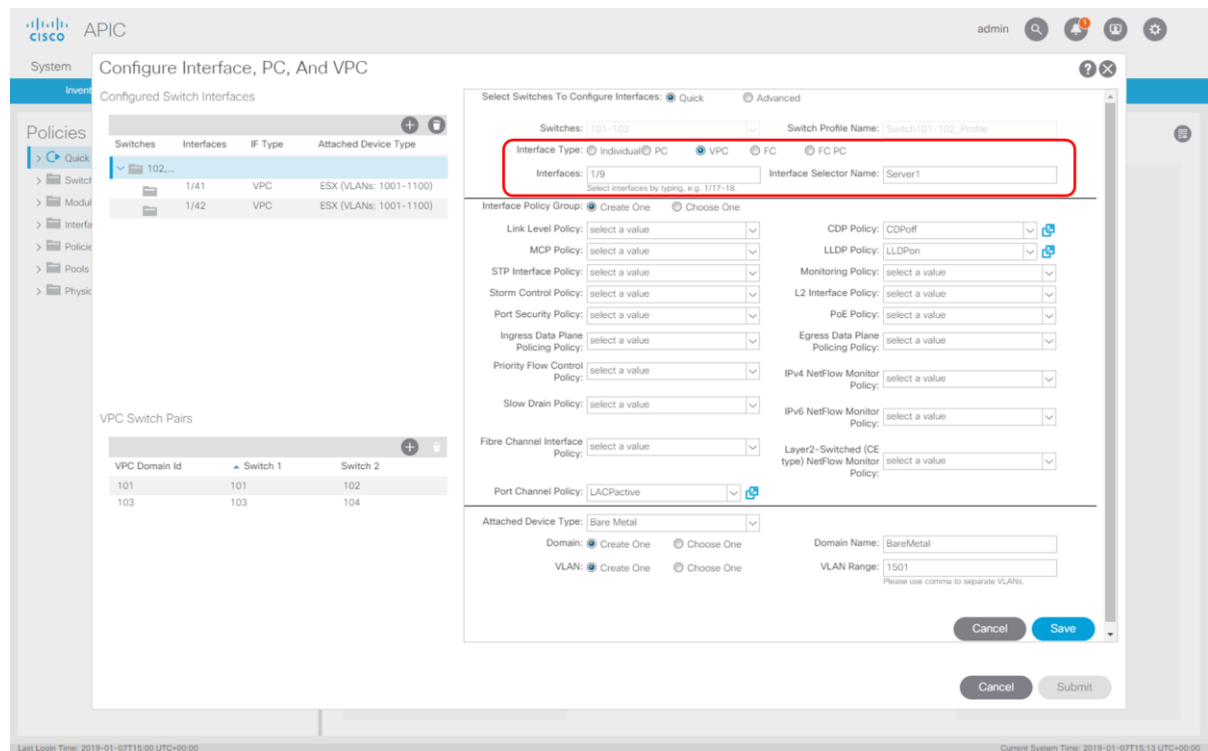
4. Select the existing CDP, LLDP and Port Channel Policy and move to the "Domain" section.

Figure 128. Bare metal with physical domain: Step 4
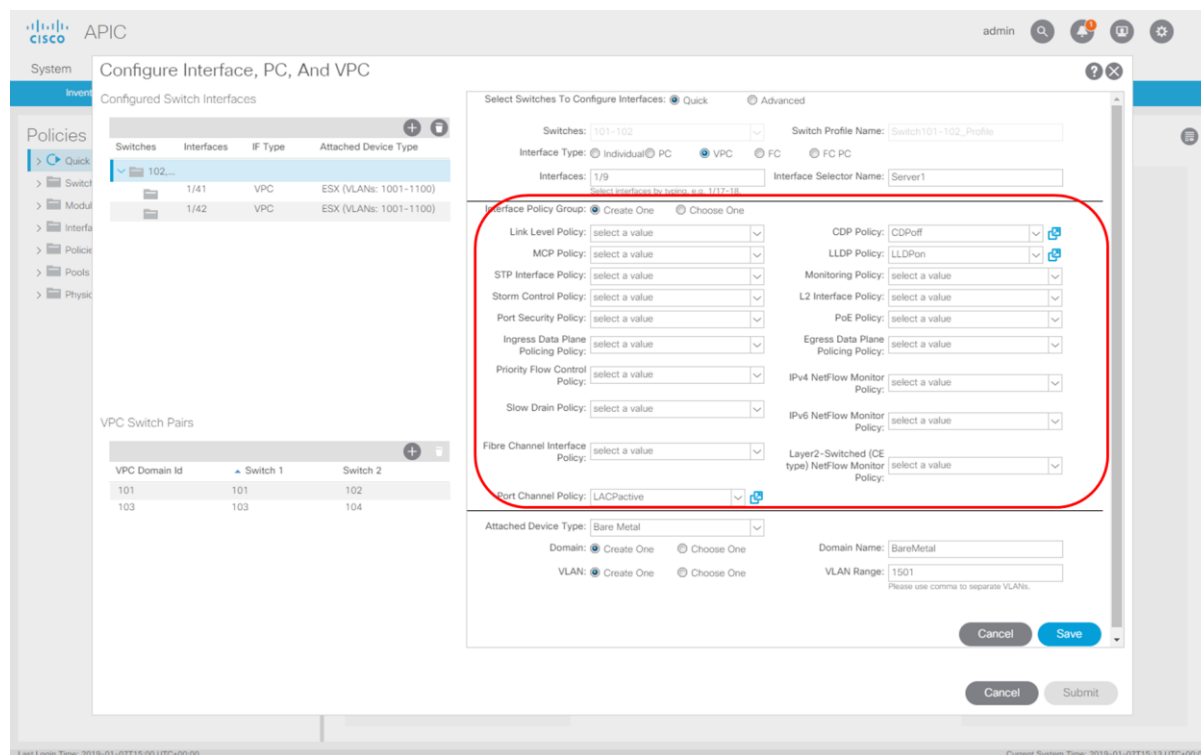
5. We will create a new "Physical Domain" called "BareMetal" and create a newly linked "VLAN Pool". After this click "Save"

Figure 129. Bare metal with physical domain: Step 5
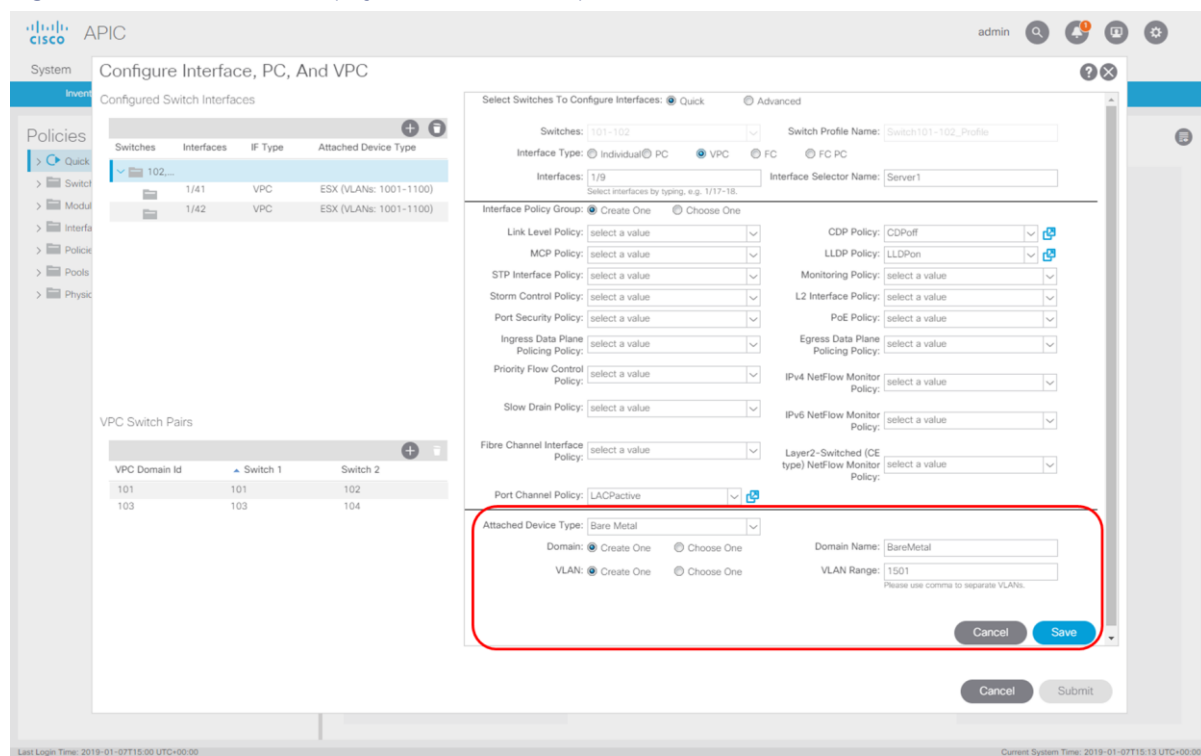
6. Click "Save"

Figure 130. Bare metal with physical domain: Step 6
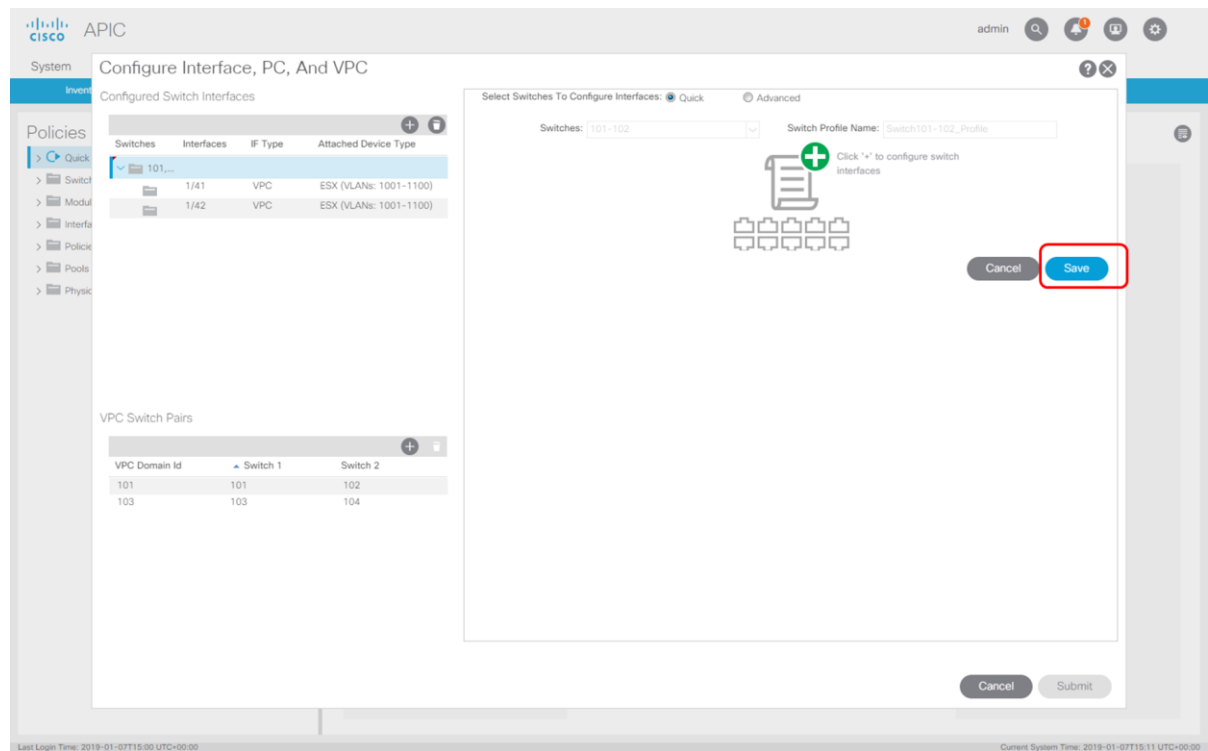
7. And click Submit.

Figure 131. Bare metal with physical domain: Step 7
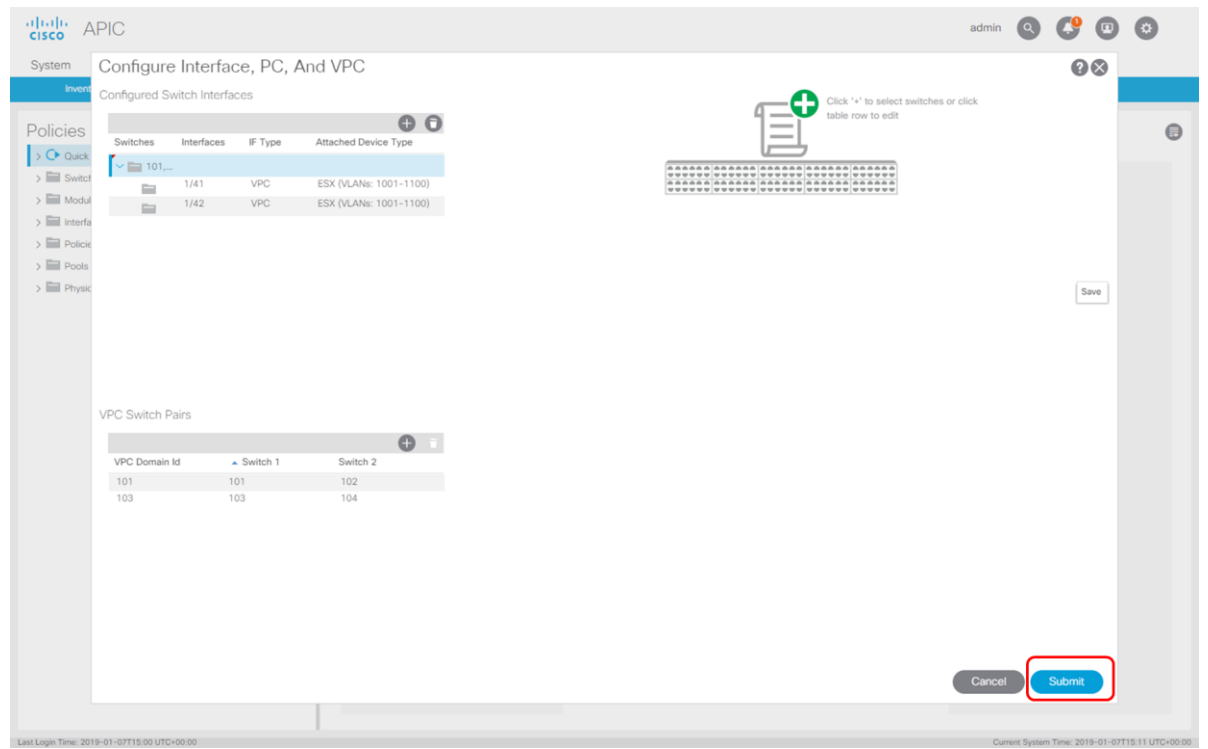
Overview of created policies

So, what happened when we have been executing this wizard? The following image shows all policies that have been created or linked to (already existing).

Figure 132. Bare metal with physical domain: Policy overview
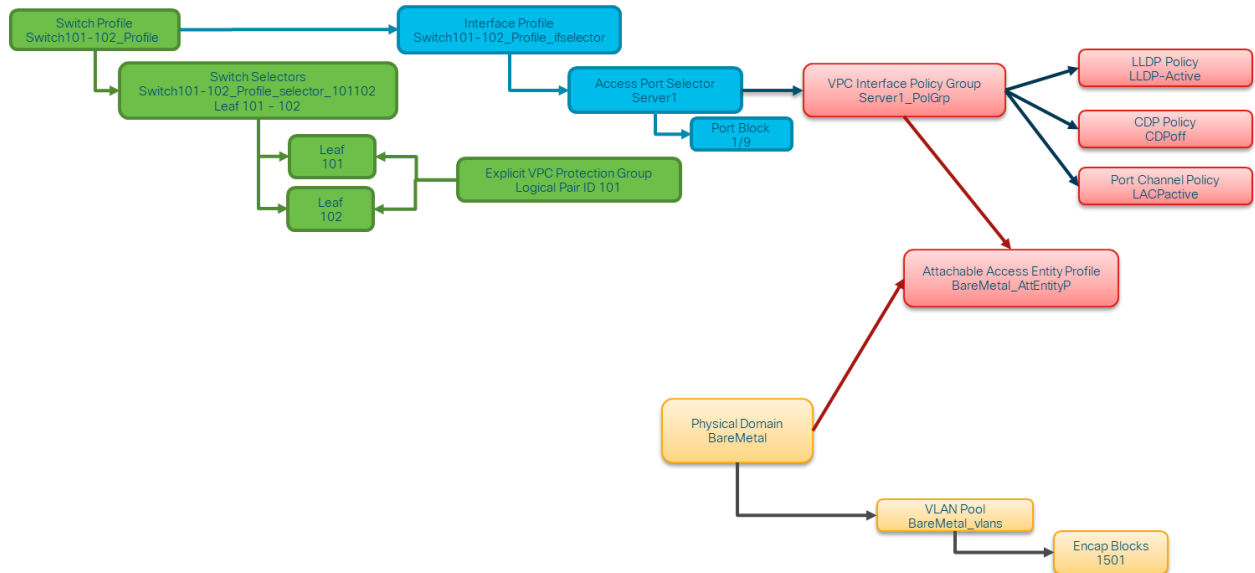
This means we now have a Physical domain that is fully equipped with access policies which can be linked to EPG's which we will later create.

The following is an overview of the created or re-used policies in detail.

Figure 133. Bare metal with physical domain: Switch Policy

Figure 134. Bare metal with physical domain: Interface Profile
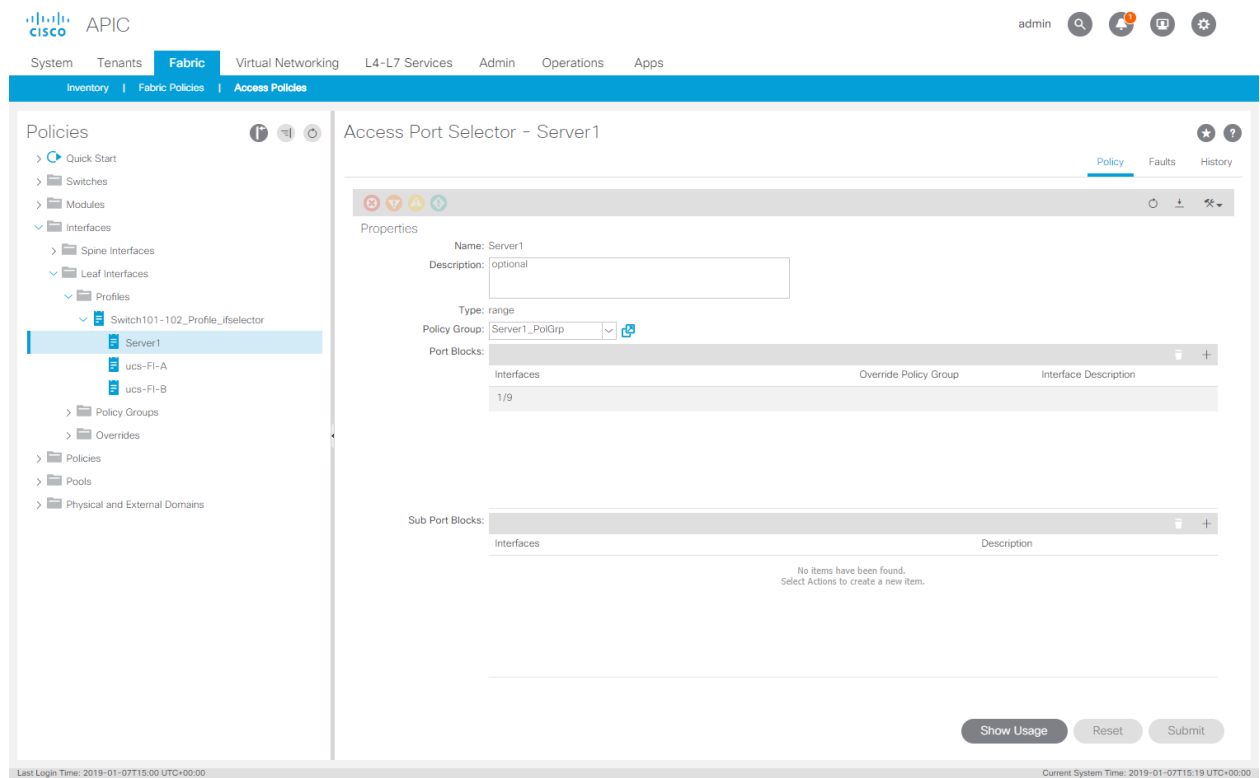
Figure 135. Bare metal with physical domain: Access port selectors
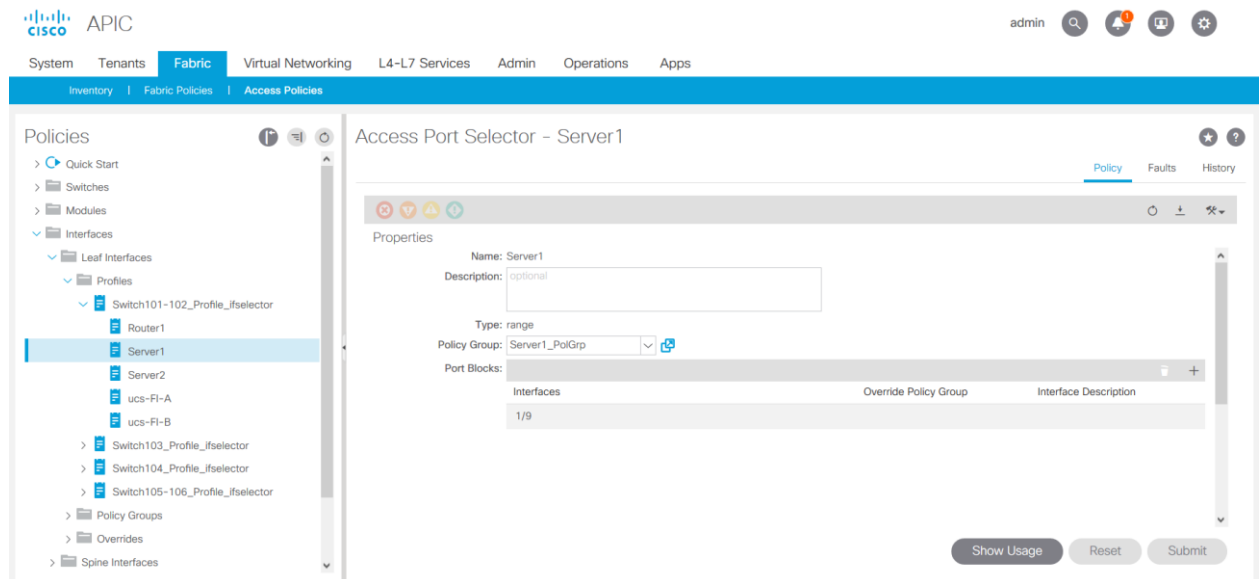
Figure 136. Bare metal with physical domain: VPC Interface Policy Group

Figure 137. Bare metal with physical domain: CDP Interface Policy

Figure 138. Bare metal with physical domain: LLDP Interface Policy
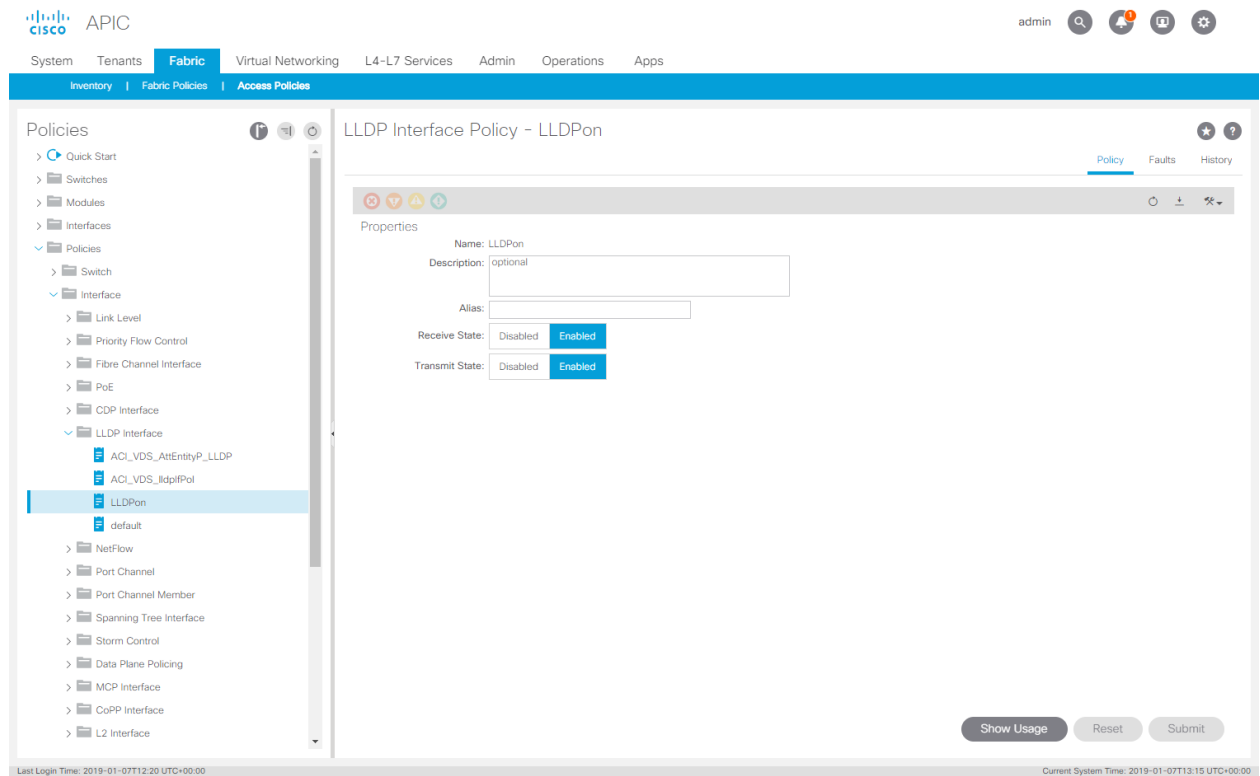
Figure 139. Bare metal with physical domain: Port Channel Policy
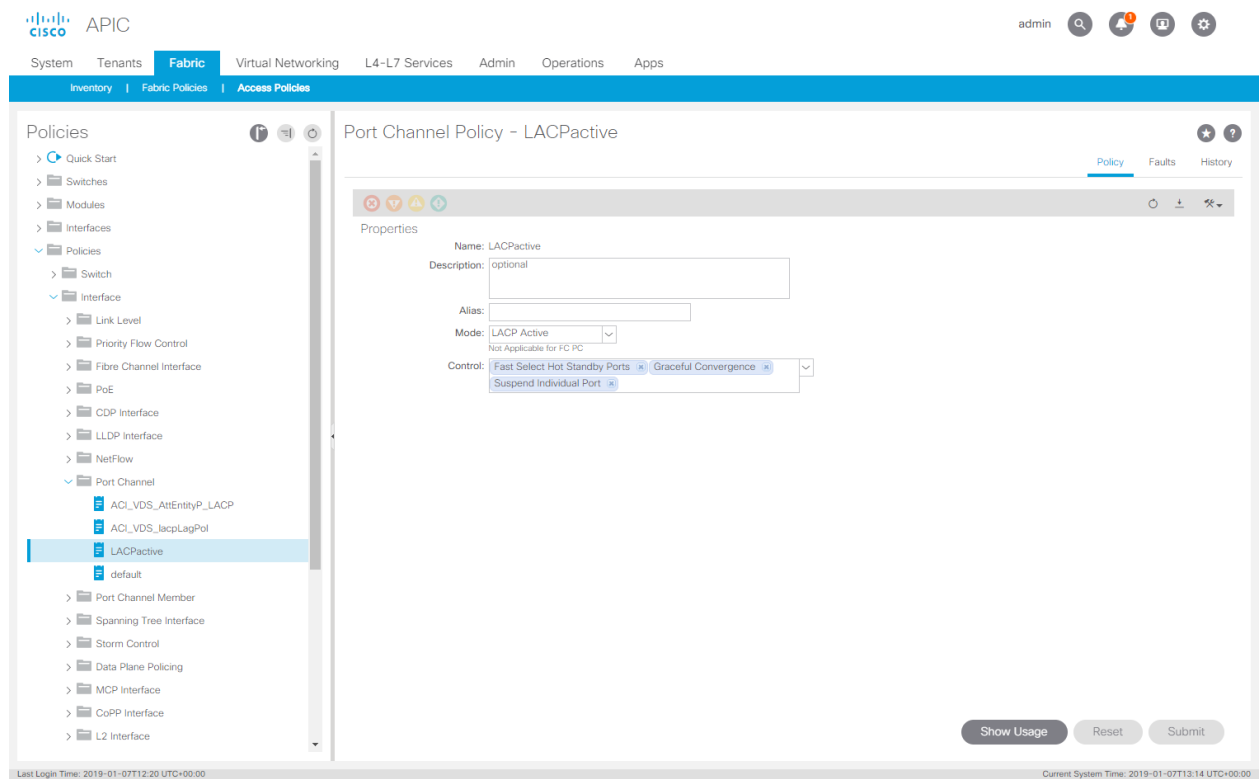
Figure 140. Bare metal with physical domain: Attachable Access Entity Profile

Figure 141. Bare metal with physical domain: VLAN Pool

Figure 142. Bare metal with physical domain: Physical Domain
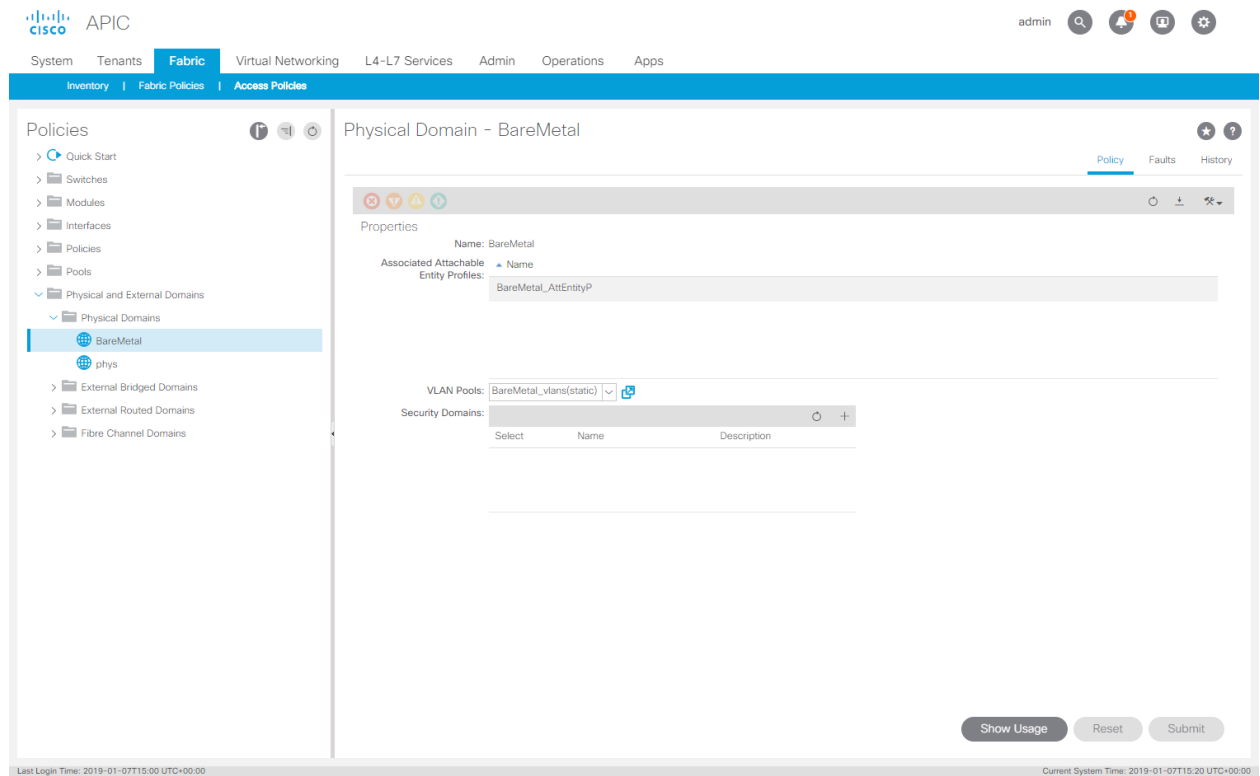
L3out to Router1

The following sections describe connecting Router1 to the fabric.

Connecting Router1

In this section we will connect router1 to the fabric which will be later on used as part of our L3out1 domain. This domain will also include router2 which we will configure in the next section. The connectivity diagram is as follows:

Figure 143. Router1 connectivity
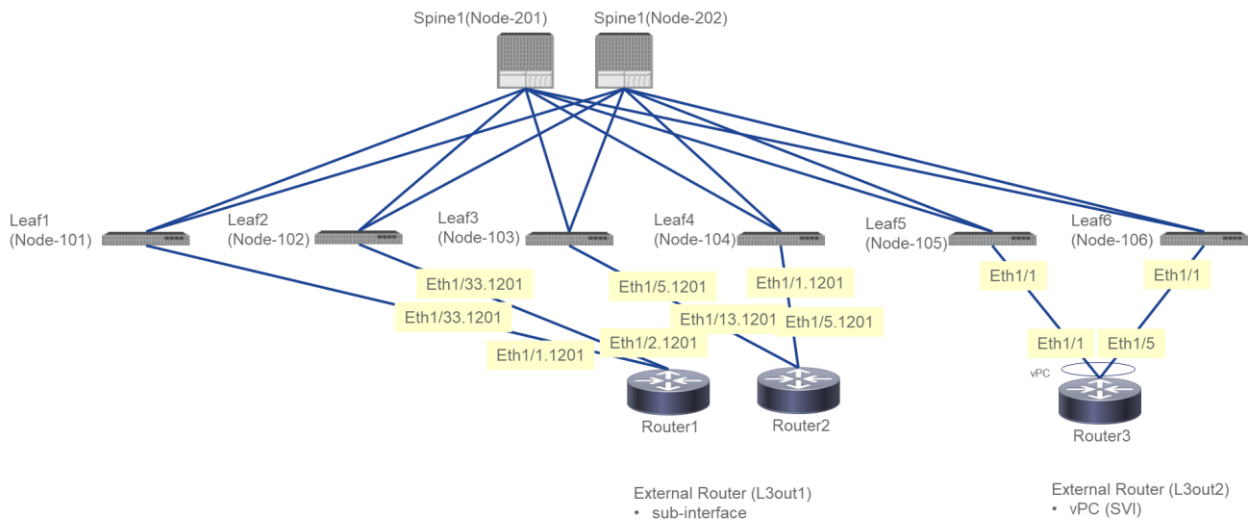
Notice we will connect router1 with 2 individual links and notice the symmetry which will allow us to use one "Switch Policy" with 2 leaves.

1. To start we will first launch the wizard.

Figure 144. Router1: Step 1
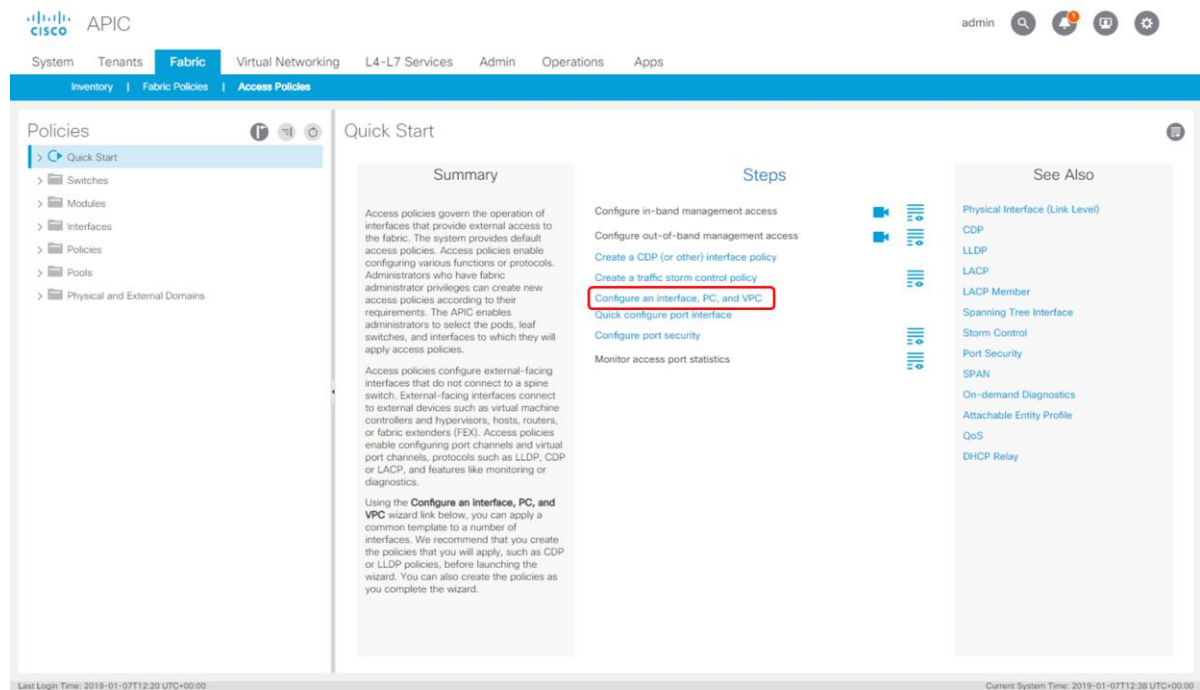
2. On the left select the "Switch Policy" for leaves 101 and 102 and click "+" on the right.

Figure 145. Router1: Step 2
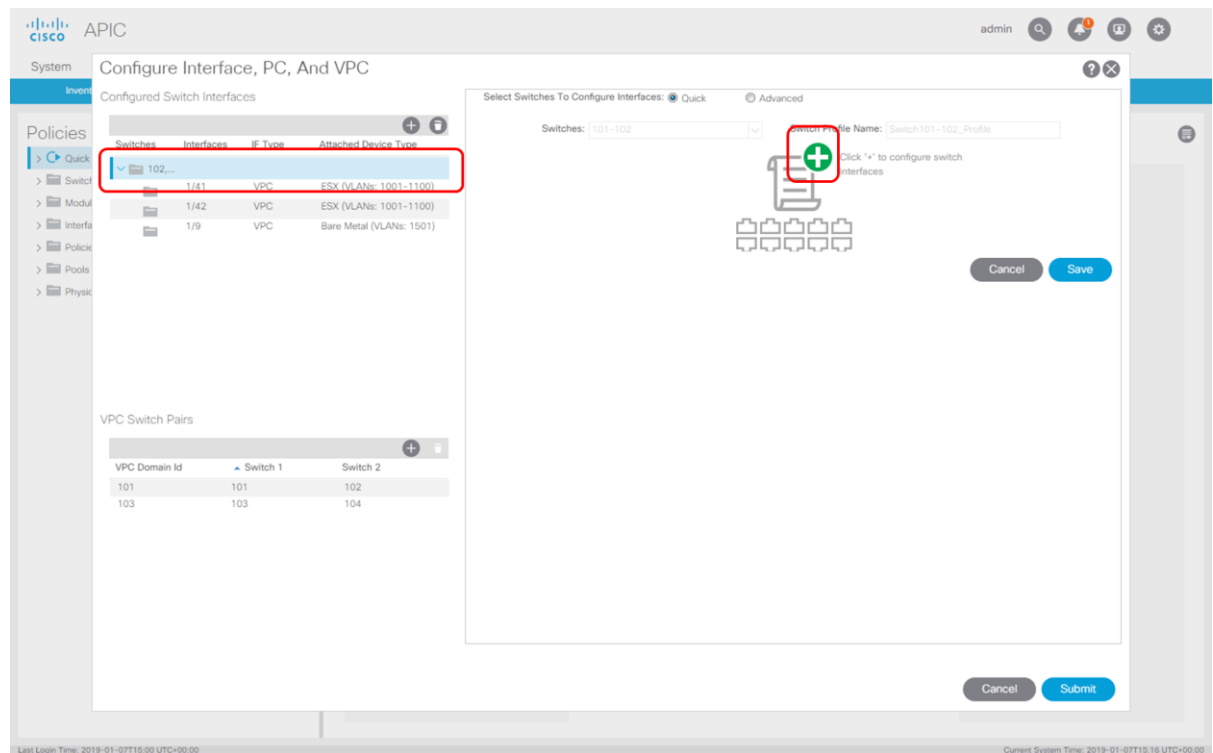
3. Select the "Interface Type" and fill in the "Interfaces". After that make sure to put a different name for the "Interface Selector Name", in our case we use "Router1".

Figure 146. Router1: Step 3
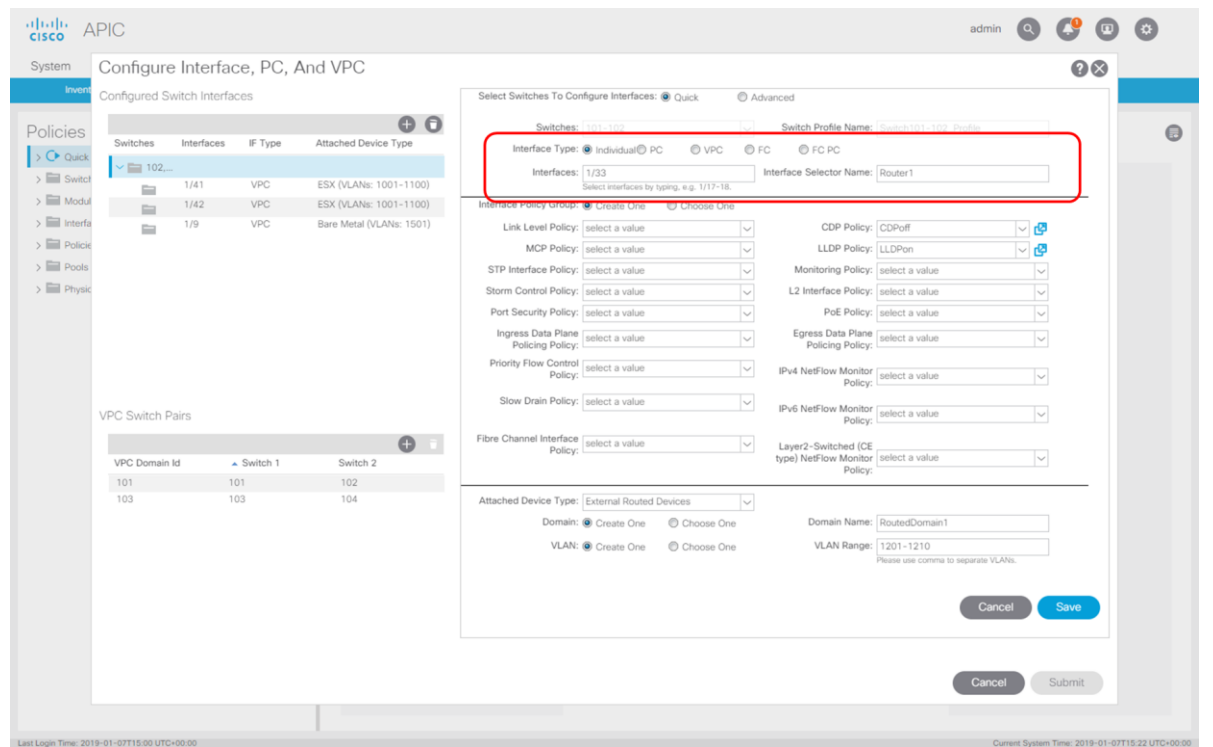
4. Now we select our existing "CDPoff" and "LLDPon" policy and last, we fill in the Domain section.

Figure 147. Router1: Step 4
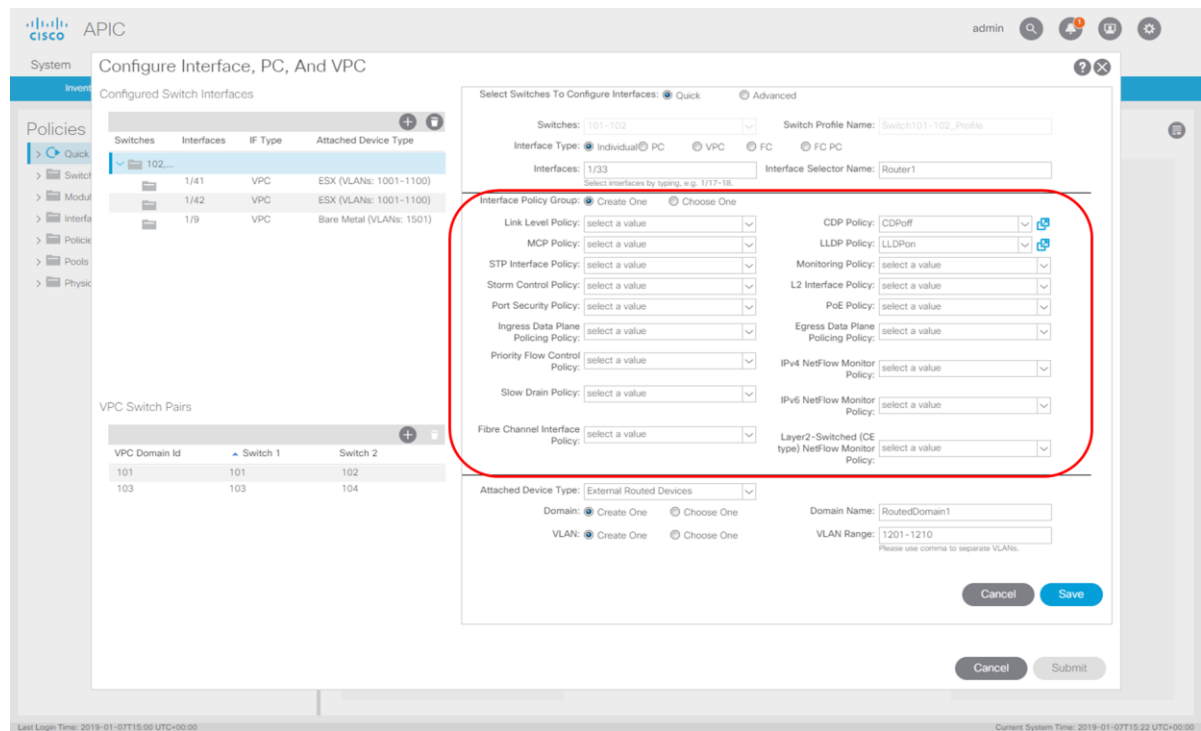
5. Here we select "External Routed Devices" and provide a name for this Routed Domain and create a new "VLAN Pool".

   Depending on the next part of the configuration in the tenant section we will need a VLAN or not, being, depending on if we will use a routed interface or SVI. As per our diagram we will use sub-interfaces so it's good to prepare a pool for this. We then click "Save"

Figure 148. Router1: Step 5
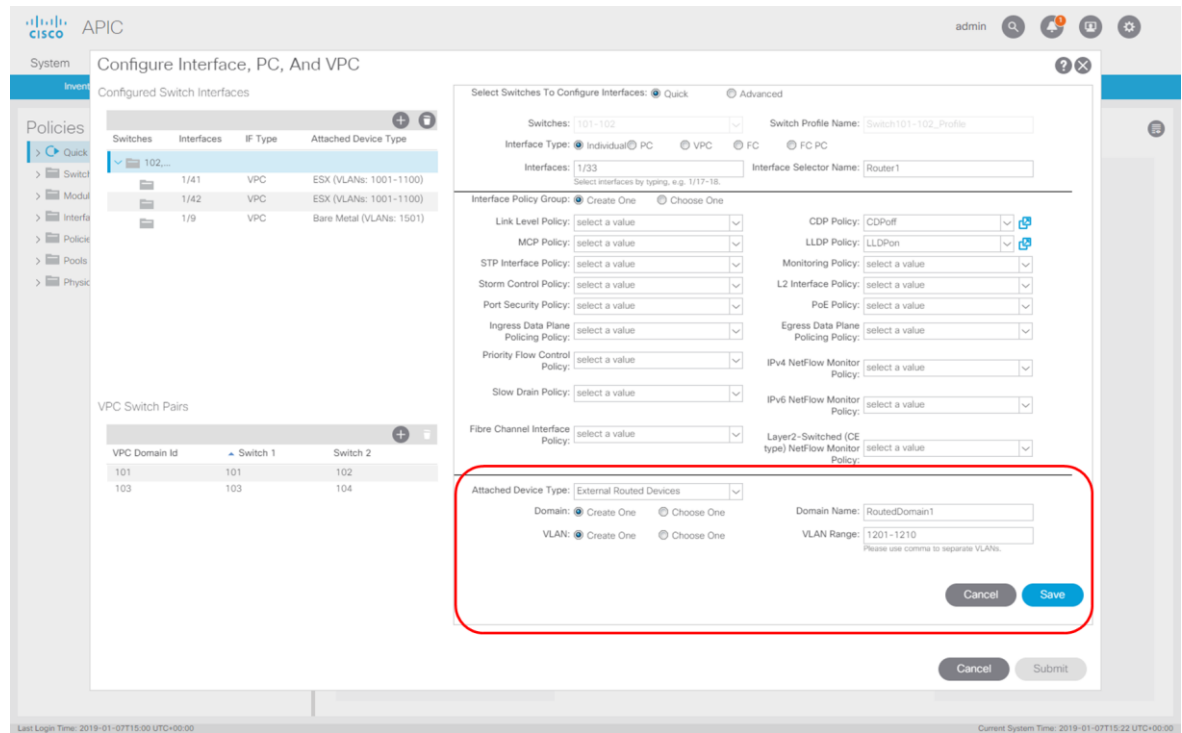
6. We click again " Save" .

Figure 149. Router1: Step 6
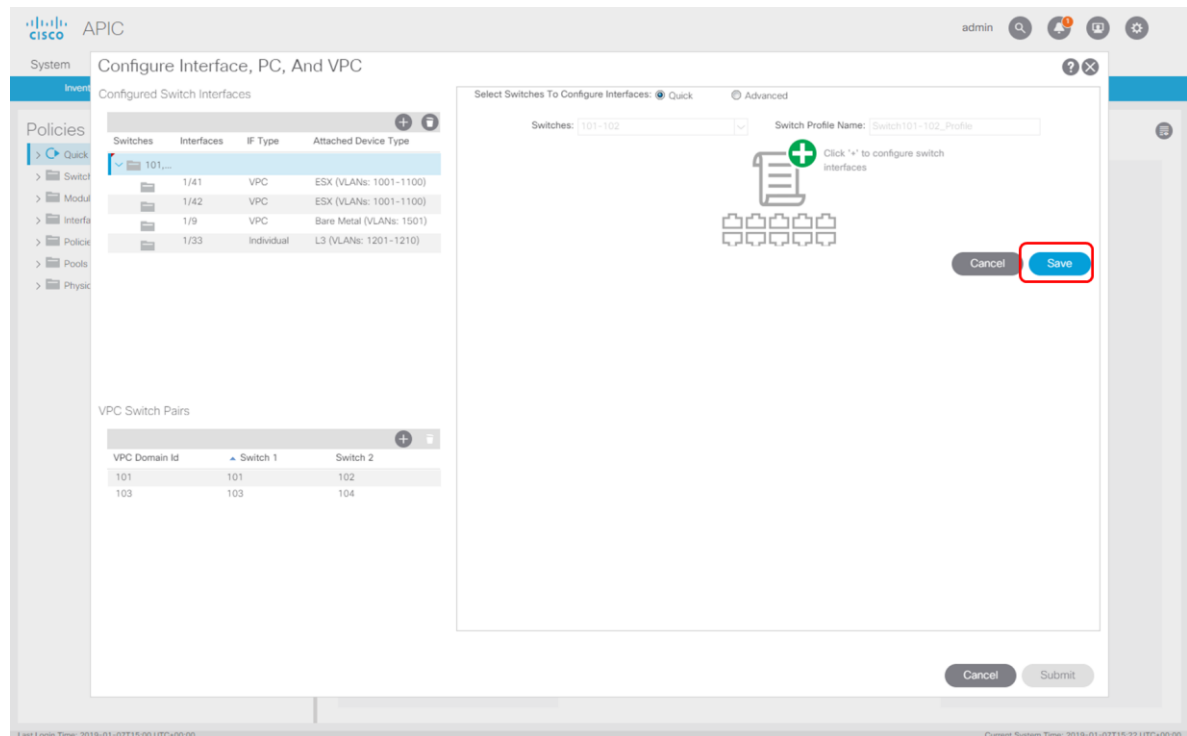
7. And finally, we click "Submit" to push all the policies.

Figure 150. Router1: Step 7
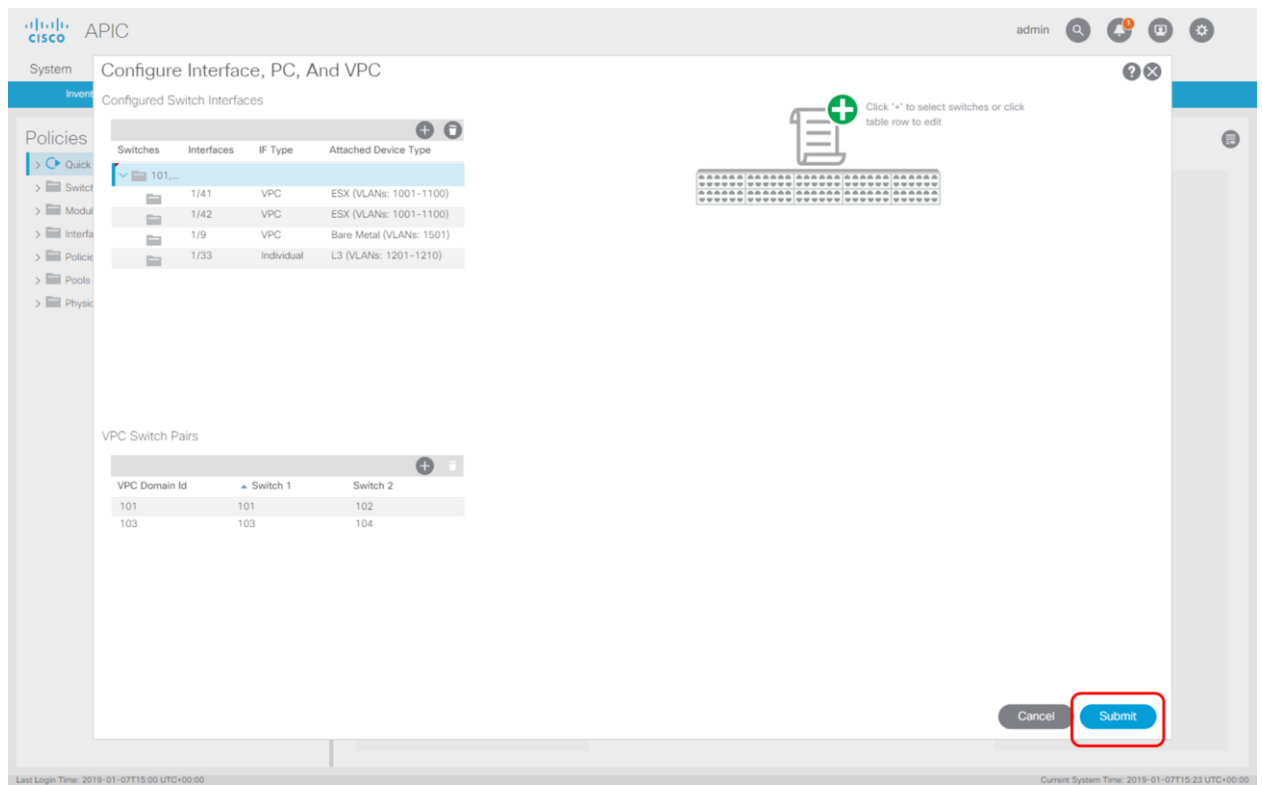
Overview of created policies

So, what happened when we have been executing this wizard? The following image shows all policies that have been created or linked to (already existing).

Figure 151. Router1: Policy overview
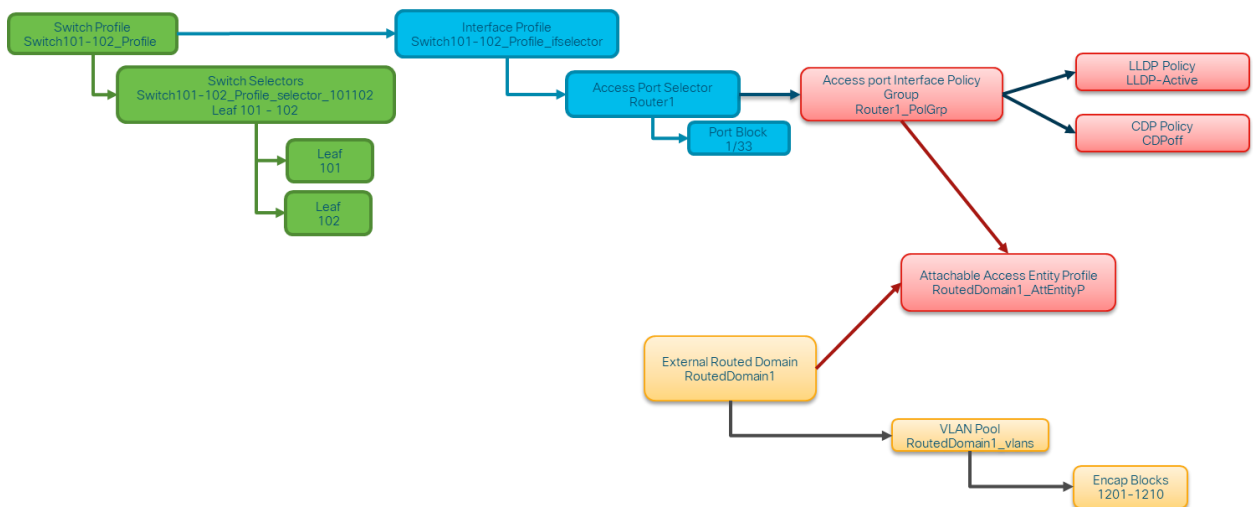
As we have symmetric interfaces in place we only created 1 "Switch Policy" and we have created an "External Routed Domain" called "RoutedDomain1" which is ready to be used in L3out configuration at tenant level in the chapter of this document. The following is an overview of the created or re-used policies in detail.

Figure 152. Router1: Switch Policy

Figure 153. Router1: Interface Profile

Figure 154. Router1: Access port selectors
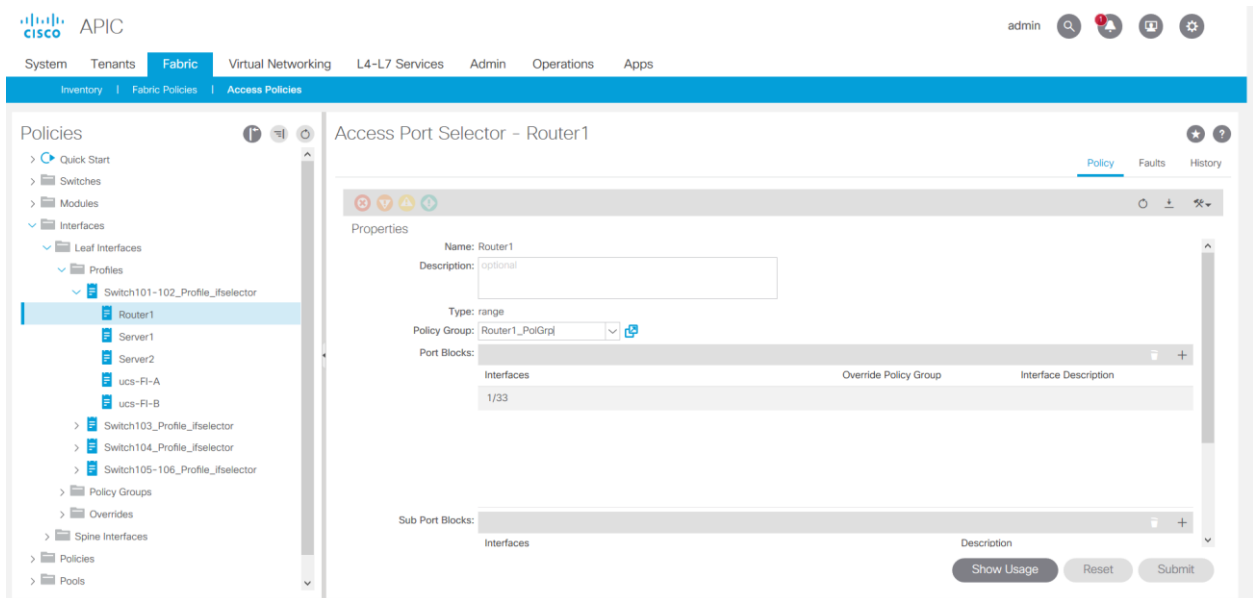
Figure 155. Router1: Leaf Access Port Policy

Figure 156. Router1: CDP Interface Policy

Figure 157. Router1: LLDP Interface Policy

Figure 158. Router1: Attachable Access Entity Profile

Figure 159. Router1: VLAN Pool

Figure 160. Router1: External Router Domains


L3out to router2 – Asymmetric policies

The following sections describe connecting Router2.

Connecting router2

In this section we will connect router2 to the fabric which will be later on used as part of our L3out1 domain. In the previous section we already created the "External Routed Domain" called "L3out1" so we will re-use that for this connectivity. The connectivity diagram is as follows:

Figure 161. Router2 connectivity
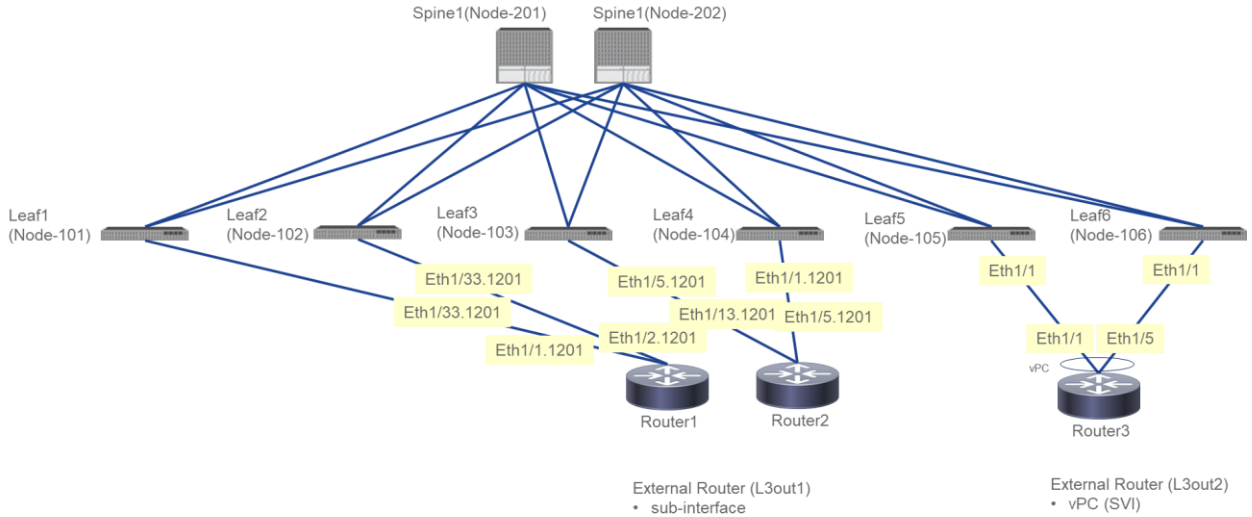
Notice we will connect router2 with 2 individual links and notice this router is asymmetric connected to leaf103 and leaf104. This means we have some more work when creating the policies through the wizard as we will need to make 2 individual "Switch Policy's" with their respective configuration linked to them.
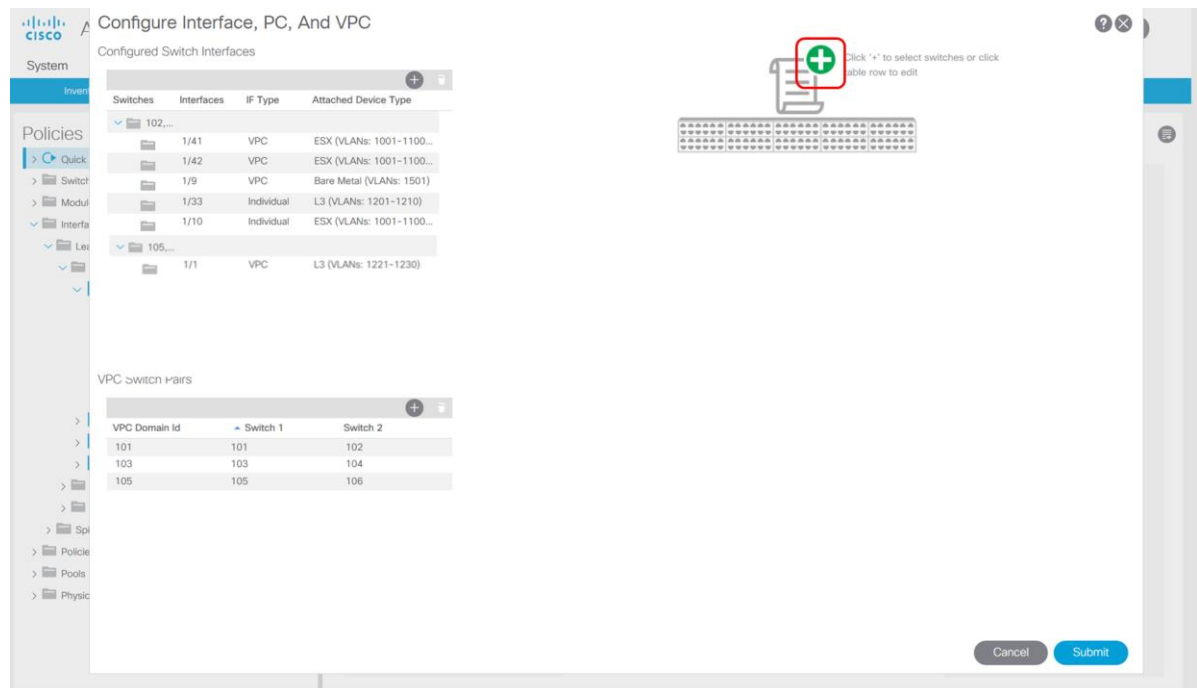
1. To start we will first launch the wizard.

Figure 162. Router2: Step 1
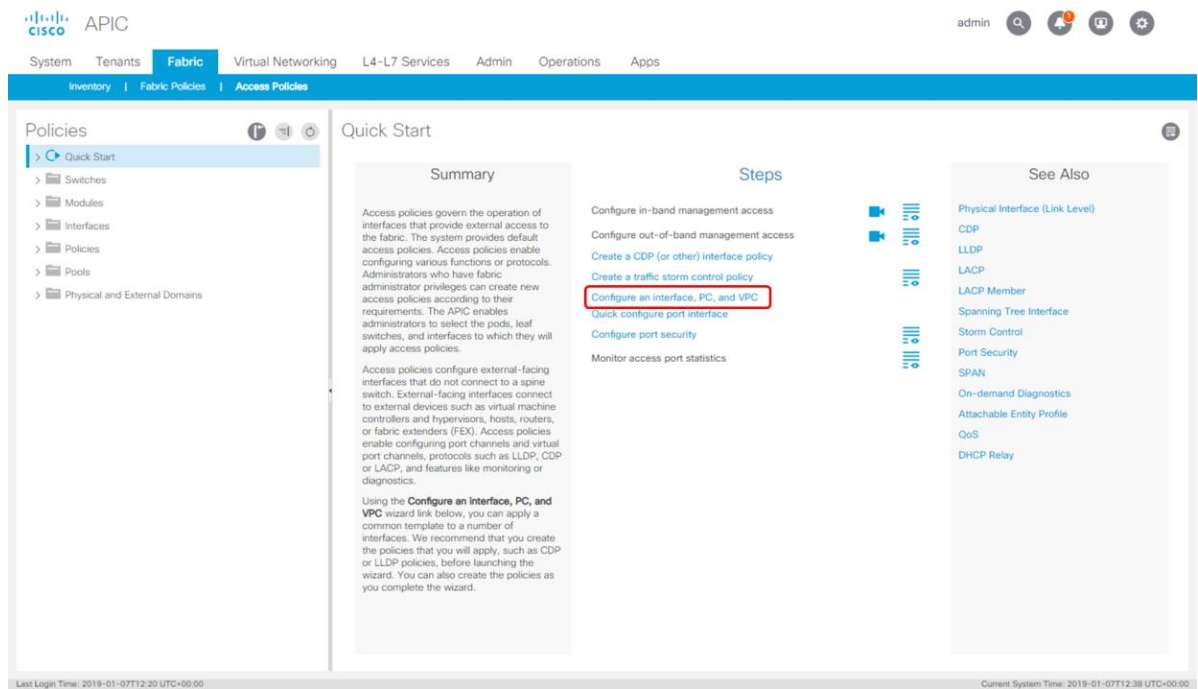
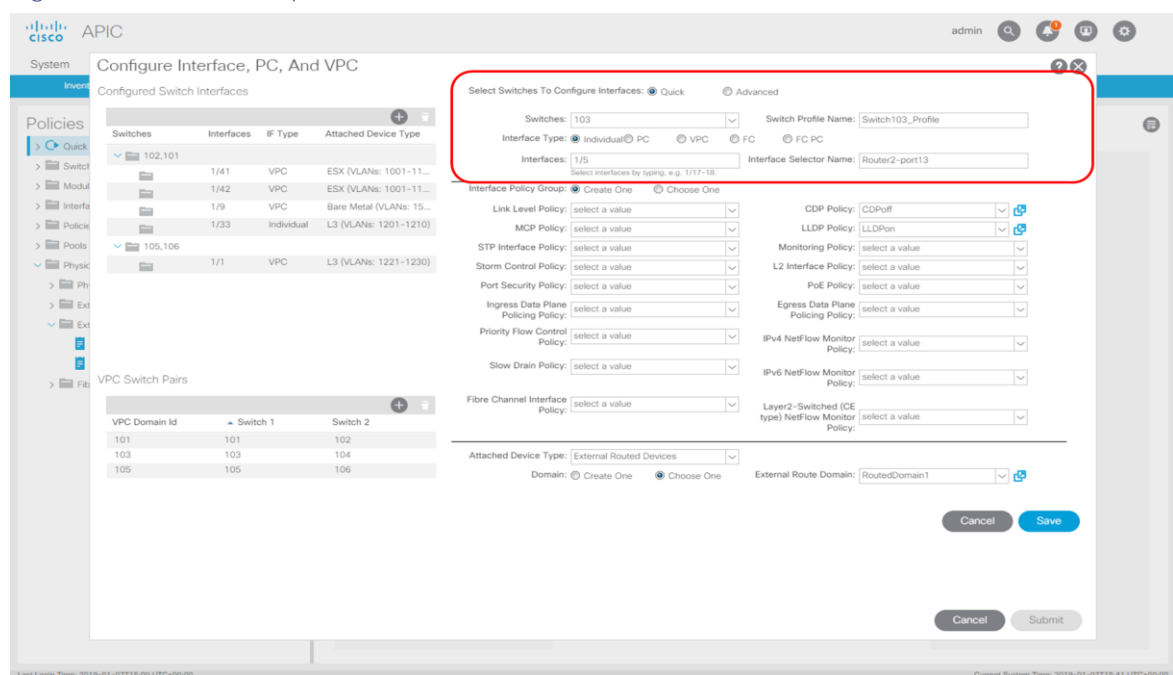2. Select "+" on the right.

Figure 163. Router2: Step 2

3. Now fill in the fields to create the "Switch Policy"

Notice we are only selecting Leaf 103 because do to the asymmetric routing in place we have to use a dedicated "Switch Policy" for each connected interface. Make sure to select Leaf 103 and click "+" to continue.

We will now provide the interface details, make sure to select as "Interface Type" "Individual" and give the interface a human understandable "Interface Selector Name".
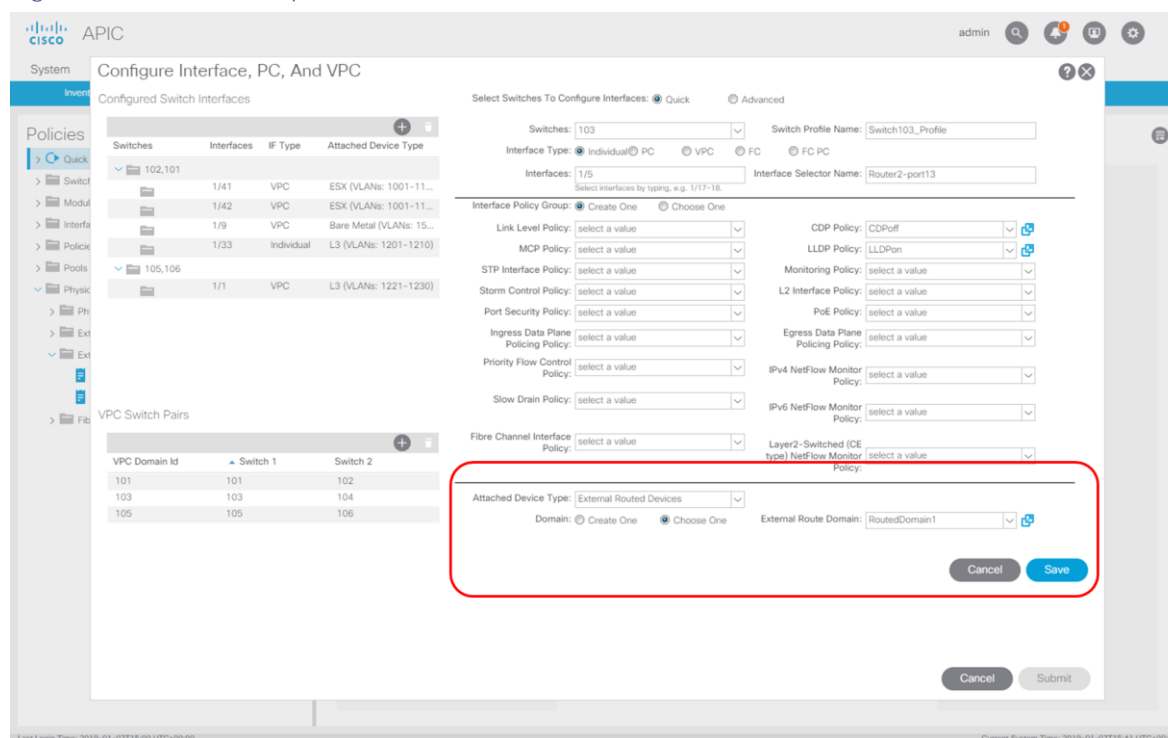
Figure 164. Router2: Step 3

Cisco Systems. All printed copies and duplicate soft copies are considered uncontrolled and the original online version should be referred to for the latest version.

Page 107 of 215

4. Next, we will link to the CDP and LLDP policies and continue to the Domain section.
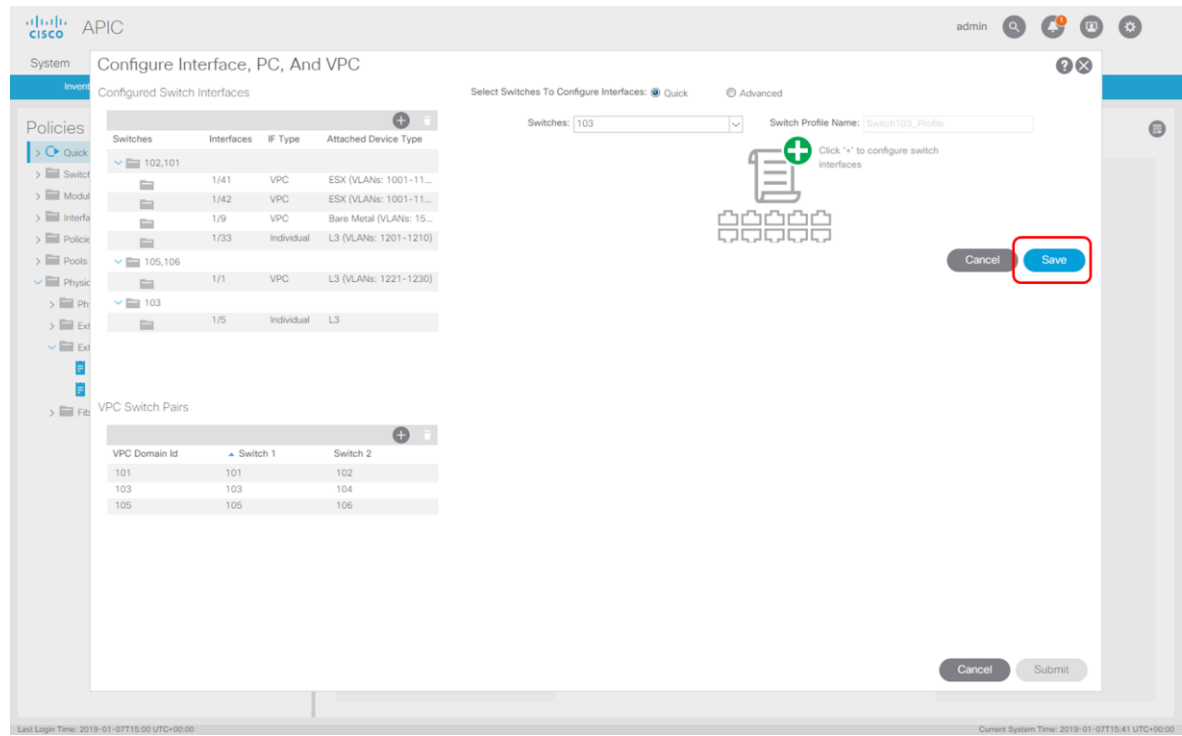
Figure 165. Router2: Step 4
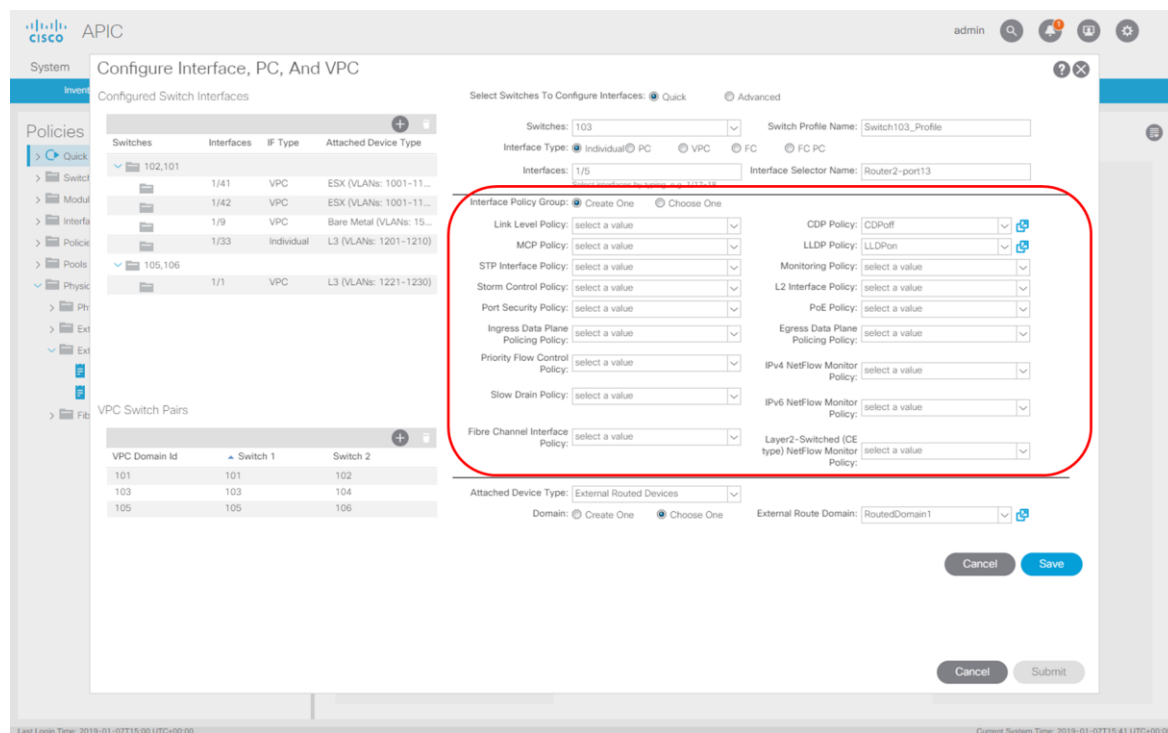
5. Notice that we are linking to an already existing "External Routed Devices" domain which is called "RoutedDomain1". The benefits of linking multiple interfaces to the same routed domain will be explained further in this document. Click "Save" to continue.
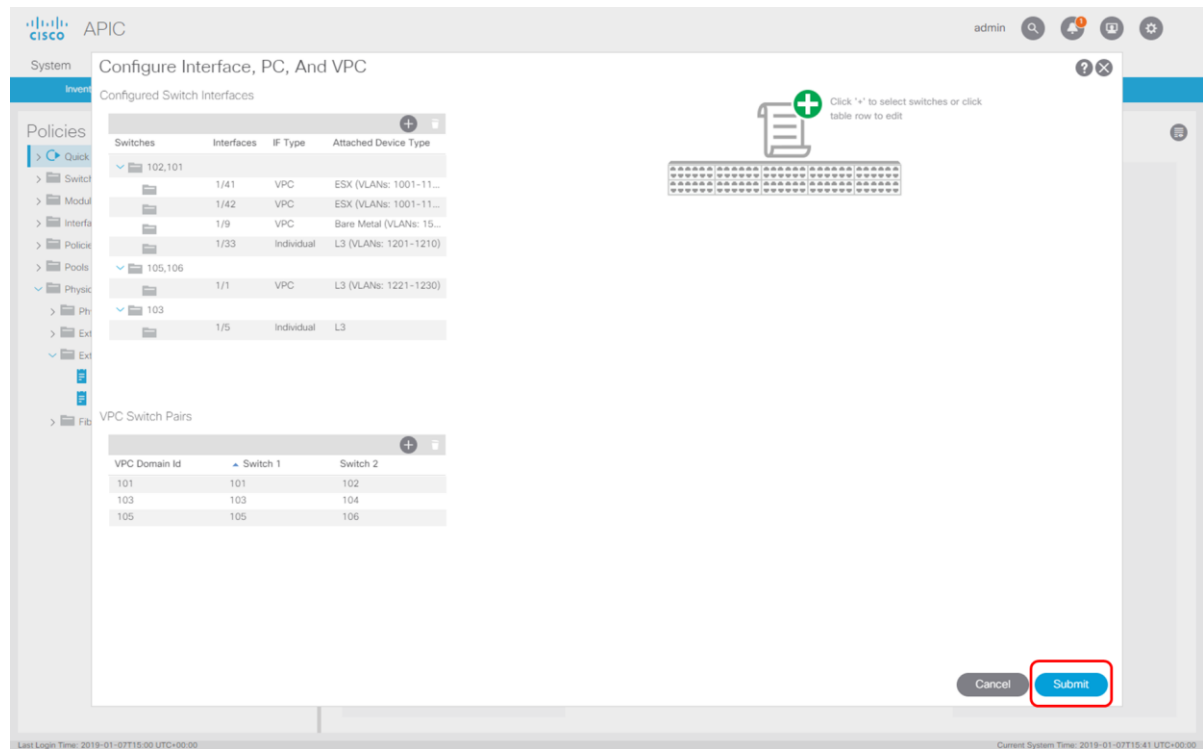
Figure 166. Router2: Step 5

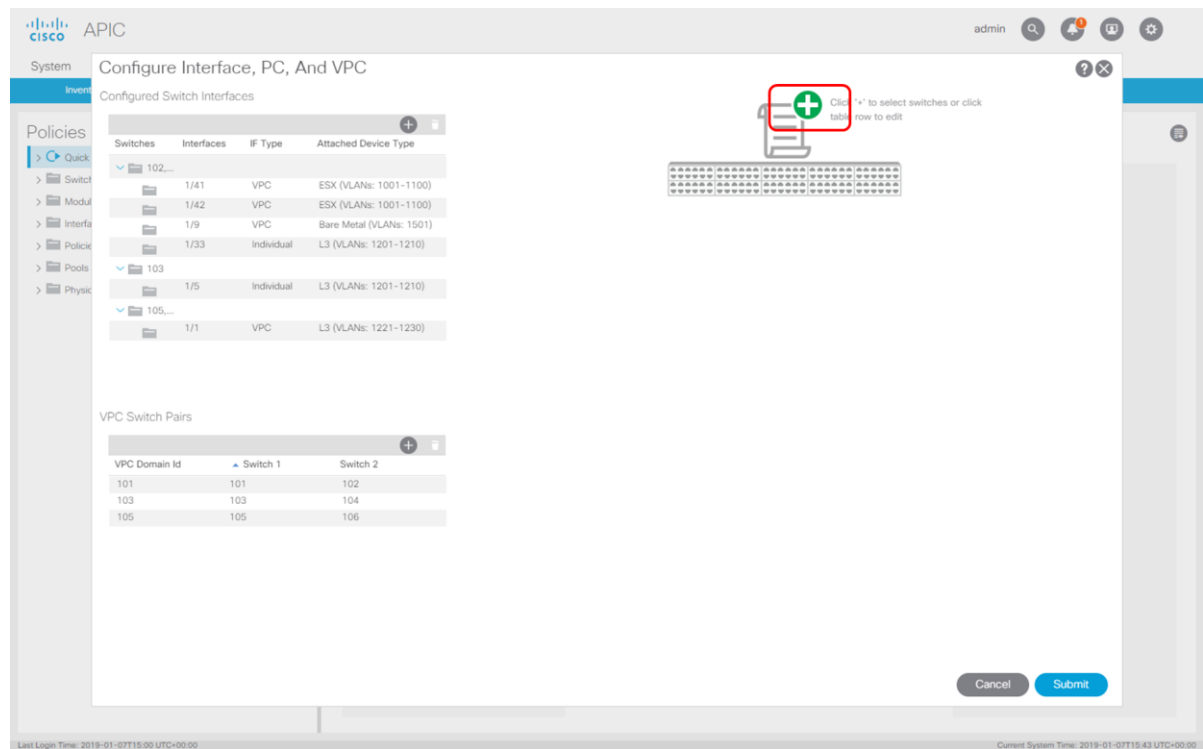6. Click "Save".

Figure 167. Router2: Step6

7. And click "Submit". Now relaunch the wizard
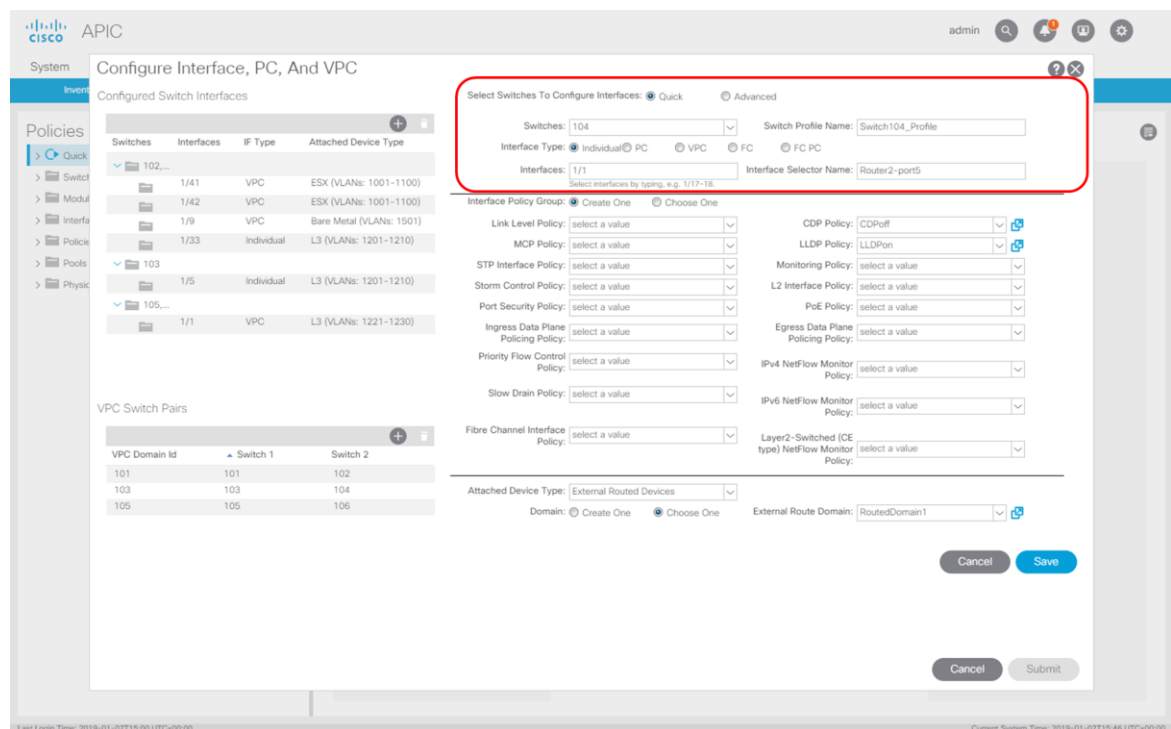
Figure 168. Router2: Step 7

8. And click "+" to add a new "Switch Policy".

Figure 169. Router2: Step 8

9. Make sure to fill in the correct switch ID. Select "Interface Type" "Individual" and fill in the "Interfaces" field with port 1/1. Make sure to put a human understandable name for "Interface Selector Name" for easy reference in the future.
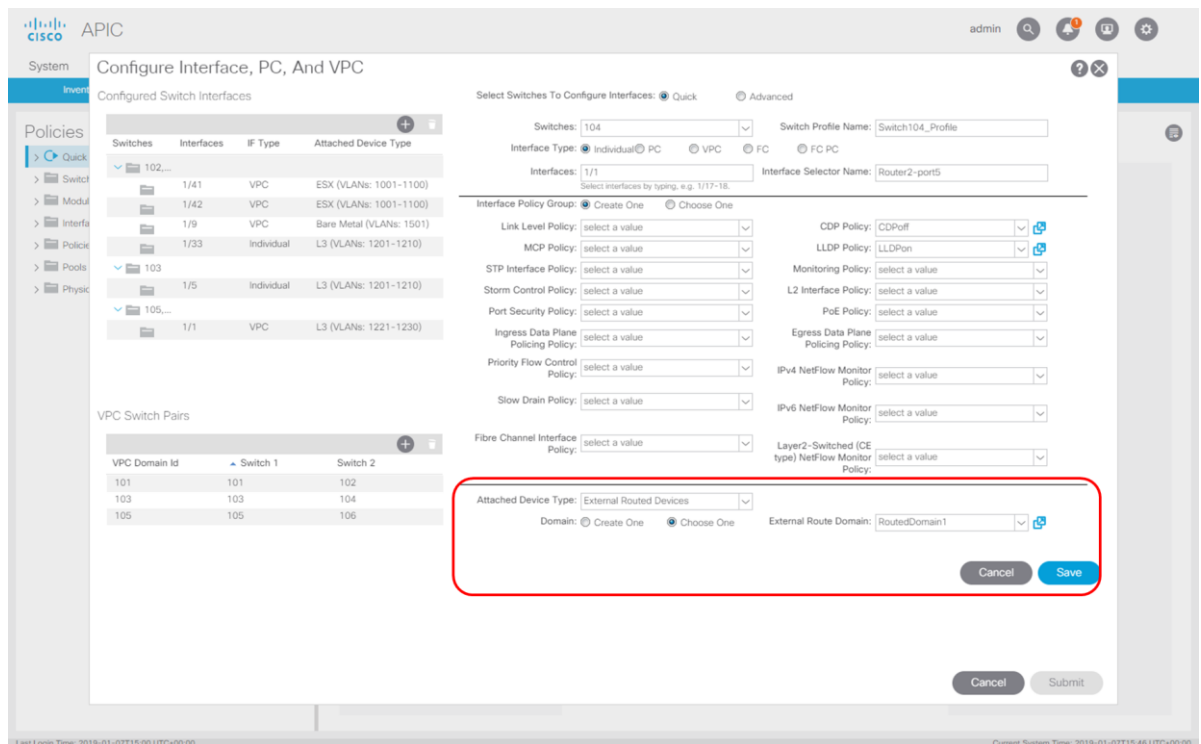
Figure 170. Router2: Step 9

10. Make sure to select the CDPoff and LLDPon policy and continue to the Domain section.
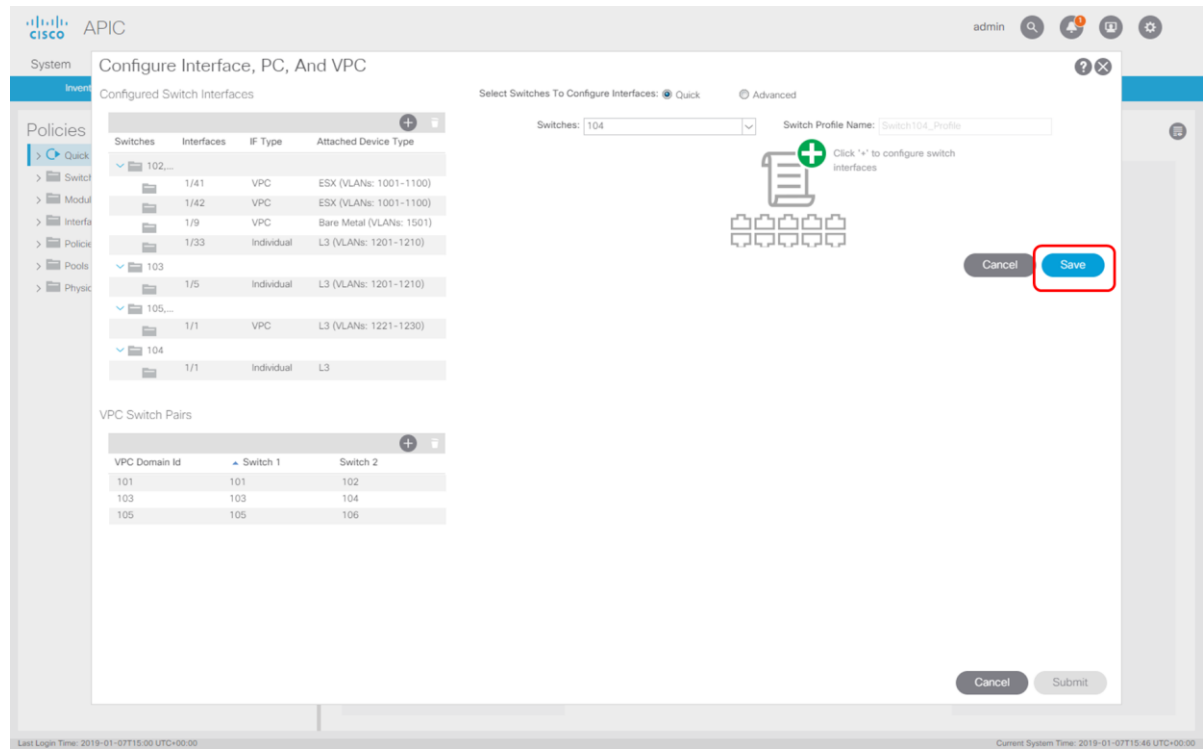
Figure 171. Router2: Step 10
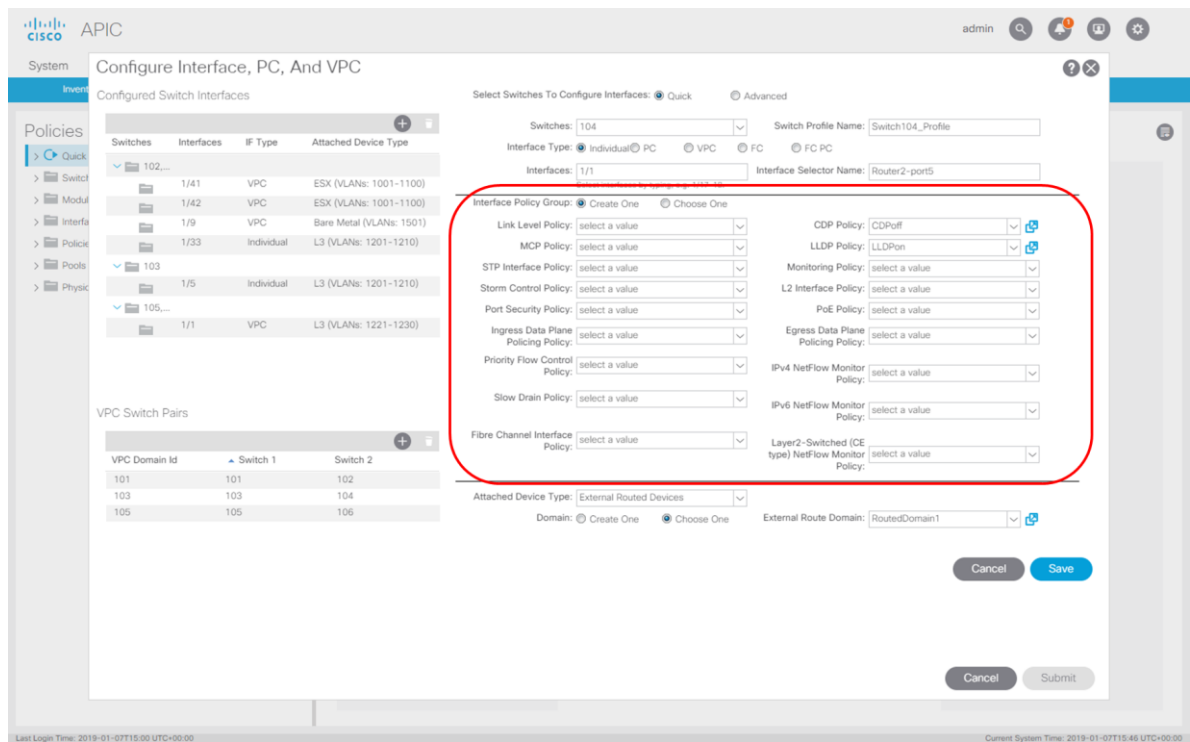
11. We will link to an existing domain called "RoutedDomain1" and click "Save".
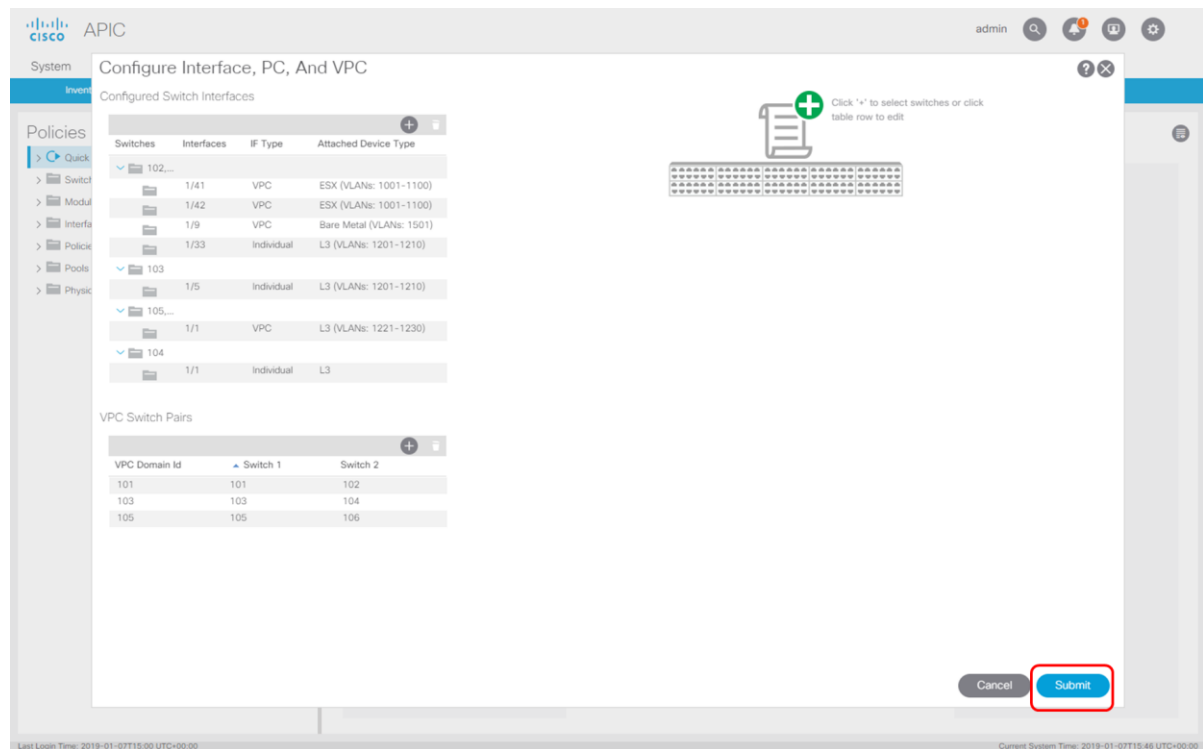
Figure 172. Router2: Step 11

12. Click "Save"

Figure 173. Router2: Step 12

13. And click "Sumit".
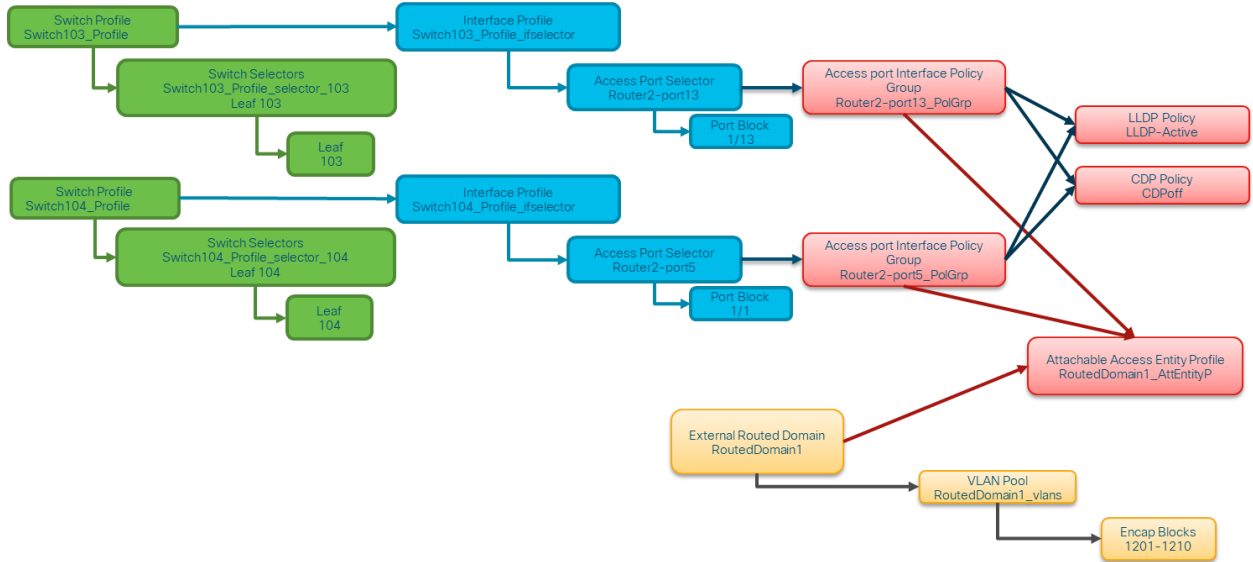
Figure 174. Router2: Step 13

Overview of Created Policies

So, what happened when we have been executing this wizard? The following image shows all policies that have been created or linked to (already existing).

Figure 175. Router2: Policy overview

As we have asymmetric interfaces in place we have two "Interface Profiles" created with their respective policies. We are linking to an existing "External Routed Domain" which can later on be used at tenant level. The following is an overview of the created or re-used policies in detail.

Figure 176. Router 2: Switch Policy

Figure 177. Router 2: Interface Profile

Router 2: Access port selectors

Figure 178. Router 2: Access Port Policy Group

Cisco Systems. All printed copies and duplicate soft copies are considered uncontrolled
and the original online version should be referred to for the latest version.

Page 117 of 215

Figure 179. Router 2: CDP Interface Policy

Figure 180. Router 2: LLDP Interface Policy

Figure 181. Router 2: Attachable Access Entity Profile

Figure 182. Router 2: VLAN Pool

Figure 183. Router 2: External Router Domains
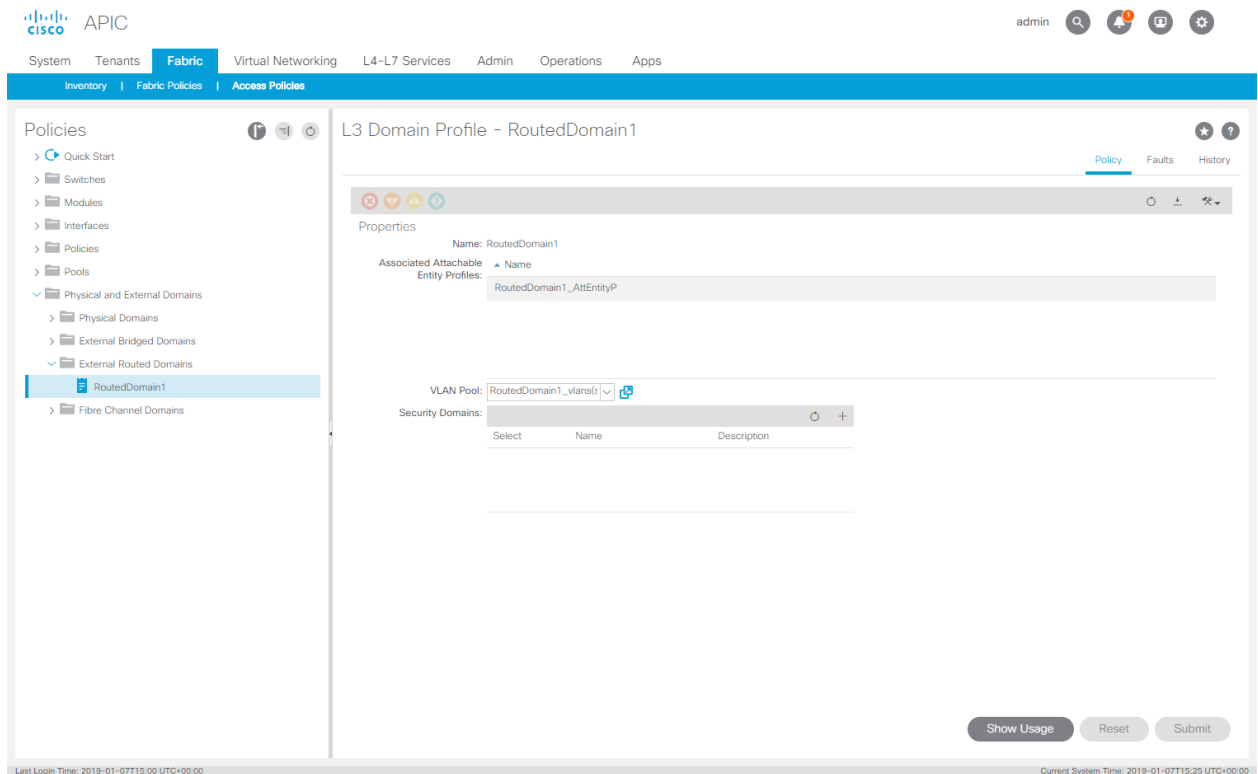
L3out to Router3

The following sections describe conneting Router3.

Connecting router3

In this section we will connect router3 to the fabric which will be later on used as part of our L3out2 domain. This router will be connected with an LACP port-channel into the fabric, this means we also need to create for leaf 105 and leaf 106 a VPC pair. The connectivity diagram is as follows:

Figure 184. Router3 connectivity

Notice we will connect router3 with 2 individual links and notice the symmetry which will allow us to use one "Switch Policy" with 2 leaves.

1. To start we will first launch the wizard.

Figure 185. Router3: Step 1
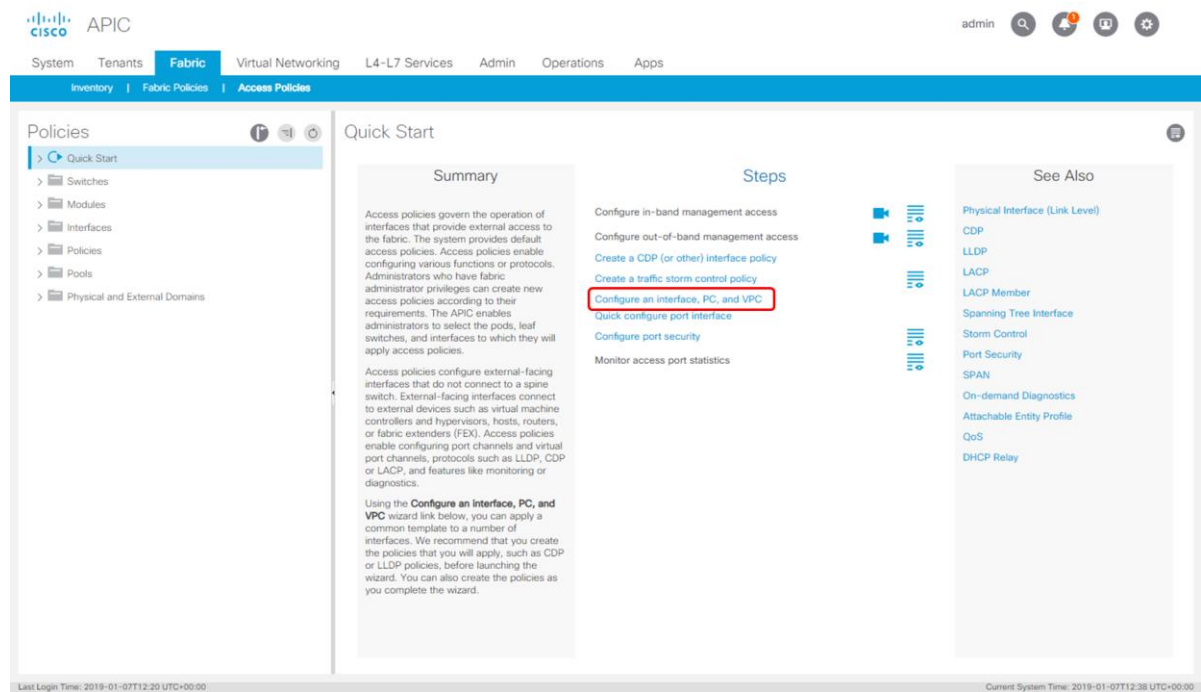
2. The first thing we will do is create a VPC policy for switch 105 and 106.

Figure 186. Router3: Step 2
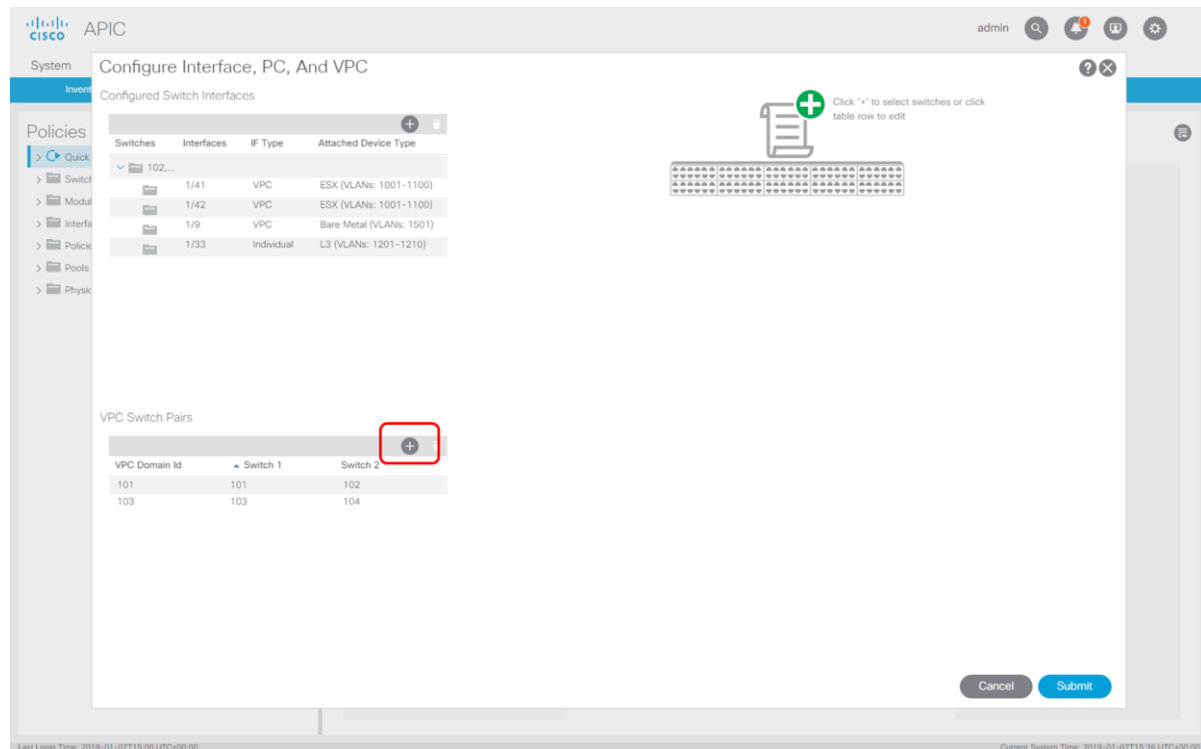
3. Select the "+" sign and launch the VPC add wizzard.

   We will use "VPC Domain ID" 105 and select both leaf 105 and 106. Now click "Save".

Figure 187. Router3: Step 3
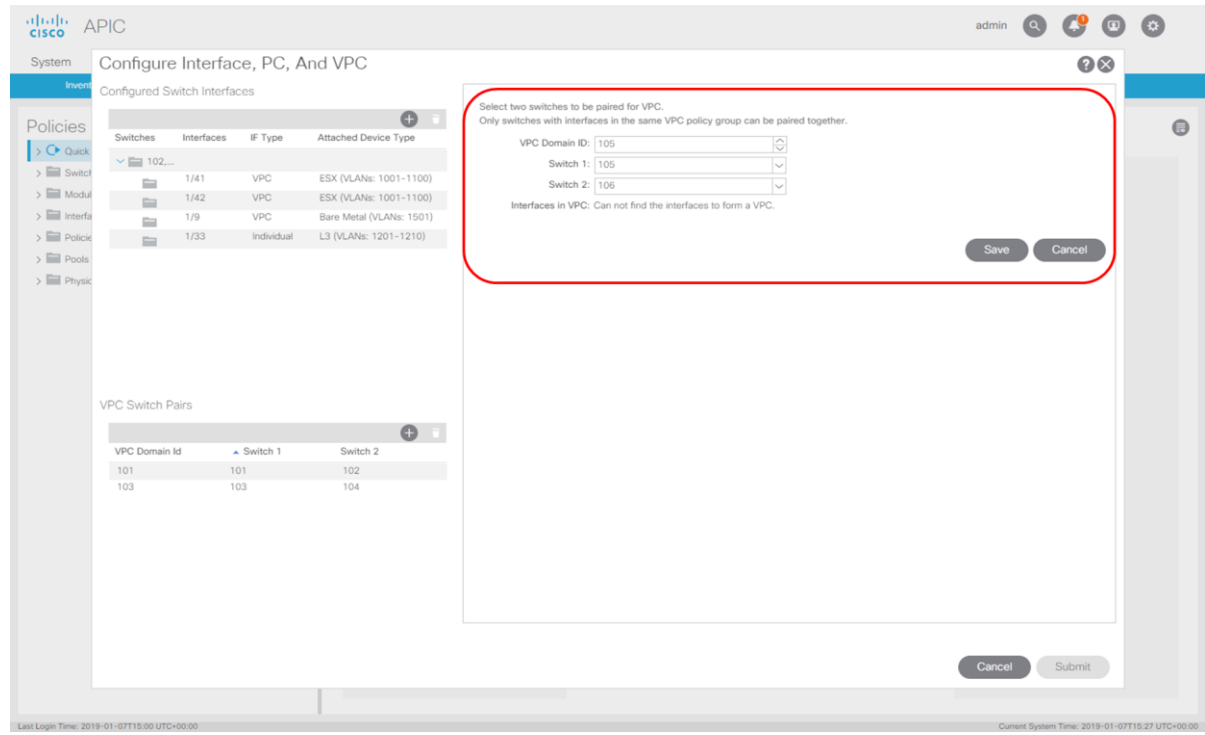
4. Now click "+" to create a new "Switch Policy"

Figure 188. Router3: Step 4
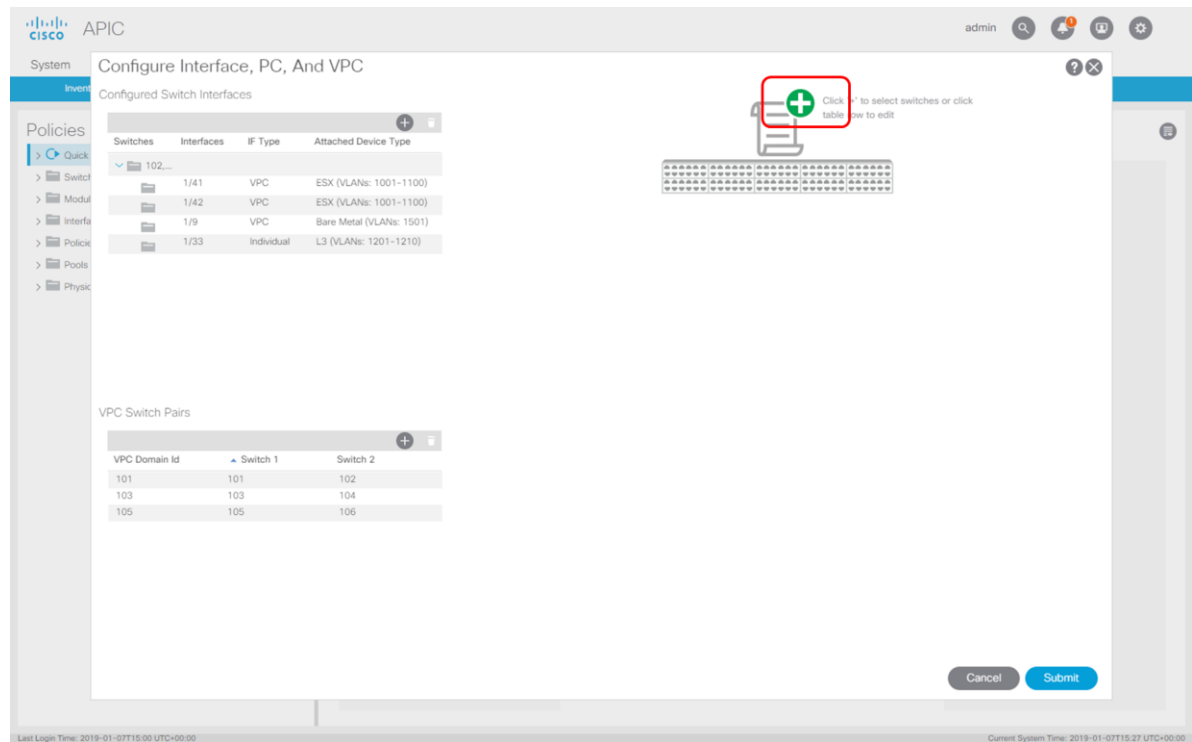
5. And select both leaf 105 and 106. Click "Save" to continue.

Figure 189. Router3: Step 5
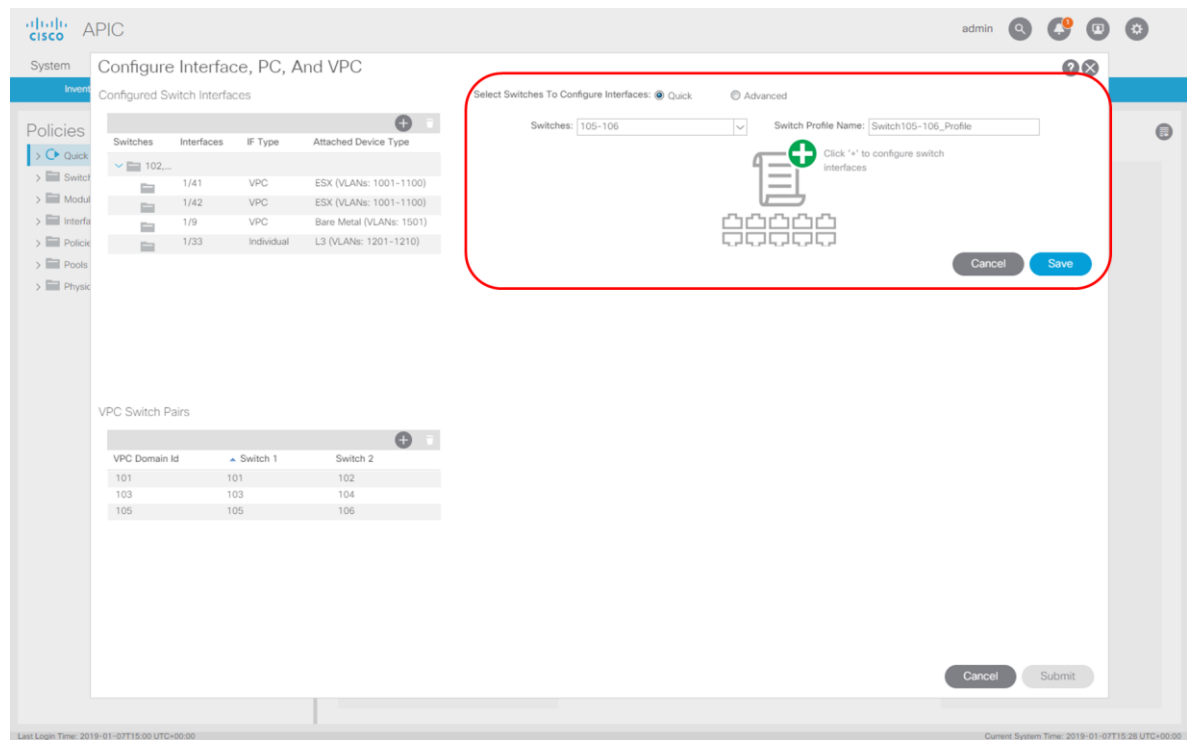
6. Select the "Interface Type" (make sure to use VPC) and fill in the "Interfaces". After that make sure to put a different name for the "Interface Selector Name", in our case we use "Router3".

Figure 190. Router3: Step 6
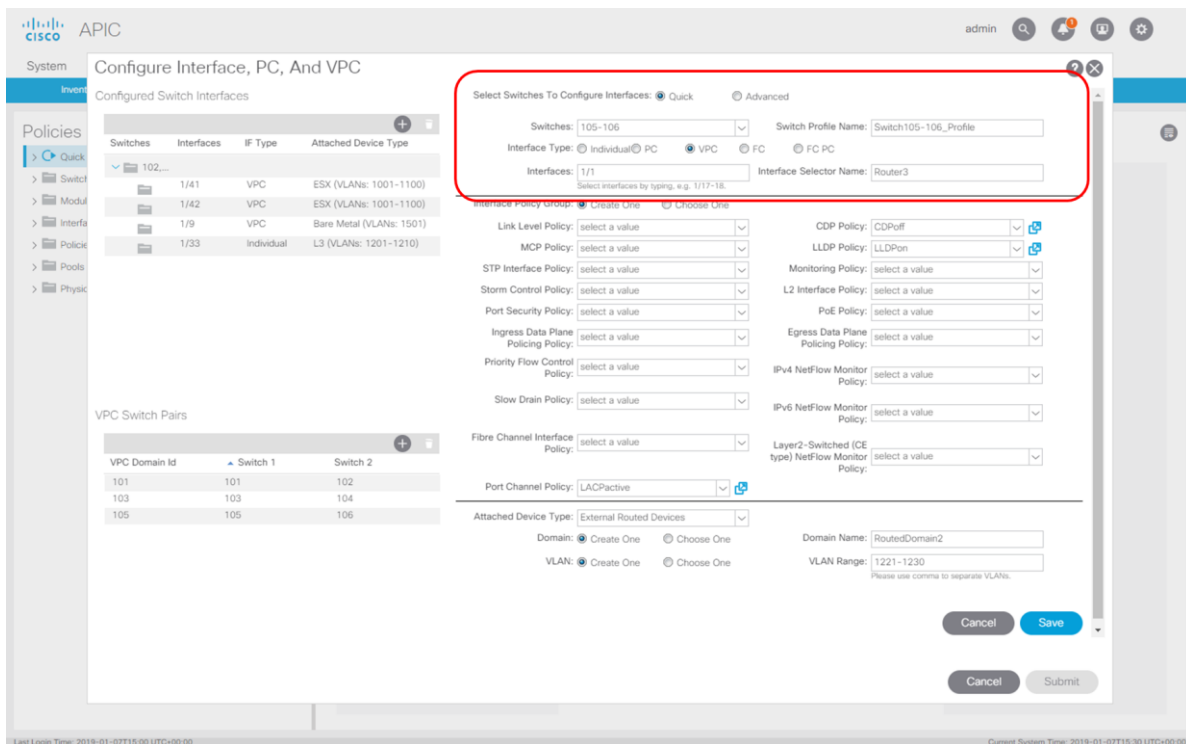
7. Now we select our existing "CDPoff", "LLDPon" and "Port Channel" policies and last, we fill in the Domain section.

Figure 191. Router3: Step 7
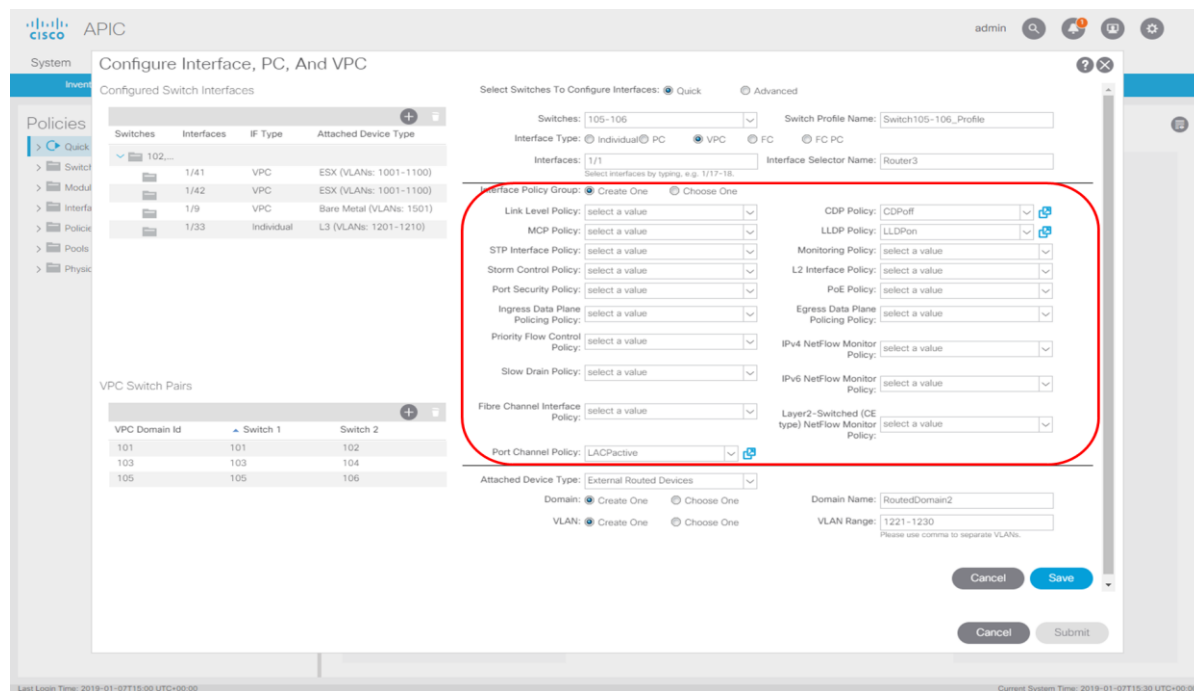
8. Here we select "External Routed Devices" and provide a name for this Routed Domain and create a new "VLAN Pool". Depending on the next part of the configuration in the tenant section we will need a VLAN or not, being, depending on if we will use a routed interface or SVI. As per our diagram we will use routed interfaces but we will create anyway already a VLAN pool if ever we need to create a subinterface or SVI in the L3out. We then click "Save"

Figure 192. Router3: Step 8
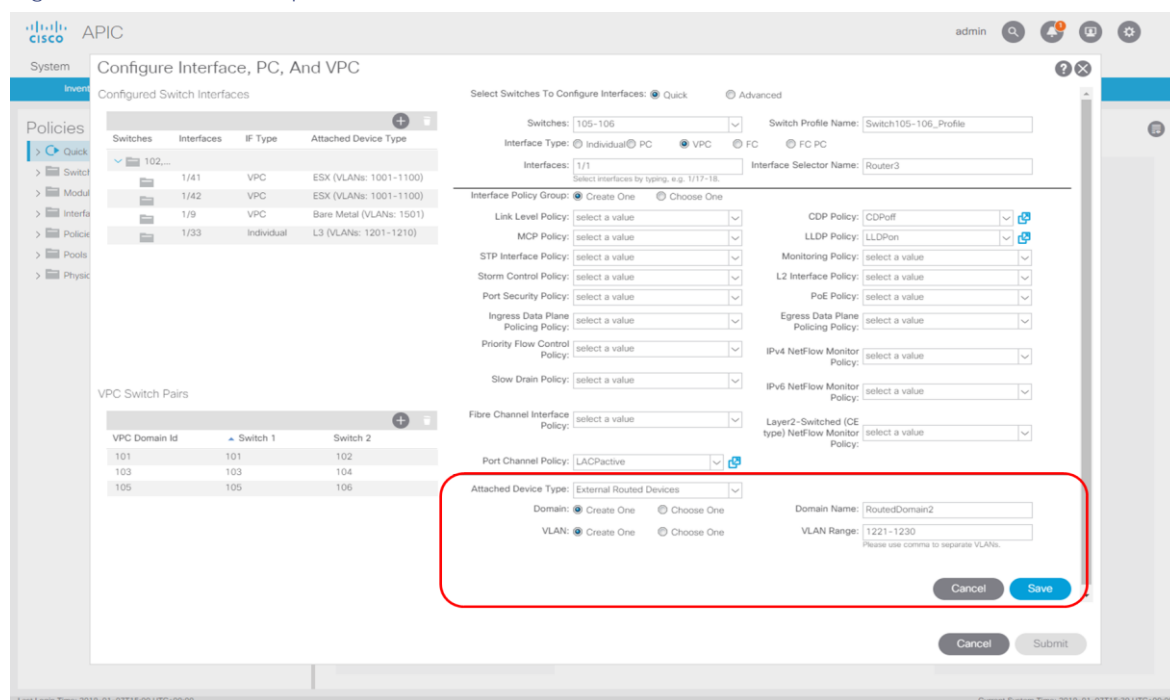
9. We click "Save" again.

Figure 193. Router3: Step 9
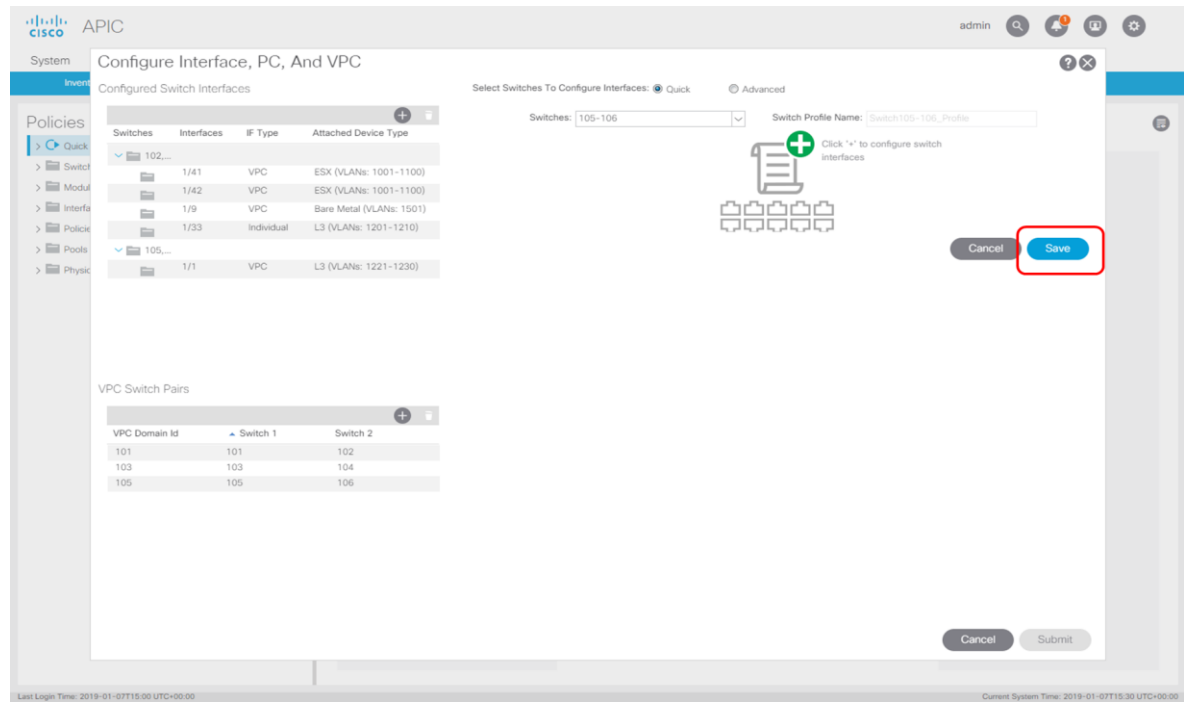
10. And finally, we click "Submit" to push all the policies.

Router3: Step 10
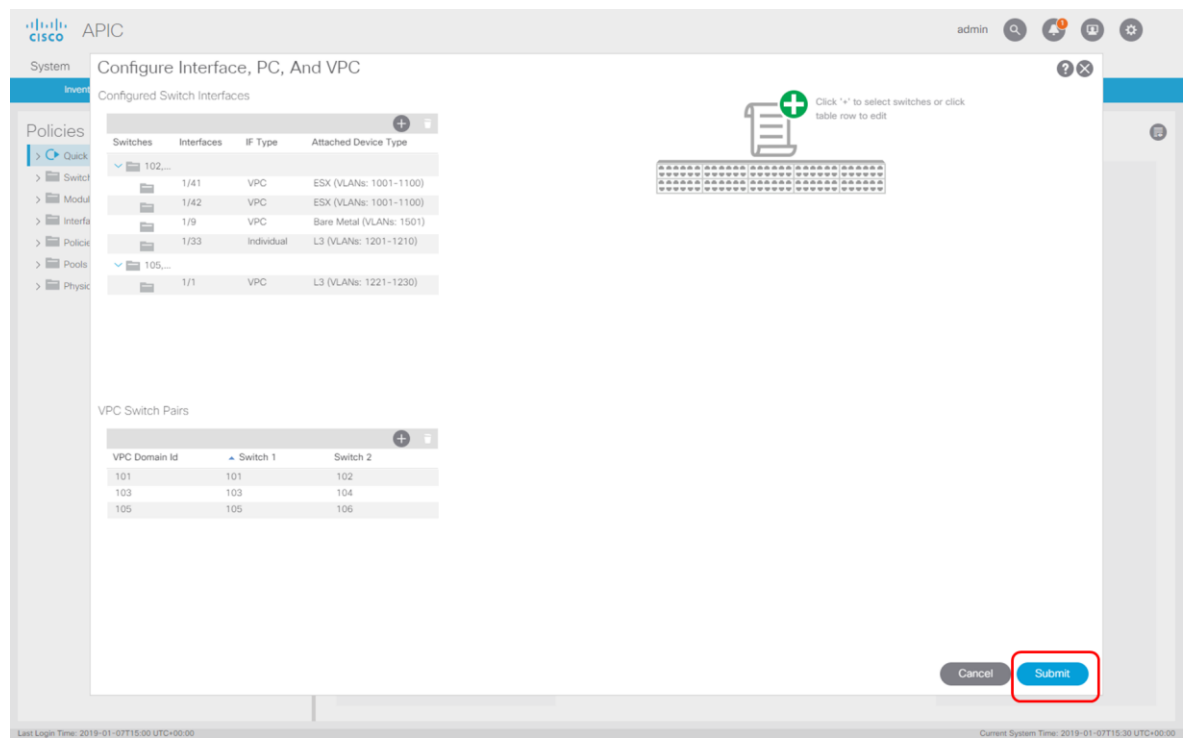
Overview of Created Policies

So, what happened when we have been executing this wizard? The following image shows all policies that have been created or linked to (already existing).

Figure 194. Router3: Policy overview
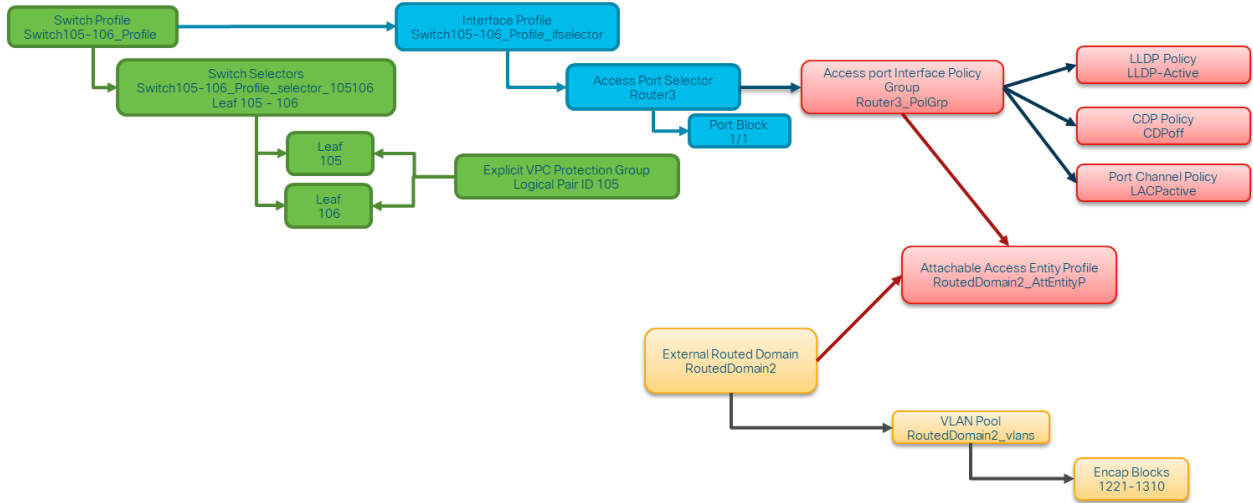
As we have symmetric interfaces in place we only created 1 "Switch Policy" and we have created an "External Routed Domain" called "RoutedDomain1" which is ready to be used in L3out configuration at tenant level in the chapter of this document. The following is an overview of the created or re-used policies in detail.

Figure 195. Router3: Explicit VPC protection group

Figure 196. Router3: Switch Policy

Figure 197. Router3: Interface Profile

Figure 198. Router3: Access port selectors

Figure 199. Router3: VPC Interface Policy Group

Figure 200. Router3: CDP Interface Policy

Figure 201. Router3: LLDP Interface Policy

Figure 202. Router3: Port Channel Policy

Figure 203. Router3: Attachable Access Entity Profile

Figure 204. Router3: VLAN Pool

Figure 205. Router3: External Router Domains

Tenant Configuration

This section describes step-by-step configuration of ACI tenant network with detailed examples. The logical construct of a typical ACI tenant network consists of tenants, VRFs, bridge domains (BD), endpoint groups (EPG), L3outs, L4-7 services insertions, as well as contracts that are applied between EPGs for white-list-based communication policy control.
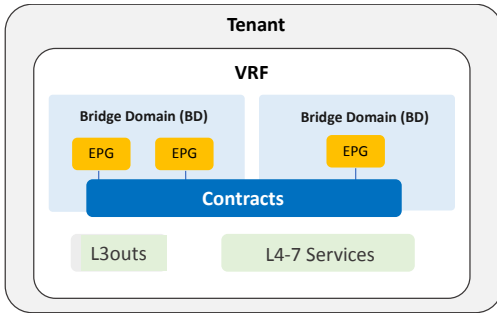
Mirroring the logical construct, the ACI tenant network configuration intuitively includes the following steps:

- Create a tenant space

- In the tenant space, create VRFs

- Create and configure bridge domains (BD) in VRFs

- Create and configure an Application Profile with end point groups (EPGs). Each EPG is associated with a BD in the overlay network construct.

- Create contracts and apply contracts between EPGs to enable white-list-based communication between EPGs.

- Create and configure L3outs

This document uses the 3-tier application shown in the following figure as an example to illustrate the tenant configuration steps.
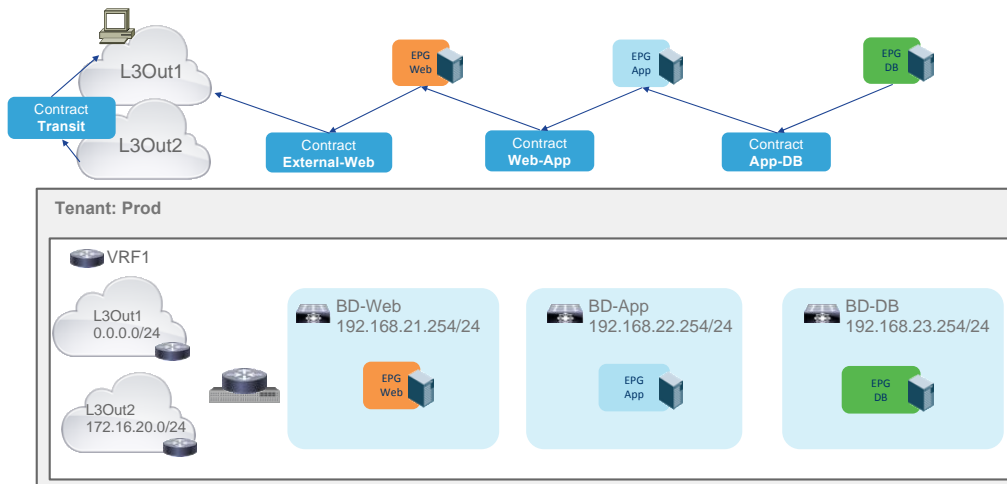
The example uses the following network design that represents a typical practice, but some other design options are discussed as alternatives throughout this session. Their configuration steps are shown as well where it is needed.

- One tenant named Prod

- One VRF in tenant Prod named VRF1

- Three bridge domains in VRF1, named BD-Web, BD-App and BD-DB

- Three EGPs named EPG-Web, EPG-App and EPG-DB, each in the bridge domain that has the corresponding name

- A contract named External-Web to allow the communication between the L3out external EPG and the internal EPG-Web

- A contract named Web-App to allow the communication between EPG-Web and EPG-App

- A contract named App-DB to allow the communication between EPG-App and EPG-DB

- Two L3out, named L3Out1 and L3Out2

The following table shows the IP addressing we use in this example.

Table 3. IP Addressing used in the example

BD	BD Subnet	EPG	VM-Name	VM IP Address
BD-Web	192.168.21.254/24	EPG-Web	VM-Web	192.168.21.11
BD-App	192.168.22.254/24	EPG-App	VM-App	192.168.22.11
BD-DB	192.168.23.254/24	EPG-DB	VM-DB	192.168.23.11
L3Out1			External-Client1	172.16.10.1
L3Out2			External-Client2	172.16.20.1

The rest of this session details step-by-step configuration of the tenant network for the 3-tired application example.

Tenant Configuration

The following sections describe tenant configuration.

Create Tenant

You can create a tenant in the APIC GUI by using the "Create a tenant and VRF" wizard from the Tenant Quick Start menu or by using the "Create Tenant" option under Tenant/All Tenants.

Figure 208. Creating tenant using tenant quick start wizard
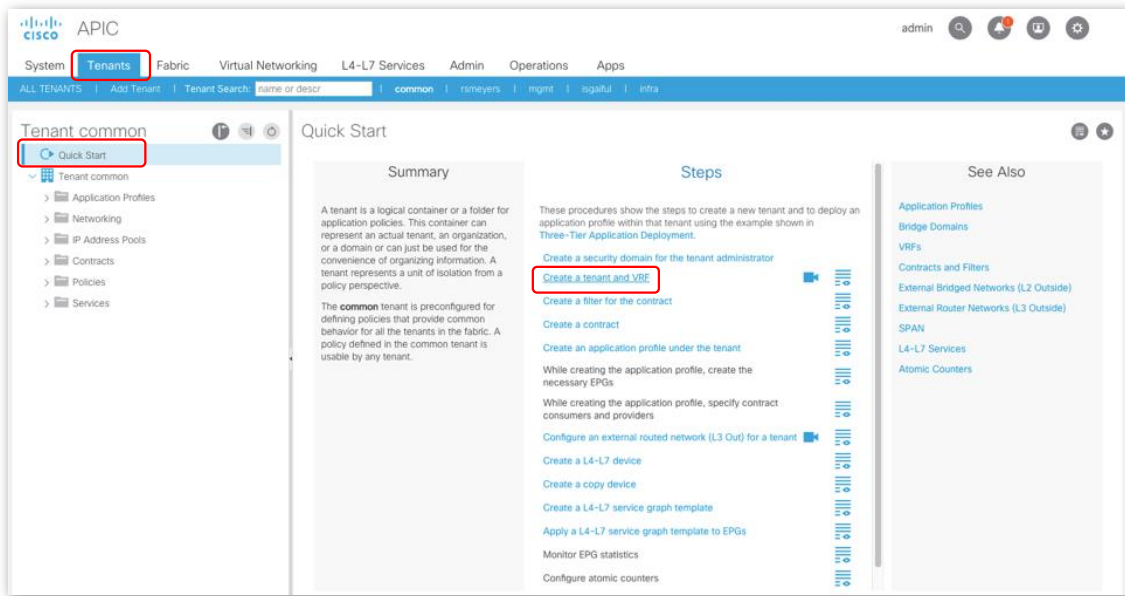
Figure 209. Creating tenant using tenant UI menu

With either of the two methods, you see a pop-up window to create a new tenant and VRF (optional for this step), as shown below.

Figure 210. Creating a new tenant named "Prod"

You need to provide the name of the tenant, and the name of the VRF if choosing to create a VRF in this step. Put in "Prod" as the tenant name, and "VRF1" for the VRF, then click on the "Submit" button. The tenant "Prod" and the VRF "VRF1" are then created.

Figure 211. New tenant "Prod" created with a VRF named "VRF1"
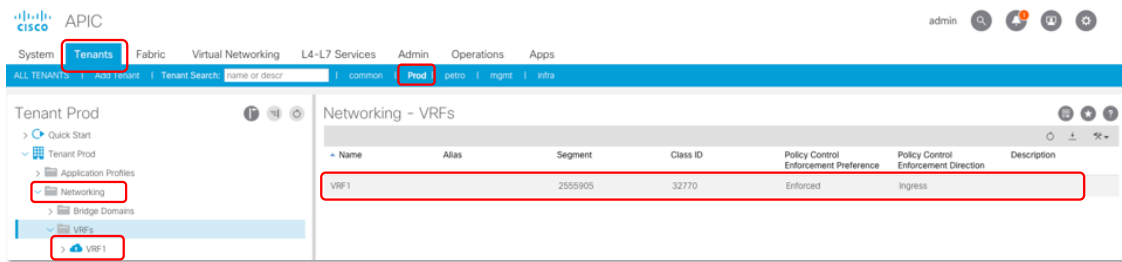
Create VRFs

In this example, we use only one VRF that is created in the previous step of creating the tenant Prod. If needed, more VRFs can be created in the tenant space under Networking > VRFs by right clicking on "VRFs" in the left panel to bring up the pop-up option for "Create VRF". Alternatively, you can use the drop-down menu in the right panel to start creating a new VRF.

Figure 212. Creating VRF (Option 1)
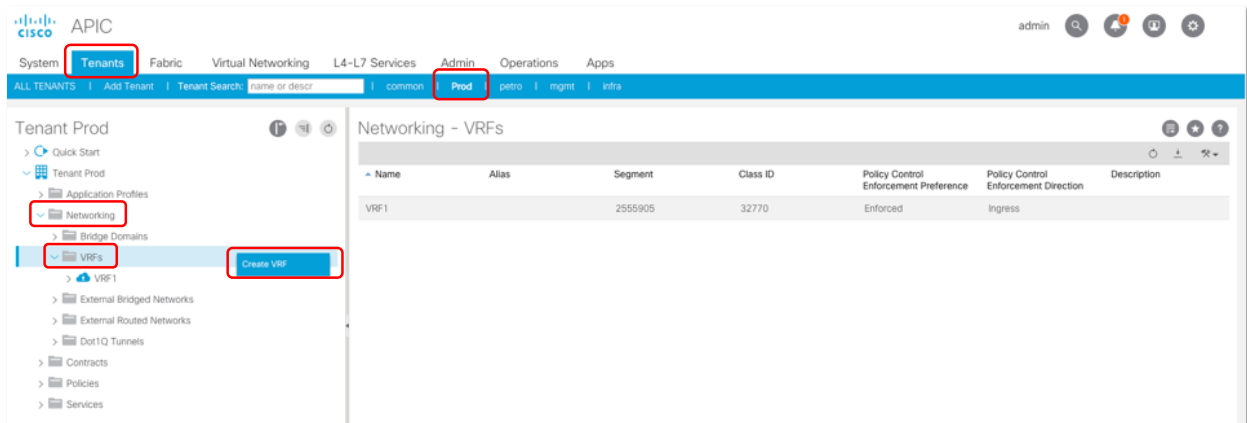
Figure 213. Creating VRF (Option 2)
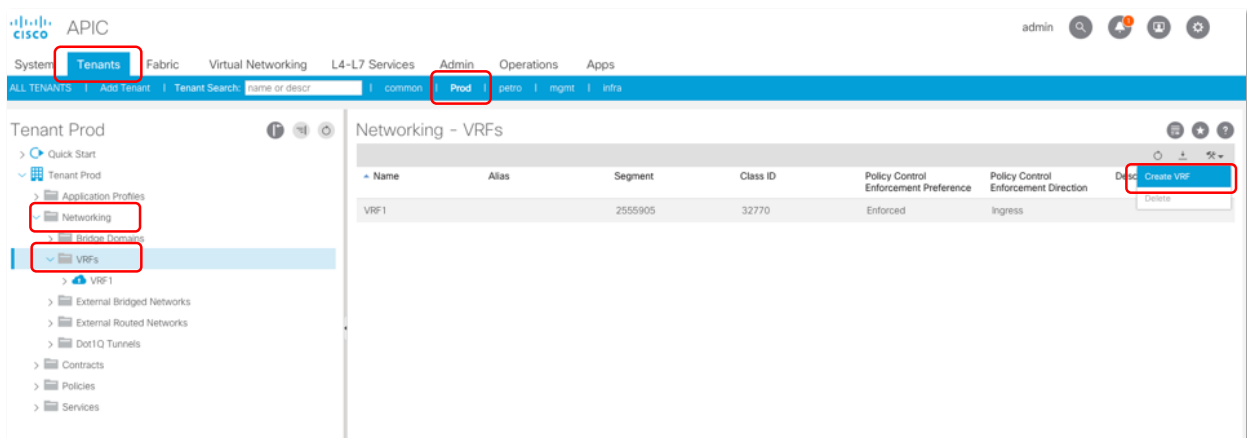
Create Bridge Domains (BD)

Bridge domains (BD) are created under Tenant > Networking > Bridge Domains.

1. You can right click on "Bridge Domains" and chose the pop-up option "Create Bridge Domain" or use the drop down menu in the right panel to select "Create Bridge Domain".

Figure 214. Starting to create bridge domains (Option 1)
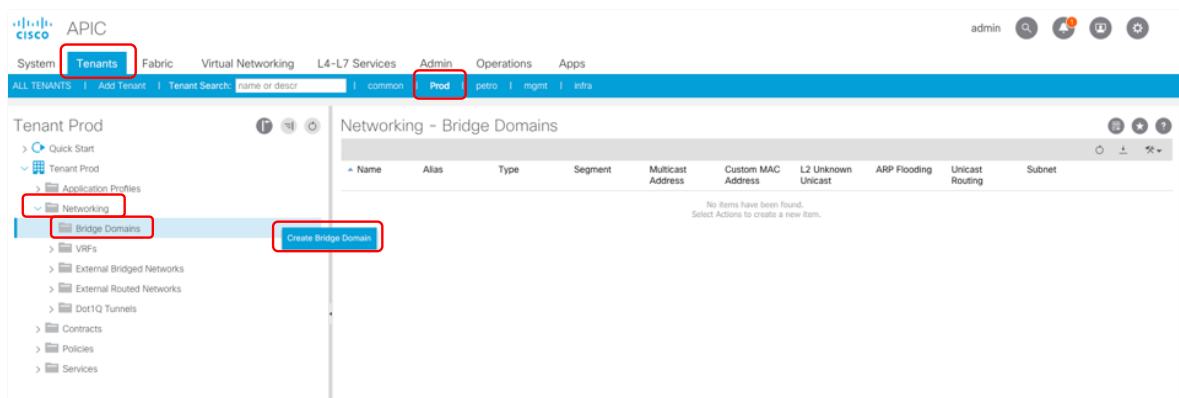
Figure 215. Starting to create bridge domains (Option 2)
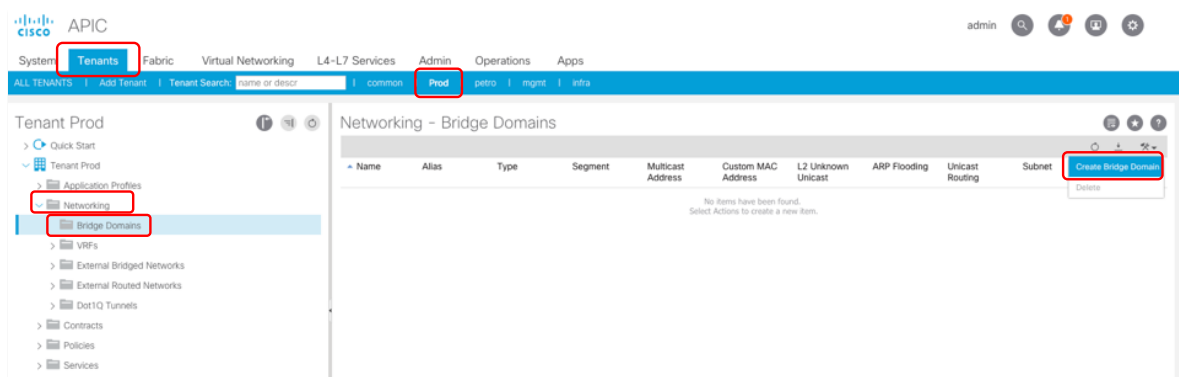
2. In the window that opens, provide the BD information, such as the name of the BD and the VRF that it belongs to. The following figure shows the example to create the BD "BD-Web" in VRF1 of the tenant Prod.

Once you provide the BD name and select the VRF you want it to belong to, you can click on the "Next" button in the window. It will proceed to the next step of BD configuration, the Layer3 Configurations, as shown below. In this example, you can use the default settings for the BD-Web.

Note: Depending on specific network scenarios, some of the settings may need to be changed, for example, ARP flooding may need to be enabled, or Endpoint Dataplane Learning may need to be disabled. But the discussion of how to choose the settings for these configurable features are not in the scope of these documentation. For more information on these features and the best practices of their settings for specific network scenarios, users can refer to the documents in the references of this document.

Figure 216. Creating bridge domain "BD-Web" in VRF1 of tenant Prod

Cisco Systems. All printed copies and duplicate soft copies are considered uncontrolled
and the original online version should be referred to for the latest version.

Page 138 of 215

3. On the next screen, click the "+" sign for "Subnets" to bring out another pop-up window for adding subnets to this BD. You can give the IP subnet 192.168.21.0/24 to BD-Web as its primary subnet with 192.168.21.254 being its IP address.

Figure 217. Starting to add L3 subnet to bridge domain "BD-Web"
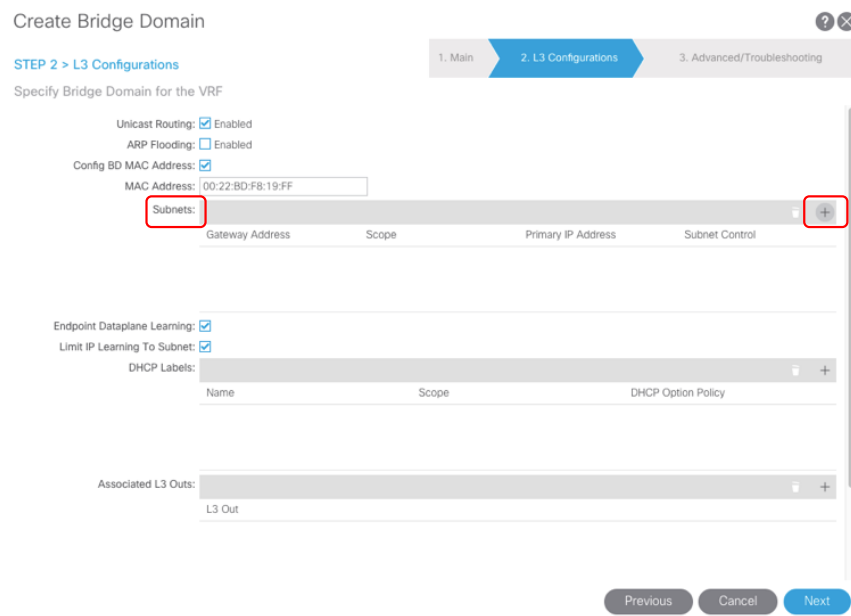
4. Provide the subnet information.

Figure 218. Adding L3 subnet 192.168.21.254/24 to bridge domain "BD-Web"
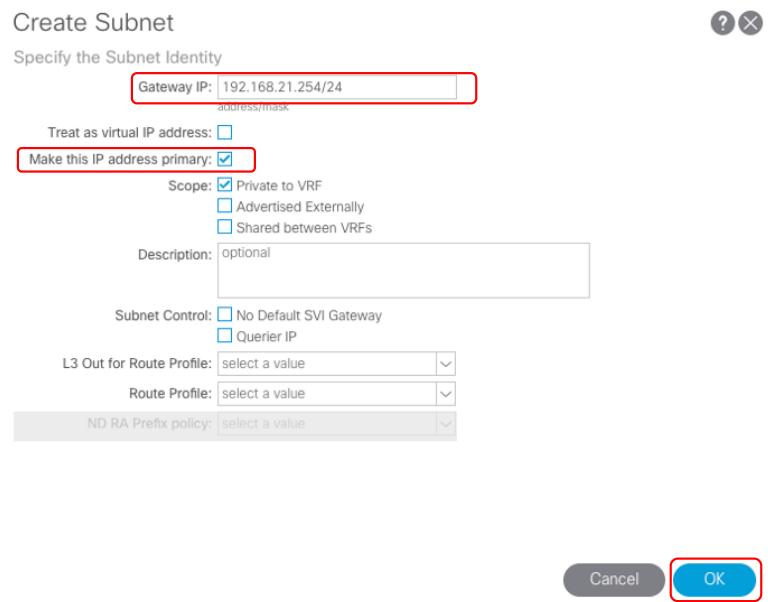
5. Clicking on the "OK" button in the subnet configuration window closes this window out and takes you back to the BD Layer 3 configuration window, with the subnet being added to the list.

Figure 219. Bridge domain L3 configuration after adding the subnet
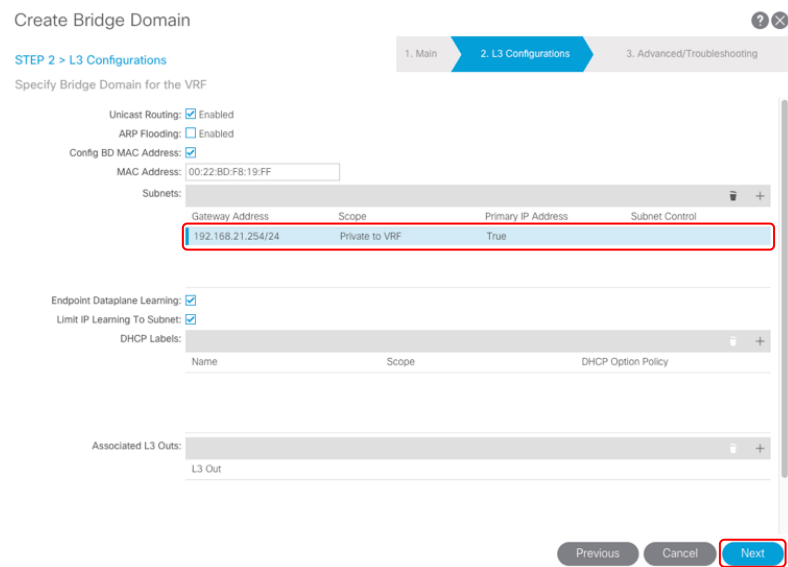
6. The next step in creating and configuring the BD is the Advanced/Troubleshooting. It is optional to configure any of the features listed in this step. In our example, the default settings are used. Clicking the "Finish" button will complete and submit the BD configuration.

Figure 220. Configuring advanced setting for bridge domain
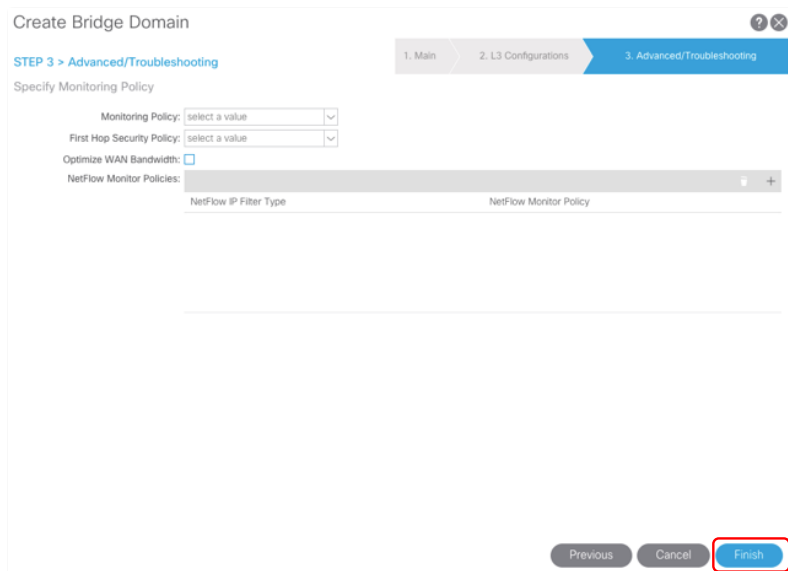
In our example, the same steps of creating BDs are repeated to create the other two BDs, BD-App and BD-DB. At the end of this step, the tenant Prod has three BDs in its VRF1:

Figure 221. Three bridge domains created for the Three-tier application
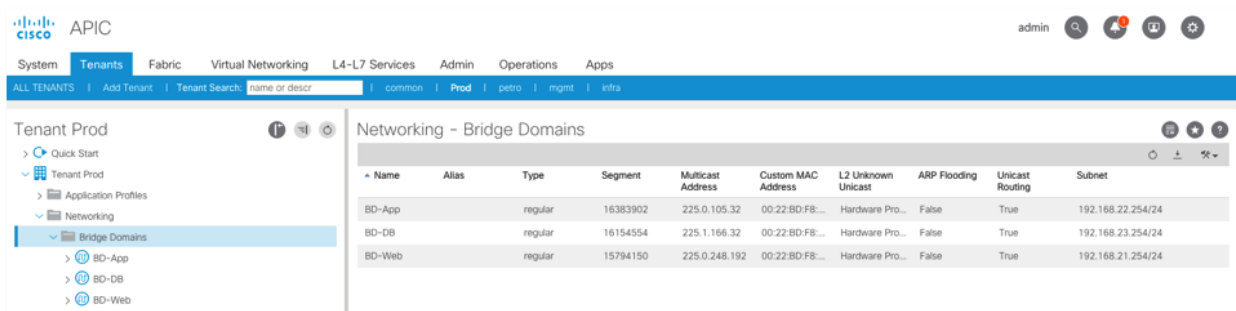
Create Application Profile and End Point Groups (EPGs)

ACI application profile consists multiple end point groups (EPGs). EPG is one of the most important concepts in ACI. An EPG is a group of end points that share the same network policies. EPGs could be defined to reflect application tiering, or simply to reflect the VLAN segments used in traditional network designs. An EGP must belong to one and only one BD while a BD can have multiple EPGs.

If you are migrating from a current traditional VLAN-based network to ACI fabric, and desire to continue using the VLAN model, you can map each VLAN to an EPG that then is associated to a unique BD on the ACI fabric. Our example uses this approach. It provides the an easy migrate path from the current VLAN-based network to ACI fabric. On the other hand, to take the most advantage of the flexibility provided by ACI network construct, multiple EPGs can be in the same BD so that their policy enforcement boundaries are separate by EPGs while sharing the same IP subnet in the BD. This can enable higher network agility, for example, it allows easy re-purposing of endpoints among the EPGs without the need of reconfiguring their IP addresses and gateway routes.

Create the Application Profile

An application profile can be created at Tenant > Application Profiles.

1. You can right click on "Application Profiles" in the left panel to bring out the option to create application profiles or use the drop down menu in the right panel to start creating a new application profile.

Figure 222. Creating application profile (Option 1)
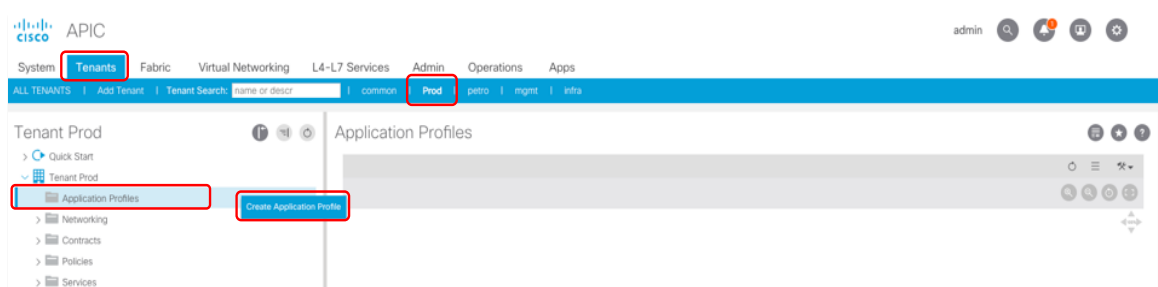
Figure 223. Creating application profile (Option 2)
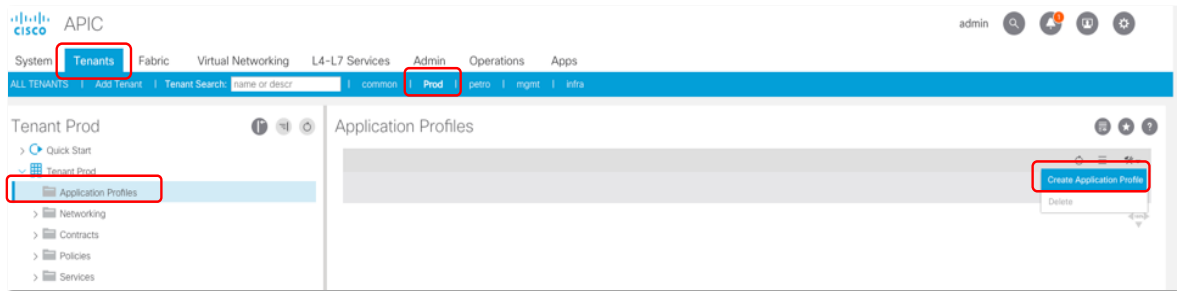
2. Either of the two methods brings a pop-up window for you to define the application profile. You need to provide the name of the application profile, then click on "Submit" to complete the creation step. In this example, an application profile with the name "APP1" is created.

Figure 224. Creating Application Profile "APP1"
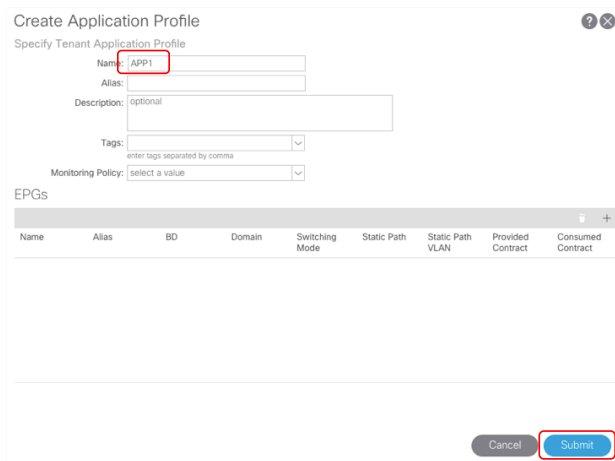
Create EPGs Under the Application Profile

EPGs can be created on the ACI UI at Tenant > Application Profiles > a specific Application.

1. Right clicking on the application profile name under Application Profiles brings up a drop-down menu, you can select "Create Application EPG" on the top to start the EPG configuration. Alternatively, you can use the drop-down menu in the right panel to create an application EPG.

Figure 225. Starting to create EPG (Option 1)
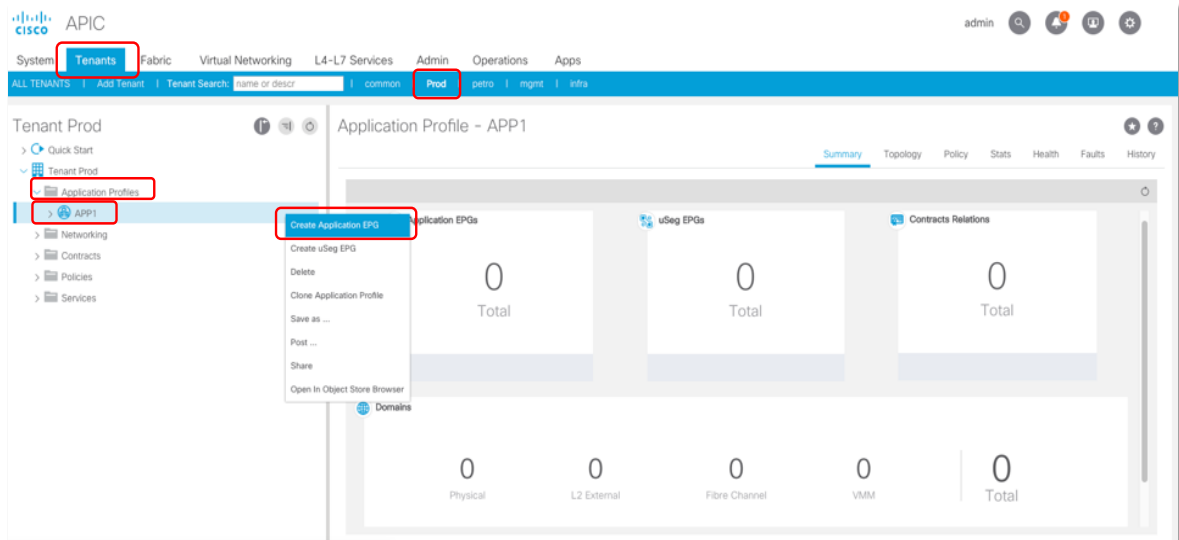
Figure 226. Starting to create EPG (Option 2)
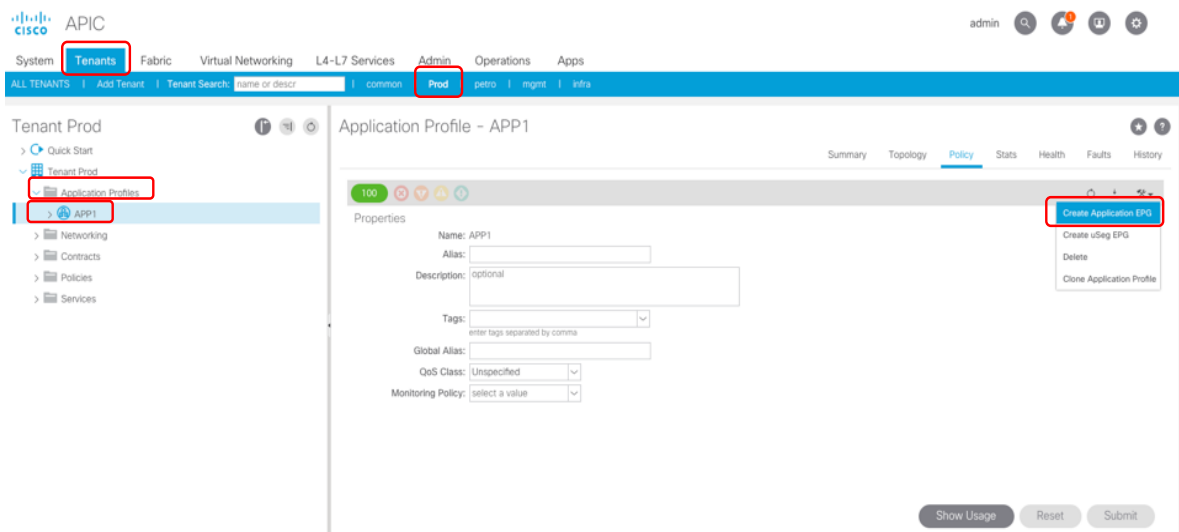
2. Either of the two methods brings out a pop-up window for EPG configuration. In the window, you need to provide the name of the EPG and the BD you want it to belong to. In this example, we created EPG-Web in BD-Web n the tenant Prod.

Figure 227. Creating EPG "EPG-Web"

3. The same steps are repeated to create the other two EPGs, EPG-App and EPG-DB. EPG-App is in BD-App while EPG-DB is in BD-DB. At the end of this step, the application profile "APP1" has 3 EPGs created.

Figure 228. Three EPGs in the application "APP1"

Associate EPGs with Physical or VMM Domains

ACI EPGs need to be associated with physical domains and/or VMM domains so that endpoints can be either statically binded to EPGs when they are connected to the ACI leaf switch ports, or dynamically classified to EPGs through ACI and VMM integration. Via physical domain association, an EPG can have static bindings to leaf ports that have bare metal servers or virtual machines connected to. For virtualized endpoints, ACI provides seamless integration with Virtual Machine Managers (VMM) to automate the virtual network provisioning in the VMM domain. In our example, we use VMM domain from EPG-Web and EPG-App, and both VMM domain and physical domain for EPG-DB.

1. Associate EPGs with VMM Domains

   A VMM domain can be associated with an EPG on the APIC UI under through Tenant > Application Profile > EPG > Domain. You can right click on "Domains (VMs and Bare Metals)" and choose the option for "Add VMM Domain association" in the pop-up menu. Alternatively, you can go to the work panel of the EPG and choose the option for "Add VMM Domain Association" from the drop-down menu under Policy > General.

   Figure 229. Starting to add VMM domain association to EPG-Web (Option 1)
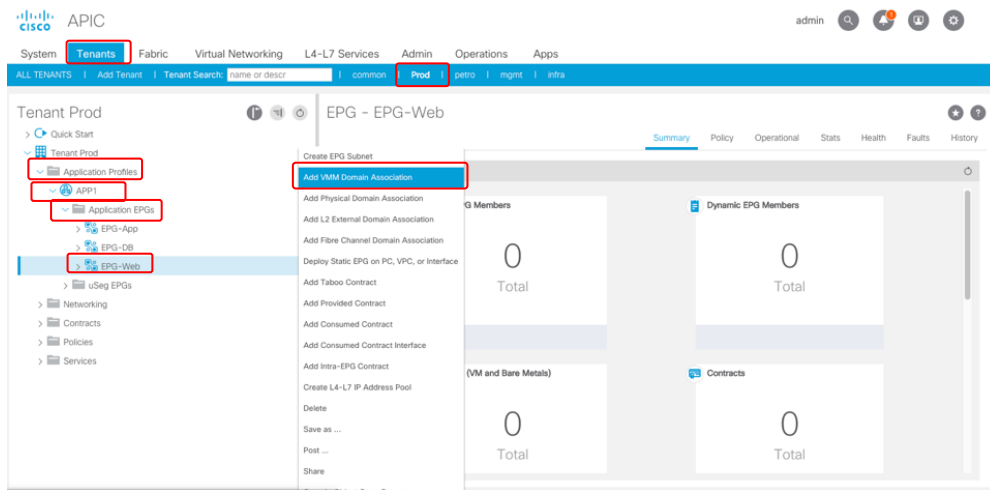
   Figure 230. Starting to add VMM domain association to EPG-Web (Option 2)
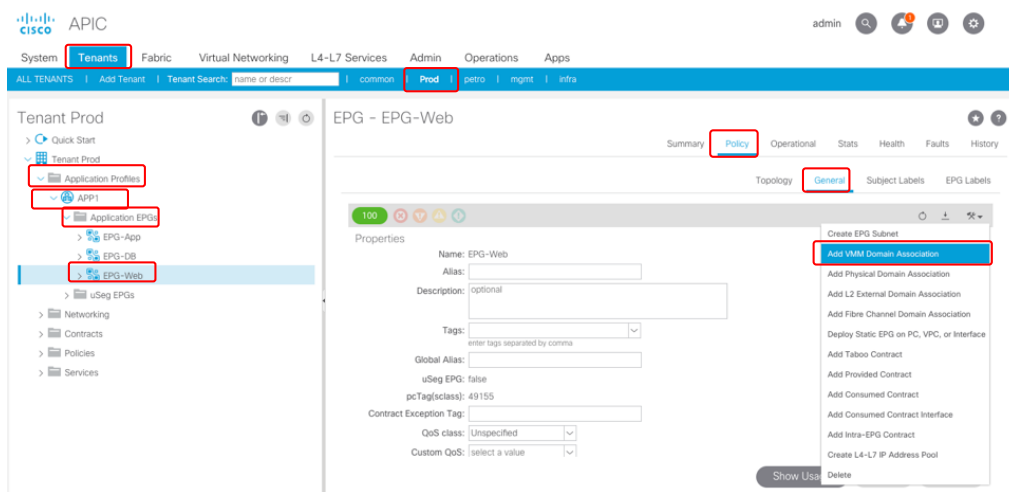
2. With either of the two methods, you get a pop-up window to configure the VMM domain association. In the pop-up window you need to choose the desired VMM domain (in our example, it is ACI_VDS) and the settings, such as deployment and resolution immediacy, and VLAN mode.

   There are two VLAN modes, dynamic and static. If EPG classification can be decoupled from the VLAN ID assignment, you can use the dynamic VLAN mode to allow ACI to pick a VLAN ID from the VLAN pool of the VMM domain. ACI will then provision a port-group on the VMM vswitch using the same VLAN ID. The dynamic VLAN mode provides the most flexibility and simplicity of VLAN ID management. In our example, we use dynamic VLAN mode. After clicking on the submit button, EPG-Web now is associated with VMM domain ACI_VDS. Repeat the same steps to associate EPG-App and and EPG-DB with the VMM domain ACI_VDS.

Figure 231. Adding VMM domain association

3. In case that you need to statically assign a specific VLAN ID to an EPG, you can use the static VLAN mode when associating a VMM domain with the EPG. In the static VLAN mode, you can specify the VLAN ID you want to be used for this EPG. This VLAN ID needs to be in the VLAN pool of the VMM domain. Figure 232 in below shows an example of creating VMM domain association with the static VLAN mode and using VLAN 100 for the EPG.

Figure 232. Adding VMM domain association with static VLAN mode

With a VMWare vCenter VMM domain, once an EGP is associated with the VMM domain, a corresponding port-group will be created on the vCenter VDS and ready to be used for virtual machine deployments.

4. At this stage, if we examine the networking configuration on the vCenter, you can see that on the VDS created via the ACI VMM integration, three port-groups are created correspondingly for the three EPGs.

Figure 233. vCenter port-groups created by the ACI VMM integration
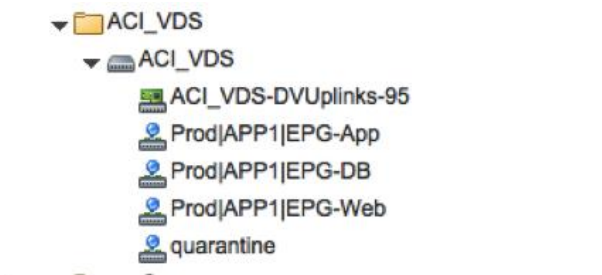
Associate EPGs with Physical Domains

Now add the physical domain and static binding to EPG-DB.

Physical domains are added to an EPG at Tenant > Application Profile > Application EPG > Domains (VMs and Bare Metals).

1. Right click "Domains (VMs and Bare Metals)", you choose the option of "Add Physical Domain Association" in the pop-up menu.

Figure 234. Starting to add physical domain association

Cisco Systems. All printed copies and duplicate soft copies are considered uncontrolled
and the original online version should be referred to for the latest version.

Page 147 of 215
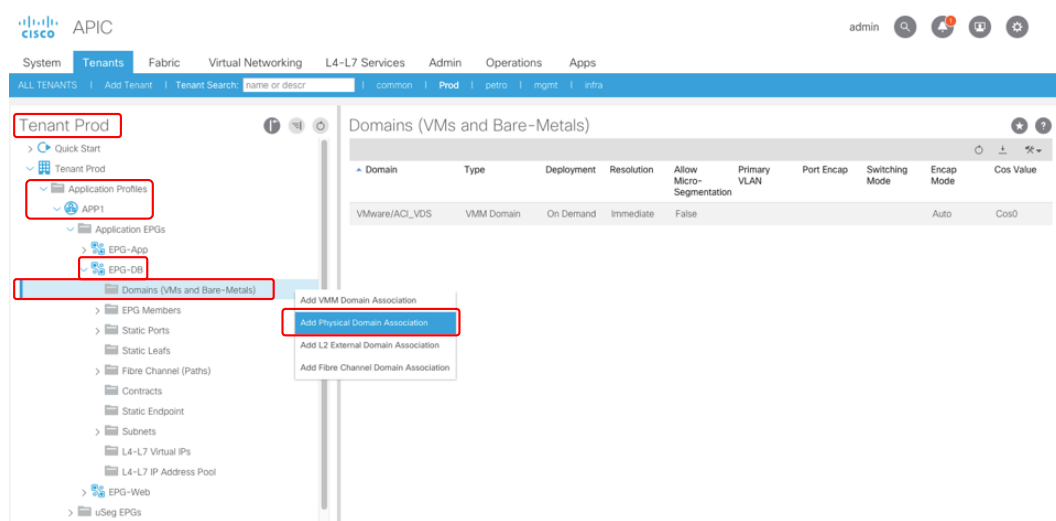
2. You then get a pop-up window to select a physical domain. The physical domain named "BareMetal" is selected. Clicking on the Submit button completes the association of this physical domain with EPG-DB.

Figure 235. Adding physical domain association
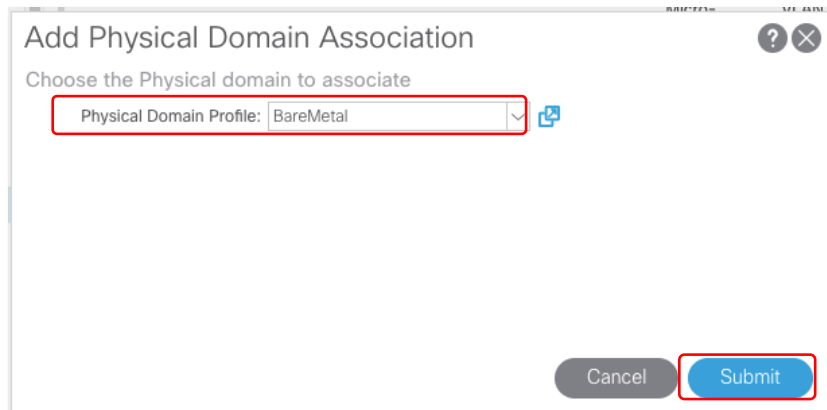
Add Static Binding for EPGs in Physical Domain

Leaf switch ports in the physical domain can be associated with EPGs through static binding.

1. To add static binding to the EPG, you can go to Tenant > Application Profiles > Application EPGs > Static Ports, then right click on "Static Ports", and select the pop-up option for "Deploy Static EPG on PC, VPC and Interface", you get the pop-up window to select the port, PC or VPC.

   In the window, you need to choose the path type, the path and the VLAN used for this path. In our example, we added VPC Server1_PolGrp as the path using VLAN 1501 for encap. Note that this VLAN ID needs to be in the VLAN pool associated with Server1_PolGrp through its AEP. Then click on the Submit button, this path binding is created.

Figure 236. Adding static port to EPG
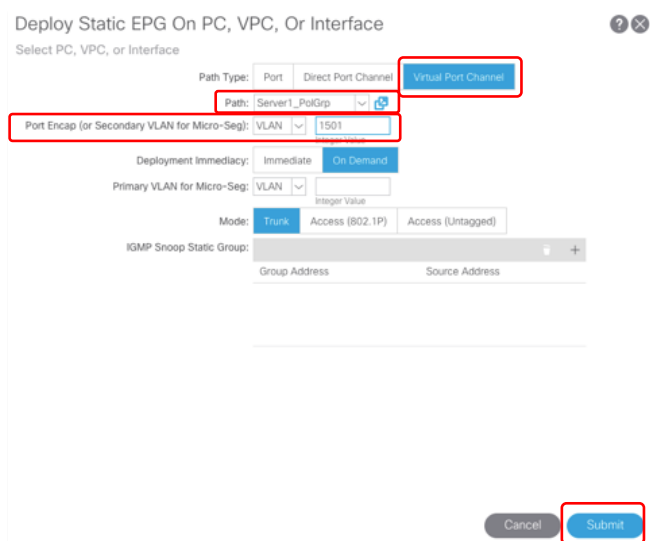
2. You can see the path is added to EPG-DB.

Figure 237. EPG with static ports added
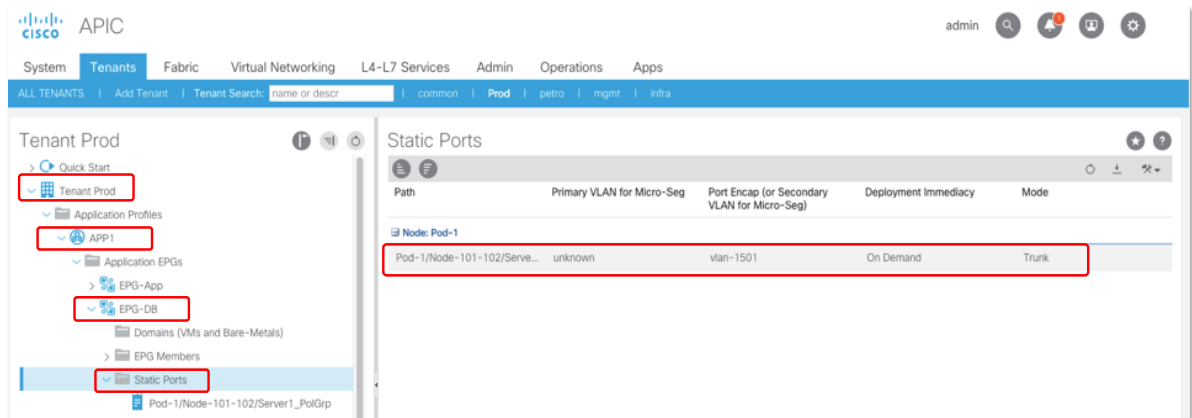
An alternative way to add static bindings to an EPG is to add the EPG to the Attachable Access Entity Profile (AEP) that has all the ports that you want to add to the EPG. This method is more convenient to add a large number of static ports to an EPG.

1. Navigate to Fabric > Access Policies > Policies > Global > Attachable Access Entity Profile, in the work panel on the right side, click on the "+" sign for "Application EPGs".

Figure 238. Starting to add EPGs to AEP
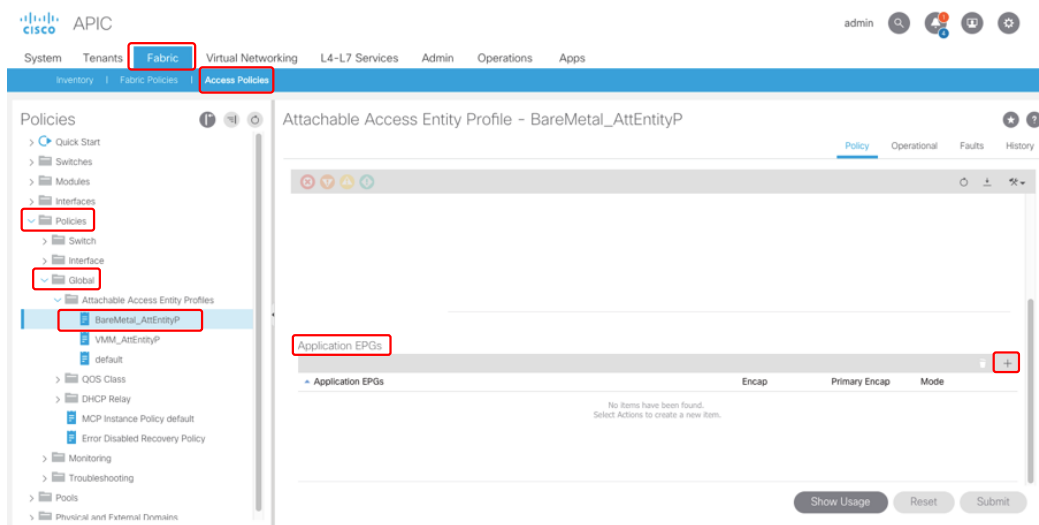
2. It will trigger the prompts for you to specify the EPGs you want to add to this AEP, including the tenant, application profile, EPG and Encap VLAN. Once you provide all the required input, click on "Update" to complete the process of adding the EPG to the AEP. You will see the EPG being added and listed in the AEP.

Figure 239. Adding EPGs to AEP
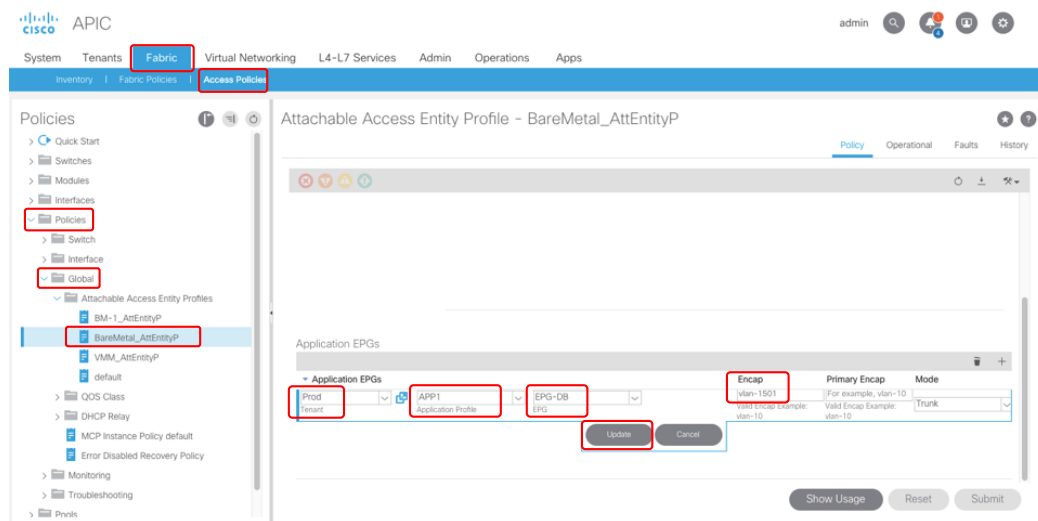
3. Effectively, the EPG is associated with all the ports in the AEP with the specified VLAN encap.

Figure 240. AEP with EPG added
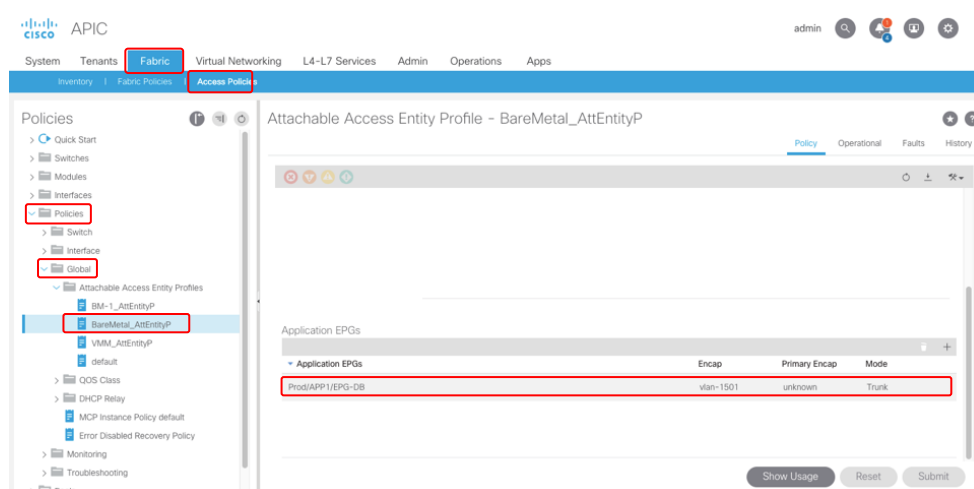
Create and Apply Contracts to EPGs

With its white-list model, by default ACI only allows communication between EPGs that is explicitly permitted with contracts. Therefore, the natural next step after creating EPGs is to create and apply contracts.

Create Filters

A contract can consist of 1 or more subjects that each contains 1 or more filters. Filters are used to classify traffic based on its Layer 2 to Layer 4 attributes.

1. Start with creating filters. Navigate to Tenant > Contracts > Filters. Right clicking on "Filters" in the left panel brings about an option for "Create Filter" or you can use the drop-down menu in the right view panel.

Figure 241. Starting to create contract filter (Option 1)
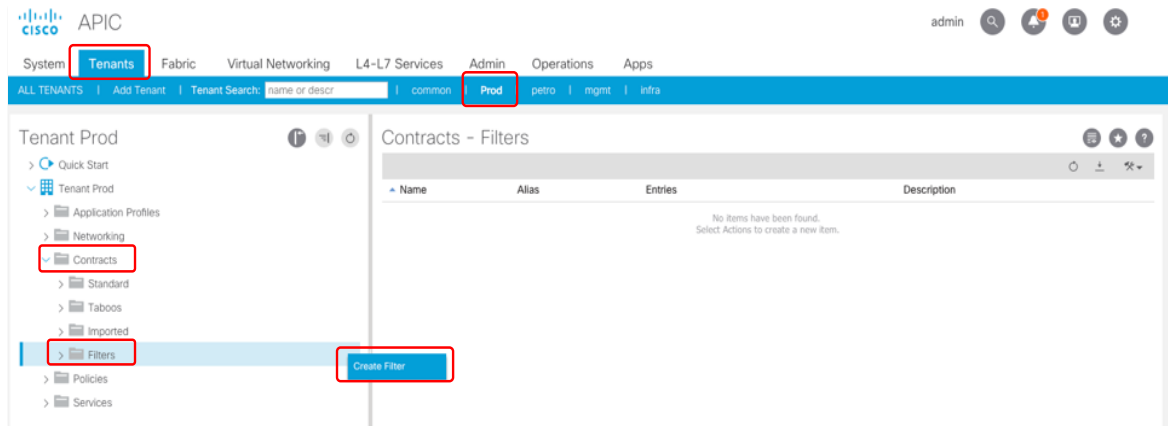
Figure 242. Starting to create contract filter (Option 2)
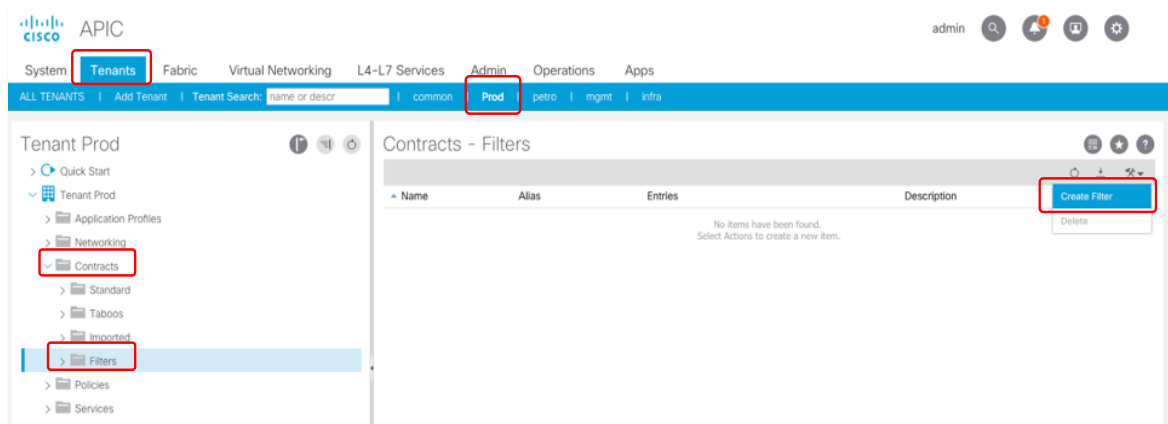
2. In the pop-up window to create filter, you need to provide a name for the filter, then add entries. Each entry defines a type of L2-4 traffic with the attributes such as Ethernet type, IP protocol, source port range, destination port range, etc. In this example, we create a filter named "fltr-permit-all" that permits all Ethernet traffic.

Figure 243. Example of creating a filter to permit all traffic (Step 1)

Figure 244. Example of creating a filter to permit all traffic (Step 2)

3. For simplification purpose, we use this permit-all filter for our contracts throughout our example. In production environments, contracts with filters to allow specific traffic types are often used.

Figure 245. Example of creating a filter to permit all traffic (finished view)
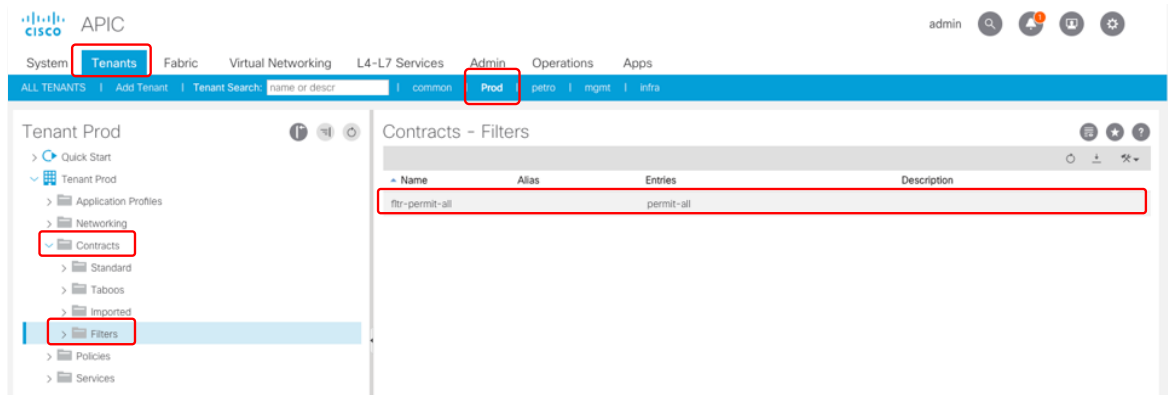
The following figures provides more example filters to match ICMP, HTTPS, or SSH traffic.

Figure 246. Example contract filter to permit ICMP traffic
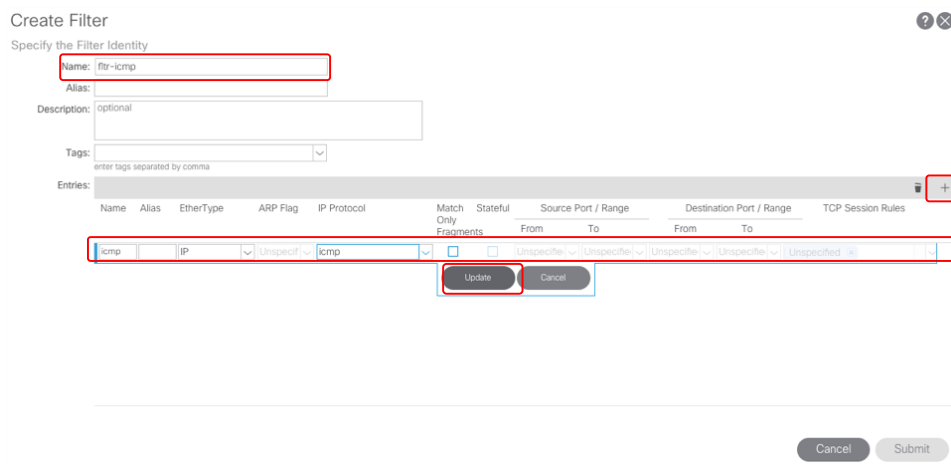
Figure 247. Example contract filter to permit HTTPS traffic
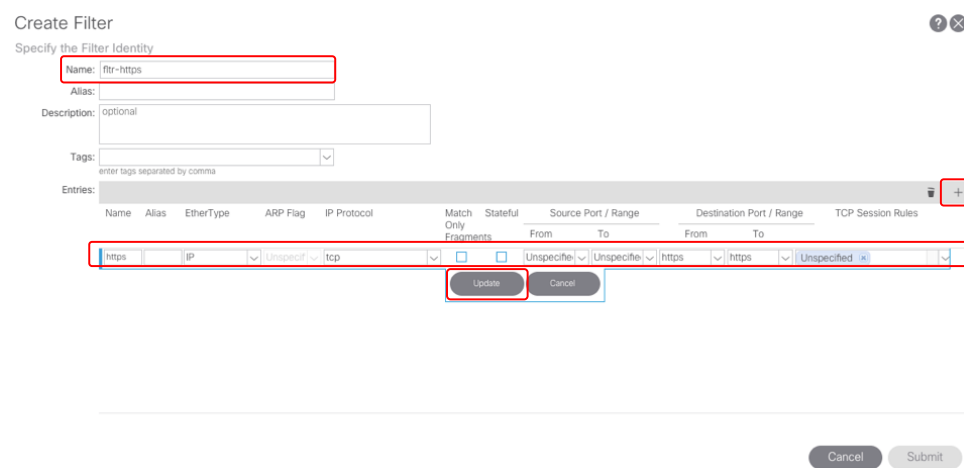
Figure 248. Example contract filter to permit SSH traffic

Create Contracts

Contracts are created at Tenant -> Contract.

1. You can right click on "Contracts" and choose "Create Contract" from the pop-up option or use the drop-down menu in the right view panel to create contract.

Figure 249. Starting to create contract (Option 1)
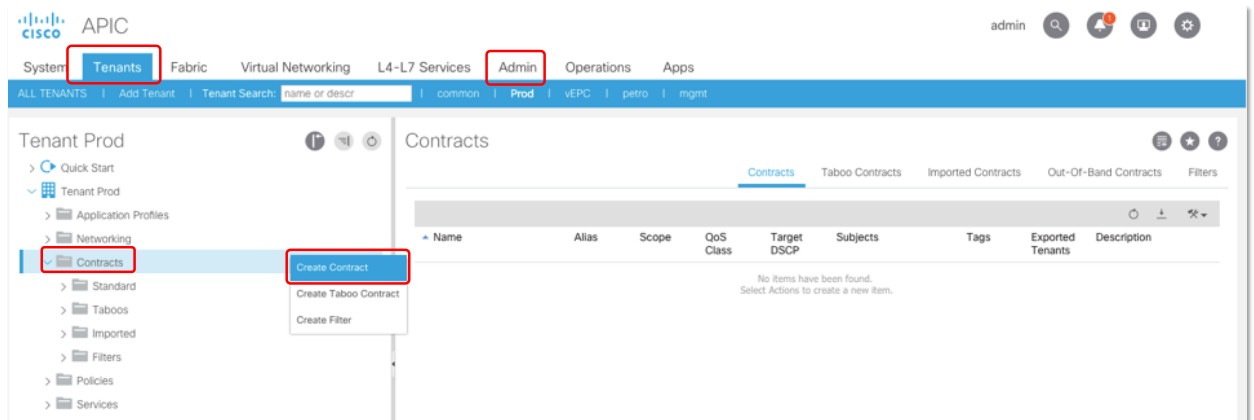
Figure 250. Starting to create contract (Option 2)
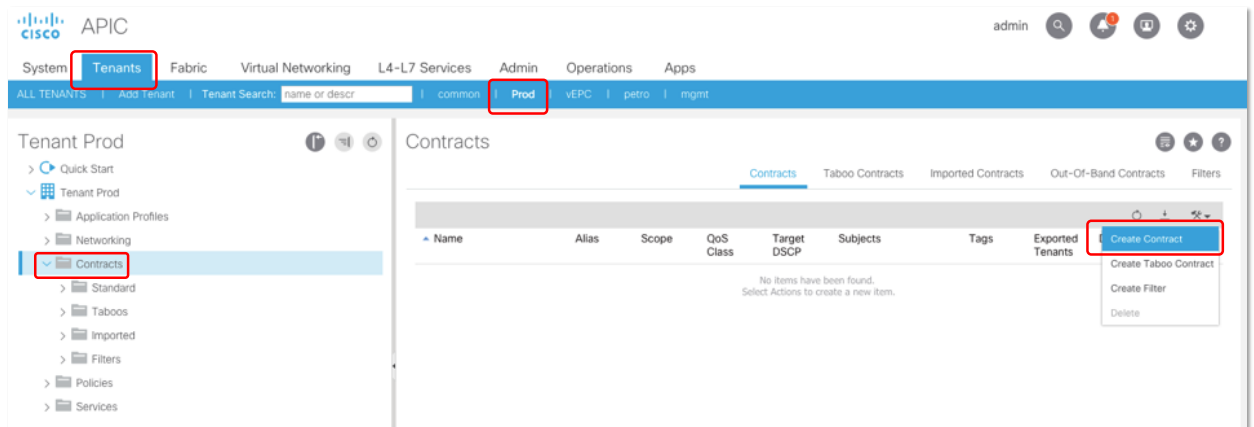
2. With either of the two methods, you get a pop-up window shown below to provide the name and the scope of the contract, as well as the option to add subjects. We then create the contract Web-App with VRF scope as it is a contract to control the communication between two EPGs in a same VRF.

Figure 251. Creating a contract

3. A sub-step of creating a contract is to create contract subject. To start this step, click on the "+" button for adding subjects in the "Create Contract" window. You are directed to a new window to create a subject.

Figure 252. Starting to create a contract subject

4. In this window, you need to add filters to the subject. In this example, we use the filter fltr-permit-all in VRF Prod to create subject sbjct-permit-all in the contract Web-App.

Figure 253. Adding filters to a subject

5. Once you provide the needed input for the filter, including the filter name, description, action and priority, you can click on "Update" to have the filer added to the subject. After all the needed filters are added, you can click on the "OK" button to complete the subject configuration and return to the contract window.

Figure 254. Completing a subject configuration

You can add multiple subjects to a contract.

6. After adding all needed subjects, you can submit the contract configuration by clicking on the " Submit" button. It will close out the contract creation window as well.

Figure 255. Submitting a contract configuration

7. Follow the same steps, we create the contract App-DB as well. The contracts show up under Tenant > Contracts> Standard.

Figure 256. Contracts created for the example three-tier application

Apply contracts to EPGs

The next step now is to apply the contract between the EPGs. When a contract is used to allow communication between two EPGs, one of the EPG is a provider of the contract whereas the other is the consumer of the contract. Therefore, when applying a contract to an EPG, you need to specify if the contract is provided or consumed by this EPG.

In our example, the contracts are defined and applied as below:

Table 4. IP Addressing

Contract Name	Contract Provider	Contract Consumer
Web-App	EPG-App	EPG-Web

App-DB	EPG-DB	EPG-App

1. To add a provided contract to an EPG, you go to Tenant > Application Profile > Application EPG -> Contract, then right click on Contract to choose the option of "Add Provided Contract" or use the drop-down menu in the right panel to choose "Add Provided Contract".

Figure 257. Adding a provided contract (Option 1)
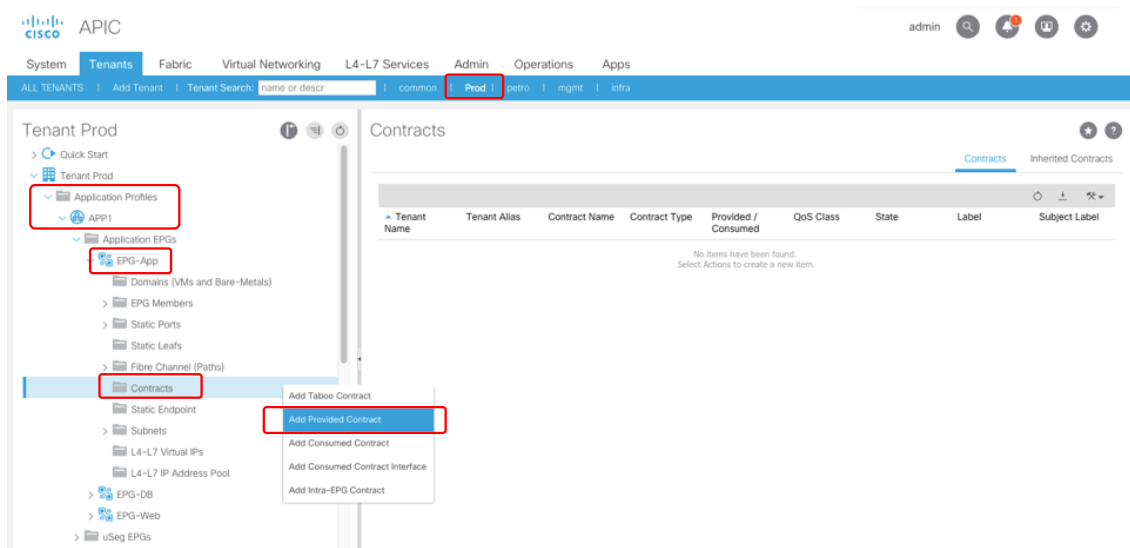
Figure 258. Adding a provided contract (Option 2)
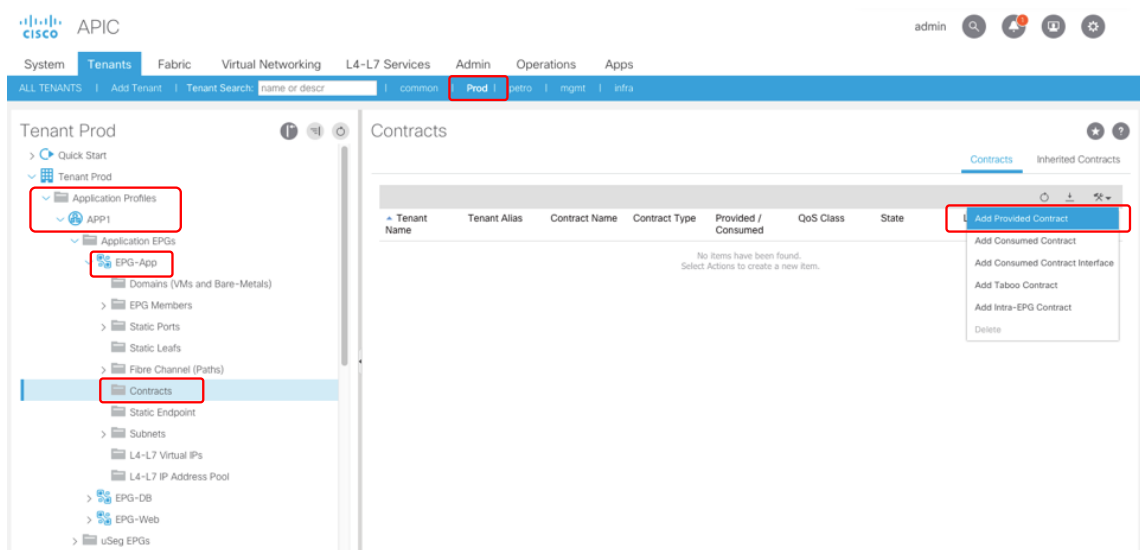
2. Either of the methods gives you the pop-up window to choose a contract. The following example adds the contract Web-App as a provided contract to EPG-App.

Figure 259. Adding provided contract to EPG

3. Choose the contract Web-App, then click on Submit. It finishes the process of adding the contract Web-App as a provided contract to EPG-App.

Figure 260. Provided contract added to EPG-App

4. Next, apply the contract Web-App as a consumed contract to EPG-Web. Following the same work flow as adding a provided contract, add a consumed contract to an EPG. In the example captured in the figures, the contract Web-App is added as a consumed contract to EPG-Web.

Figure 261. Starting to add a consumed contract to EPG-Web

Figure 262. Adding a consumed contract to EPG-Web

Figure 263. Consumed contract Web-App is added to EPG-Web

Repeat the same steps to add the contract App-DB as a provided contract to EPG-DB and a consumed contract to EPG-App.

5. Now, you have an application network profile for the three-tier App1 deployed with the topology shown in the following figure.

Figure 264. Logical topology of App1
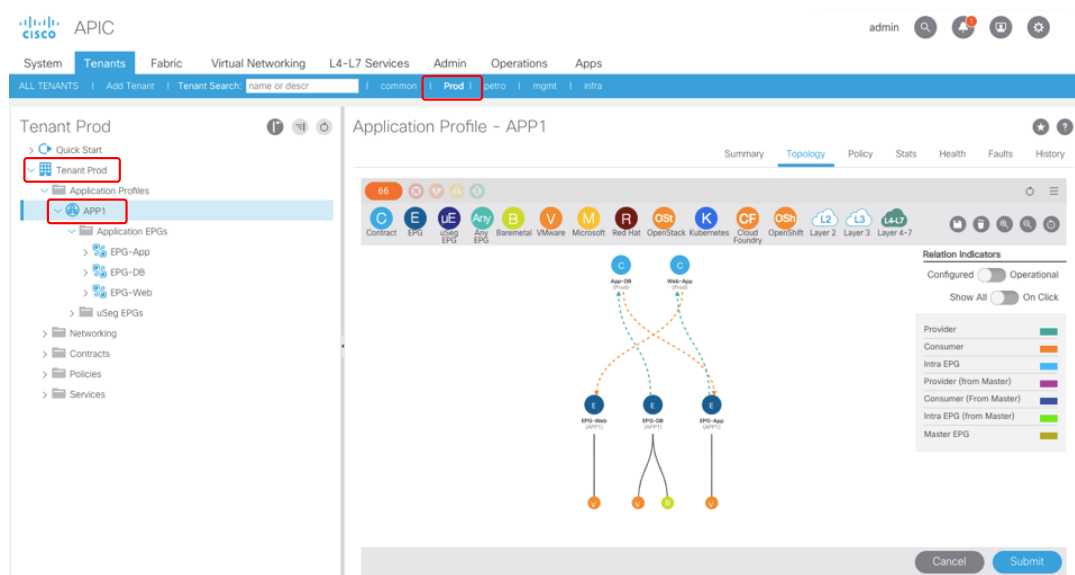
With this network profile, the communication between EPG-Web and EPG-App is enforced by the contract Web-App, whereas the communication between EPG-App and EPG-DB is enforced by the contract App-DB.

Alternative Options for Security Policy Configuration

Our example leverages the default behavior of ACI white-list-based policy model to govern the east-west communication among application EPGs by using contracts. It offers a way to best control and enforce security policies among EPGs. ACI also provide a few alternative ways to provide simpler policy configuration by either reducing the requirements of inter-EPG contracts or simplifying contract relationship when a set of common security policies are applied to multiple EPGs. These options can be useful for the cases where the interdependency among application tiers are not clearly identified, or the network is designed using the traditional VLAN-based model. A design with VLAN-based EPGs is often used when migrating existing applications from a traditional data center network onto ACI fabric. In this migration, you can use VLAN IDs as EPG classification criteria, BD and EPG are simply mapping to existing subnets and VLANs. You can gradually port your current VLAN based network design and operation to the ACI fabric so that it is working in a familiar, well understood, and classical manner. Since most traditional networks have east-west communication being mostly open either among all VLANs or a selected set of VLAN, when migrating the same design onto ACI fabric, the alternative contract simplification methods can be useful to get an easier migration path. In the rest of this section, we use examples to explain how to configure some of these options, including:

■ Unenforced VRF

■ vzAny

■ Preferred Group

Unenforced VRF

By default, security policies are enforced among EPGs in a VRF, but this enforcement can be disabled for the entire VRF. When it is disabled for a VRF, all EPGs within this VRF can communicate with one another without the requirements of contracts. We disable policy control enforcement for VRF1. It is done under Tenant > VRF, on the right panel, under Policy, choose "Unenforced" for "Policy Control Enforcement Preference".

Figure 265. Disabling policy control enforcement in VRF1
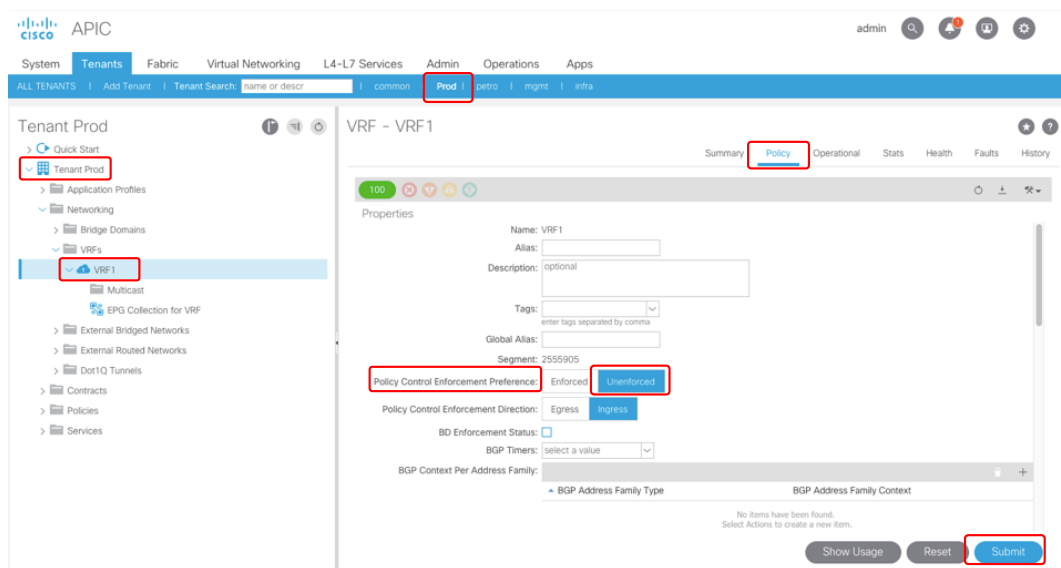
vzAny

vzAny, referred to in the Cisco ACI GUI as "EPG collection for context," is a special object that represents all EPGs associated with a given VRF instance, including the Layer 3 external EPG (L3out EPG). This concept is useful when a configuration has contract rules that need to be applied across all the EPGs in the same VRF.

vzAny is a per-VRF object. You can apply a contract to vzAny of a VRF by adding it as a provided contract or a consumed contract to the vzAny. When a contract is applied to the vzAny, all the EPGs in the VRF assume the configured role of the contract in terms of policy enforcement based on the contract. To add contracts to the vzAny, you go to Tenant > Networking > VRF > EPG Collection for VRF, under Policy > General in the right panel, you can add provided contracts or consumed contracts. In our example, we make the VRF1 vzAny as both the provider and the consumer of our permit-all contract, which effectively allows open communication among all our EPGs in VRF1.

Figure 266. Contracts applied to VRF1 vzAny
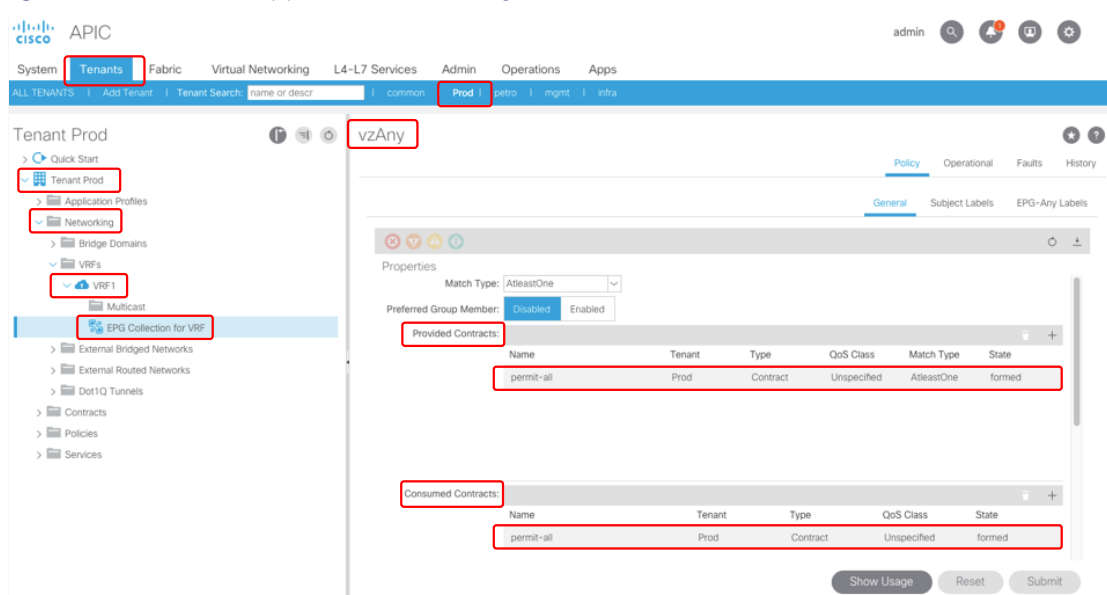
vzAny is also an effective tool to reduce TCAM resource consumption for a design that has a set of common policies applied to all the EPGs in a VRF. For example, you add a set of management tools to our example, and

put them in a new EPG named EPG-Mgmt. Yu need all the tools to communicate with all the three EPGs in our example, so you create a contract named Contract-Mgmt, and you want to make the three application EPGs as the providers of the contract, and the EPG-Mgmt as the consumer. Without using vzAny, you need apply the contract configuration to all the EPGs, and the contract rules will be programmed three times in a leaf switch TCAM table. Now consider using vzAny in VRF1. You can make the VRF1 vzAny as the provider of Contract-mgmt while EPG-Mgmt as the consumer. It automatically make all the EPGs in VRF1 as the provider of Contract-mgmt. You get the same communication policy relationship, but the contract rules are only programmed once in the leaf TCAM table. The TCAM consumption is reduced to 1/3 of what is needed without vzAny. The figure below shows the comparison of with and without using vzAny. In production, the number of EPGs in a VRF can be much greater than 3. To apply this kind of common contract to all the EPGs by using vzAny, you reduce the TCAM utilization to 1/N where N is the number of EPGs in the VRF. Figure 267 shows the comparison between applying contracts to individual EPGs and using vzAny.

Figure 267. Comparison of individual contracts and use of vzAny
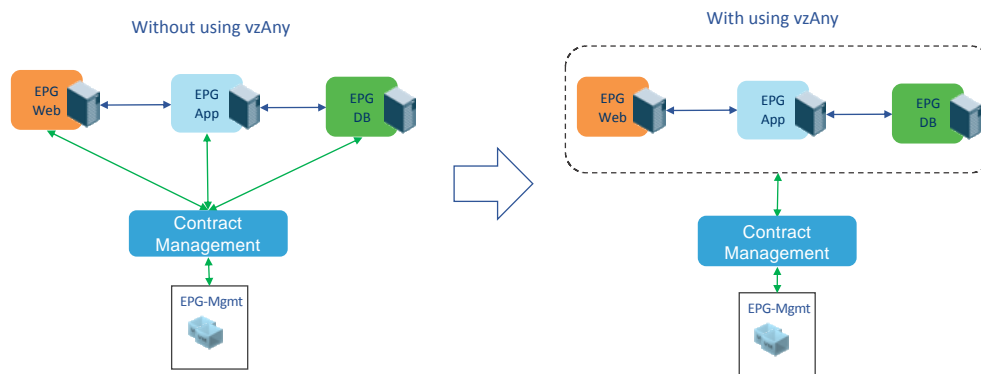
For more details about vzAny, please refer to the following document:
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/kb/b_KB_Use_vzAny_to_Automatically ApplyCommunicationRules_toEPGs.html

Note: When using vzAny with shared services contracts, vzAny is supported only as a shared services consumer, not as a shared services provider.

Contract Preferred Group

The contract preferred group feature enables simpler control of communication between EPGs in a VRF. If most of the EPGs in the VRF should have open communication, but a few should only have limited communication with the other EPGs, you can configure a combination of a contract preferred group and contracts with filters to control inter-EPG communication precisely.

When contract preferred group is enabled for a VRF, each EPG in the VRF is either included or excluded in the preferred group. EPGs that are included in the preferred group can freely communicate with one another without contracts. EPGs that are excluded from the preferred group EPGs require contracts to communicate with one another and require contracts to communicate with any EPGs that are in the preferred group (the default behavior of ACI policy model). The following figure shows the contract requirement rules for preferred group.

Figure 268. Contract enforcement rules for preferred group

Open Communication within Preferred Group. EPG-VLAN10, EPG-VLAN20, EPG-VLAN30 and External EPG can all talk to each other in this example

Default white-list model behavior. In this example, EPG-APP01 and EPG-APP02 needs a contract to talk to each other, and they need contracts to talk to any EPGs in Preferred Group.
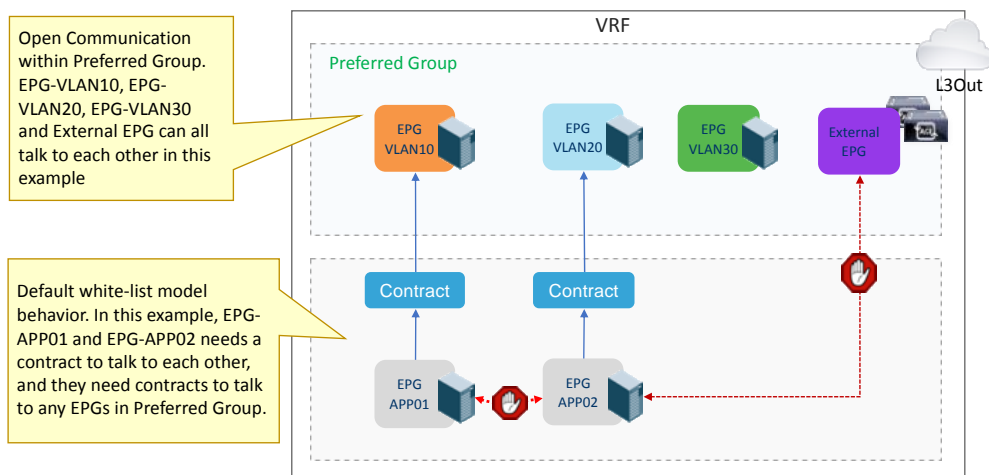
Configurating contract preferred group takes two steps:

1. Enable preferred group in the VRF's vzAny

    It is done at Tenant > Networking > VRF > EPG Collection for VRF. Under Policy > General in the right panel, you can see the option to either disable or enable "Preferred Group Member". The default setting is "Disabled". To enable it, you need to select "Enabled". The following figure shows preferred group enabled for VRF1.

Figure 269. Enabling preferred group for VRF1
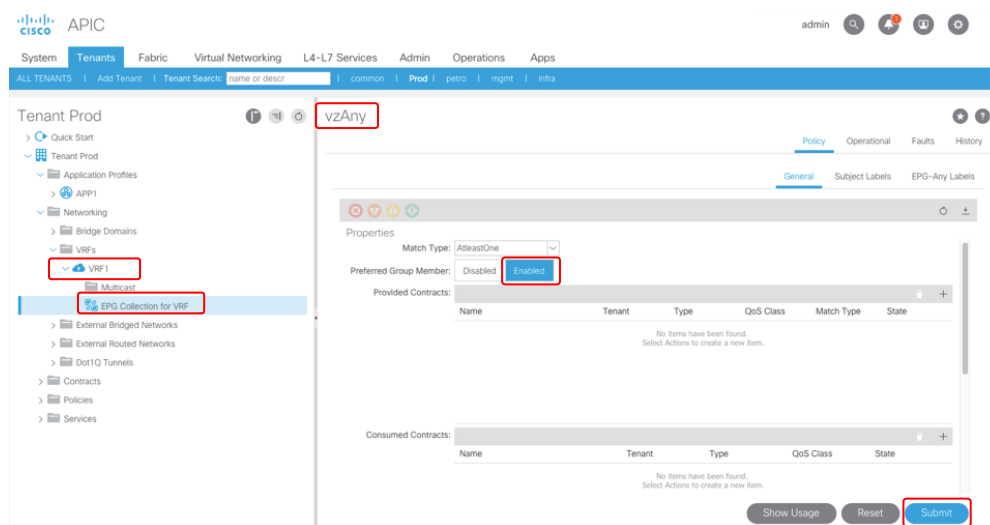
2. Include EPGs in preferred group

    It is done under Tenant > Application Profile > Application EPG. Under Policy > General in the right panel for the selected EPG, you see "Exclude" and "Include" as options for "Preferred Group Member". The default setting is "Exclude". If you want to include the EPG in the preferred group, you need to select "Include". We add EPG-Web and EPG-App into the preferred group, while leaving EPG-DB excluded from the preferred group.

Figure 270. Adding EPG-Web into the preferred group
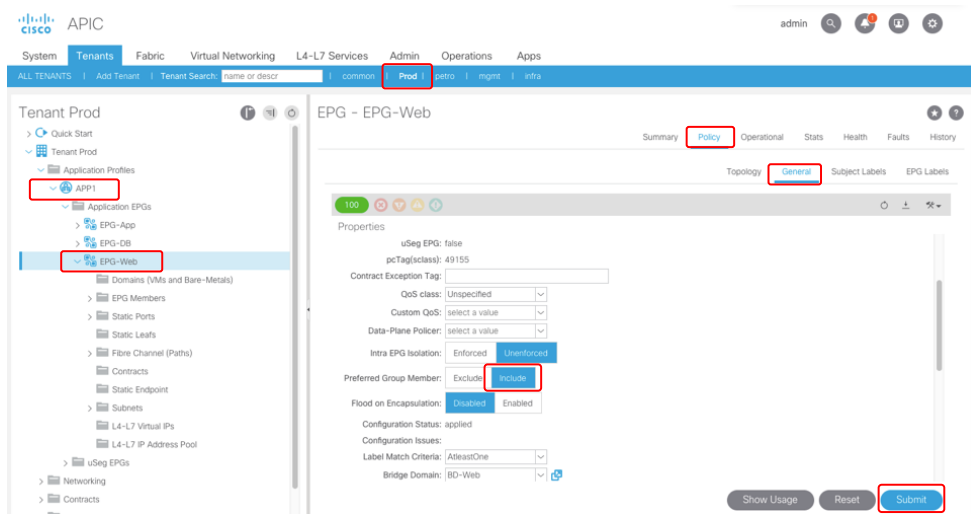
Figure 271. Adding EPG-App into the preferred group
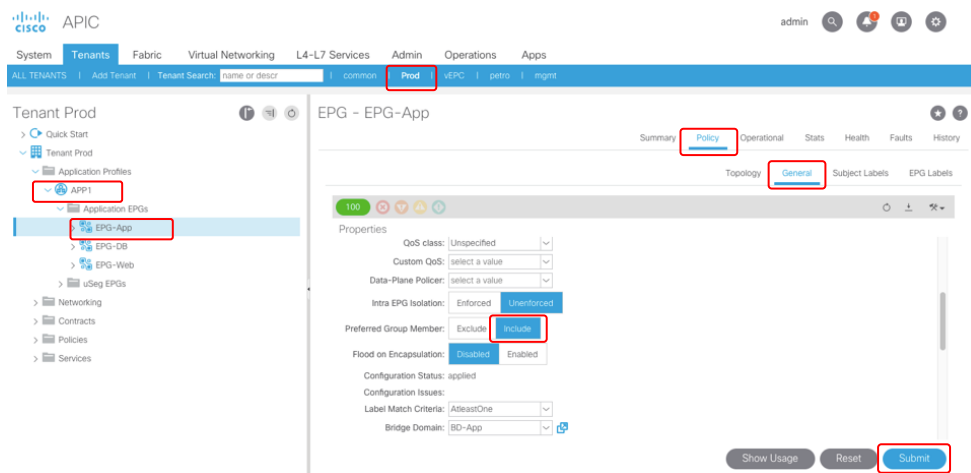
With this configuration, EPG-Web and EPG-App can talk with each other without any contract whereas EPG-DB is still under white-list-based contract control. In order to allow EPG-App to talk to EPG-DB, you still need to apply a contract between them. Figure 272 depicts the communication relationship among the three EPGs with the preferred group configuration in this example.

Figure 272. Communication relationship among the three EPGs in App1
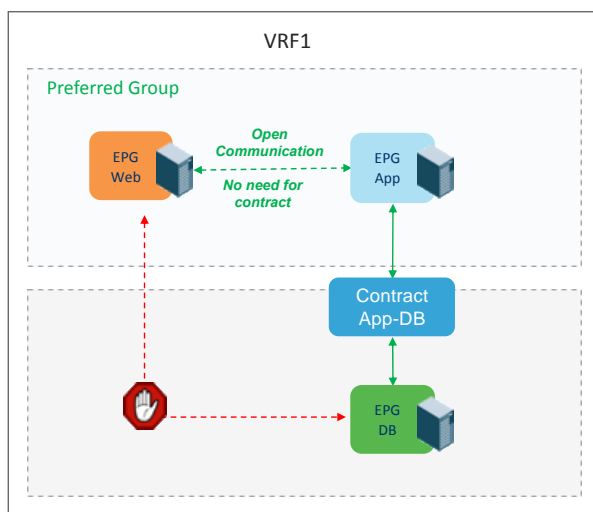
L3out

Cisco ACI refers to external Layer 3 connectivity as a L3out, which allows you to use standard Layer 3 technologies to connect to external network. These can be Layer 3 connections to an existing network, WAN routers, firewalls, mainframes, or any other Layer 3 device. This section describes design and step-by-step configuration example of L3out. This document focuses on two common design use cases:

■ North-South L3out Design: Contract between L3out EPG and regular EPG for North-South traffic. This to connect ACI fabric to external network

■ Transit L3out Design: Contract between L3out EPGs. This is to use ACI fabric as a transit network.

For each design, the following are discussed:

■ Design overview

■ Configuration steps

In this section, you are going to add L3outs and contracts for the L3outs to the tenant design in previous section. If you use unenforced, preferred group or vzAny contract, you might not need a contract.

North-South L3out Design

Assuming you already have EPGs and contracts for East-West communication within a tenant, you will likely need an external connectivity to get your application accessible from outside of ACI fabric. In that case, a contract between L3out EPG and a regular EPG is required. The following figure shows the example of North-South L3out design that is a contract "External-Web" between L3Out1 EPG as consumer and EPG-Web as provider in addition to East-West communication. If you want to get EPG-App and EPG-DB accessible as well, you need to have them provide the contract "External-Web".

Figure 273. Contract External-Web between L3Out1 and EPG-Web
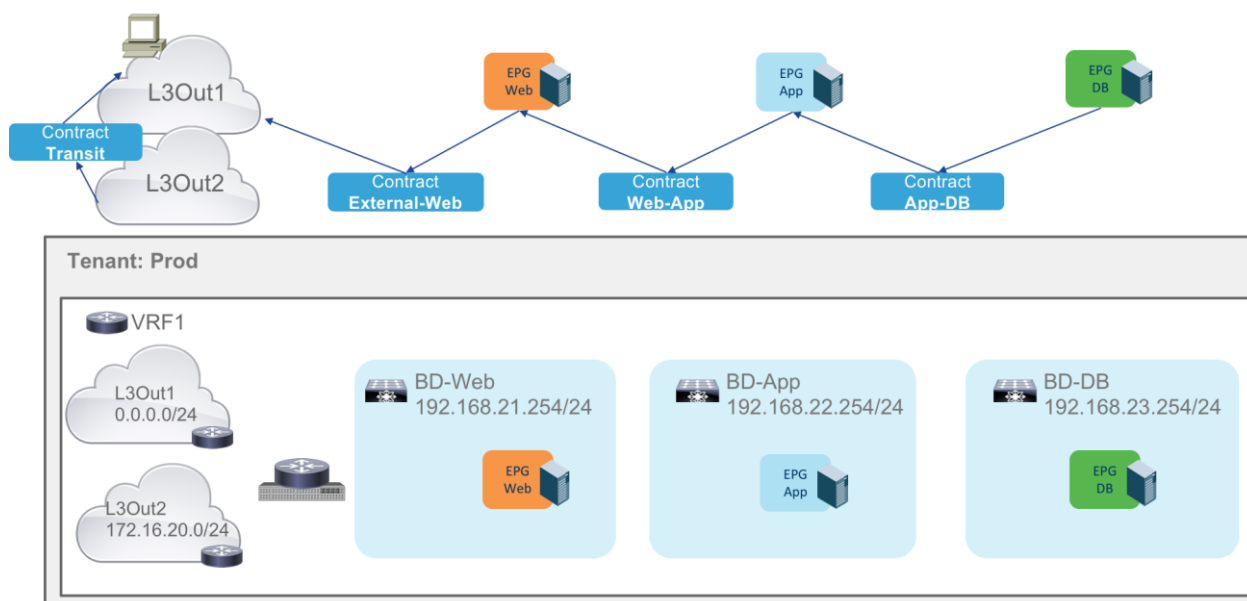
In addition to contract design, a connectivity to external network should be taken into consideration. The following choices are available for external connectivity:

- Use a pair of leaf nodes as both the computing and VRF-lite L3Out border leaf nodes (or border leaf for short). These leaf nodes are used to connect endpoints and to connect to WAN or campus routers.

- Use a dedicated pair of border leaf nodes. In this case, no servers connect to the leaf. Connections are only to WAN or campus routers.

- Use Layer 3 EVPN services through the spine (GOLF) instead of using a border leaf. (It's not covered in this document)

Then, the following interface options are available for a connectivity between border leaf and external device

- Routed interface: This is common when each router physical interface is dedicated to a VRF.

- Routed sub-interface: This is common when router physical interfaces are shared by multiple tenants or VRFs.

- SVI (Switch Virtual Interface): This is common to connect service devices, for example Firewall and Load-balancer, via port-channel or vPC and to share service device physical interfaces.

ACI supports Layer 3 connections using static routing (IPv4 and IPv6) or the following dynamic routing protocols:

- OSPFv2 (IPv4) and OSPFv3 (IPv6)

- BGP (IPv4 and IPv6)

- EIGRP (IPv4 and IPv6)

The following figure illustrates what is going to be covered in this document: Use of pairs of border leaf nodes and routed sub-interface connectivity to external router by using OSPF network type point-to-point. From border leaf nodes, ACI fabric advertises the subnet of EPG-Web (BD-Web: 192.168.21.0/24) to external. From external routers via border leaf nodes, ACI fabric learns 0.0.0.0/0 external network. 0.0.0.0/0 is used for L3Out1 external EPG classification.

- Border leaf nodes

- Routed sub-interface

- OSPFv2 (IPv4)

- 0.0.0.0/0 as L3Out1 EPG subnet

Figure 274. External connectivity
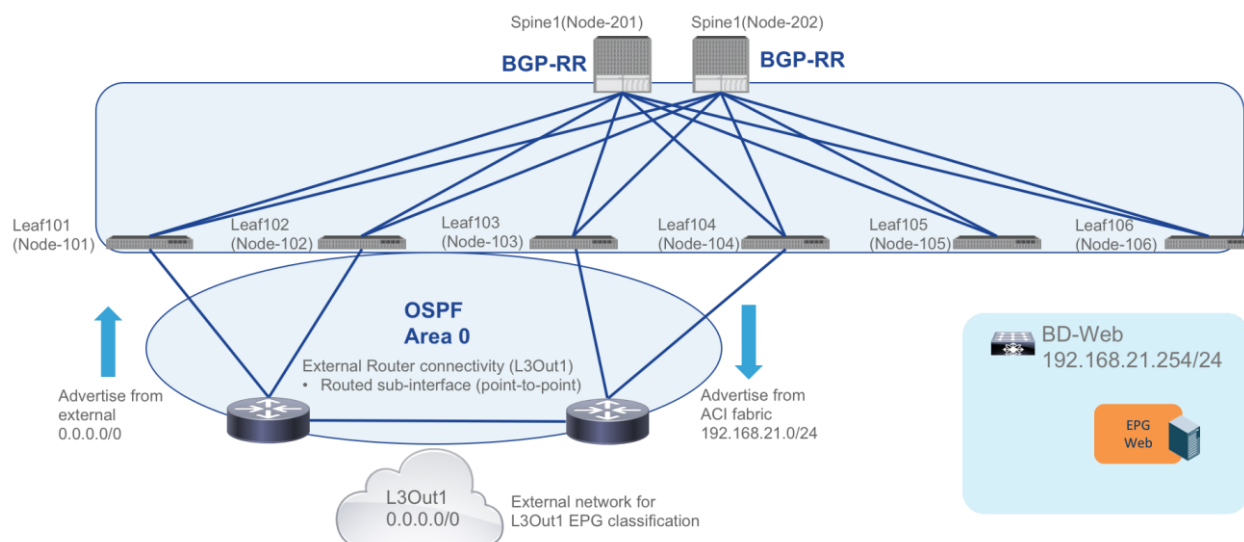
Note: L3out is used to configure interfaces, protocols and protocol parameters necessary to provide IP connectivity to external routing devices. Part of the L3Out configuration involves also defining an external network (also known as an external EPG) for the purpose of access-list filtering. The external network is used to define which subnets are potentially accessible through the L3out connection. In above figure, 0.0.0.0/0 is used for L3Out1 external EPG classification, which means the networks 0.0.0.0/0, all possible routes, is accessible through an L3Out connection. If you want to let only particular external network subnet access to an EPG in ACI fabric, more specific subnet configuration for L3out external EPG is required. For example, if you have 10.0.0.0/24 as L3Out1 external EPG classification, only 10.0.0.0/24 external network is accessible from EPG-Web through L3out connection.

North-South L3out Configuration

The North-South L3out configuration in ACI includes the following steps:

1. Create a L3Out

   - Create Routed Outside

   - Create Node Profile

   - Create Interface Profile

   - Create External EPG Network

2. Configure a BD to advertise BD subnet

3. Configure a contract between the L3Out and an EPG

The rest of this section detailes the steps for the L3Out1 configuration.

The following table and figure show the IP addressing used in this document.

Cisco Systems. All printed copies and duplicate soft copies are considered uncontrolled and the original online version should be referred to for the latest version.

Page 168 of 215

Table 5. IP Addressing

Device, Interface	IP address/subnet mask
Leaf101 router-id	10.1.1.1
Leaf101 Eth1/33.1201	172.16.21.1/30
Leaf102 router-id	10.1.1.2
Leaf102 Eth1/33.1201	172.16.21.5/30
Leaf103 router-id	10.1.1.3
Leaf103 Eth1/5.1201	172.16.21.9/30
Leaf104 router-id	10.1.1.4
Leaf104 Eth1/1.1201	172.16.21.13/30

Figure 275. IP addressing
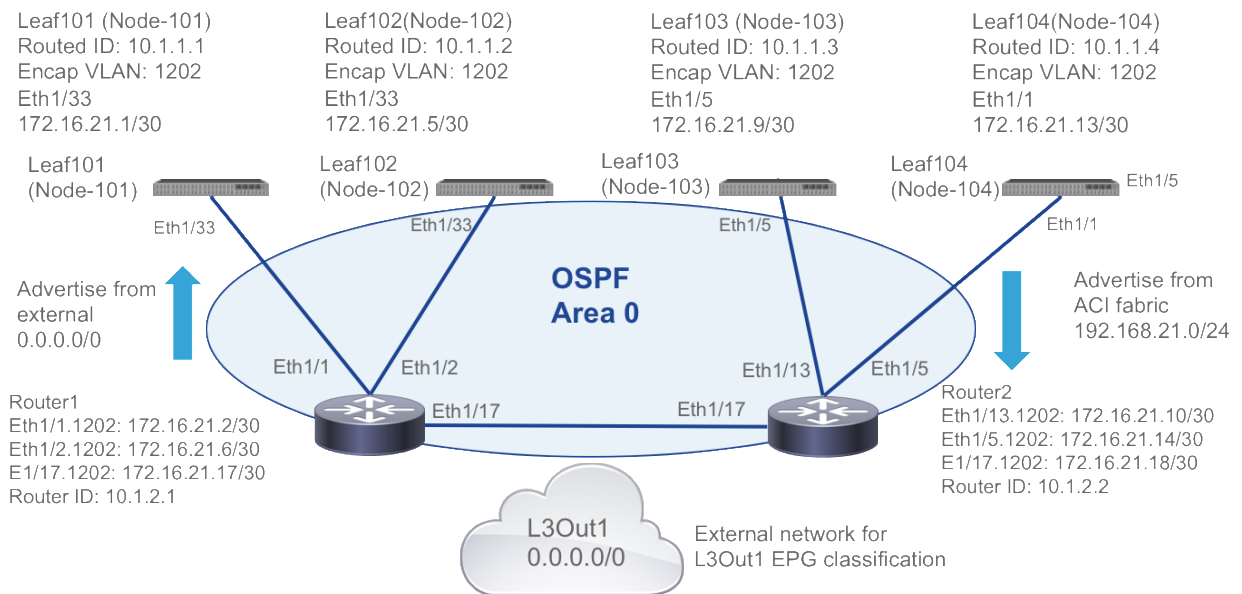
Create a L3Out

You are going to create L3Out1 with routed sub-interface and OSPFv2.

1. Navigate to Tenant > Networking > External Routed Networks > Create Routed Outside

Figure 276. Create Routed outside for L3Out1
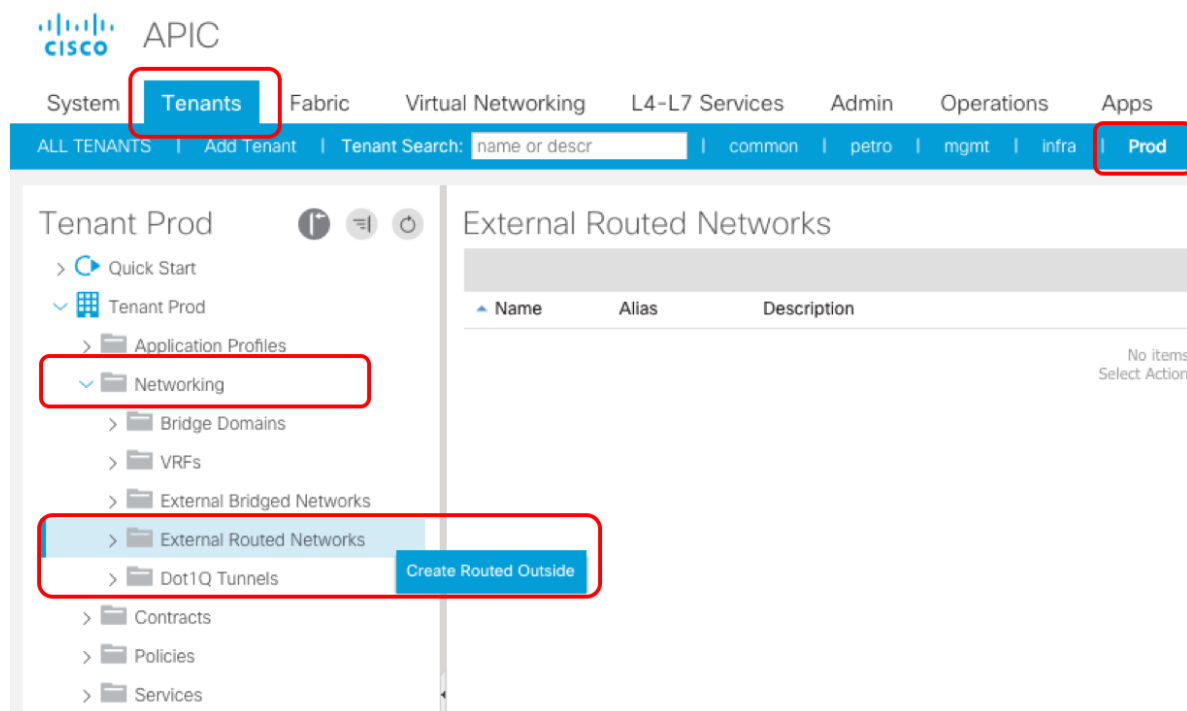
2. In the Create Routed Outside wizard, select VRF and Routing Protocol.

   ■ Name: L3Out1

   ■ VRF: VRF1

   ■ External Routed Domain: Not necessary (If it's routed interface or routed sub-interface, External Routed Domain is not required)

   ■ Check OSPF

      ■ OSPF Area id: 0

      ■ OSPF Area Type: Regular area

   ■ Add Node and Interfaces Protocol Profiles by clicking + icon on the bottom.

Figure 277. Select VRF and Routing Protocol

3. Create Node Profile pop-up appears. Then, add Nodes.

   ■ Name: L3Out1-NP

   ■ Add Nodes by clicking + icon

Figure 278. Create Node Profile

4. Select Node pop-up appears. Then, add Node Leaf101 and repeat same steps for Leaf102, 103 and 104.

- ■ Name: Leaf101(Node-101)

  - ● Router ID: 10.1.1.1

- ■ Name: Leaf102(Node-102)

  - ● Router ID: 10.1.1.2

- ■ Name: Leaf103(Node-103)

  - ● Router ID: 10.1.1.3

- ■ Name: Leaf104(Node-104)

  - ● Router ID: 10.1.1.4

Figure 279. Add Node

Cisco Systems. All printed copies and duplicate soft copies are considered uncontrolled
and the original online version should be referred to for the latest version.

Page 172 of 215

5. Verify Nodes are added and add OSPF Interfaces Profiles by clicking + icon.

Figure 280. Add OSPF Interface Profiles

6. Create Interface Profile wizard appears.

   ■ Name: L3Out1-IP

   ■ Check Config Protocol Profiles

Figure 281. Create Interface Profile

7. Then, create OSPF Interface policy. Select Create OSPF Interface Policy from the OSPF Policy pull-down menu.

Figure 282. Create OSPF Interface Policy

8. Create OSPF Interface Policy pop-up appears. Select Point-to-point and submit.

   ■ Name: Point-to-point

   ■ Network Type: Point-to-point

Figure 283. Create OSPF Interface Policy

Cisco Systems. All printed copies and duplicate soft copies are considered uncontrolled and the original online version should be referred to for the latest version.

Page 174 of 215

9. Verify Point-to-point is selected and click Next.

Figure 284. Select OSPF Policy

10. Specify the interfaces for L3Out1.

   ■   Select Routed Sub-interface

   ■   Add Routed Sub-Interfaces by clicking + icon.

Figure 285. Specify the Interfaces

11. Select Routed Sub-interface pop-up appears. Select Leaf101 interface and repeat same steps for Leaf102, 103 and 104 interfaces.

- ■ Path Type: Port

- ■ Node: Leaf101

  - ● Path: eth1/33

    - o Encap: 1202

    - o IPv4 Primary Address: 172.16.21.1/30

- ■ Node: Leaf102

  - ● Path: eth1/33

    - o Encap: 1202

    - o IPv4 Primary Address: 172.16.21.5/30

- ■ Node: Leaf103

  - ● Path: eth1/5

    - o Encap: 1202

    - o IPv4 Primary Address: 172.16.21.9/30

- ■ Node: Leaf104

  - ● Path: eth1/1

    - o Encap: 1202

    - o IPv4 Primary Address: 172.16.21.13/30

Note: The default MTU configuration is inherit that means 9000 bytes. You may need to change it based on MTU of your external router interface.

Figure 286. Select Router Sub-Interfaces

Cisco Systems. All printed copies and duplicate soft copies are considered uncontrolled and the original online version should be referred to for the latest version.

Page 176 of 215

12. Verify Routed Sub-interfaces are added and click OK.

Figure 287. Add Router Sub-Interfaces
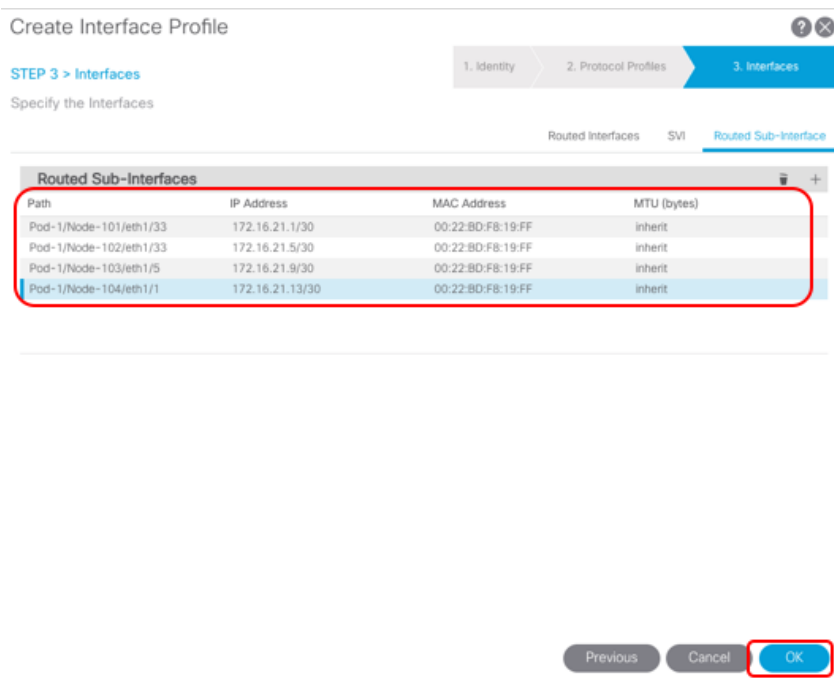
13. Verify OSPF Interface Profile is added and click OK.

Figure 288. Create Node Profile

Cisco Systems. All printed copies and duplicate soft copies are considered uncontrolled
and the original online version should be referred to for the latest version.

Page 177 of 215
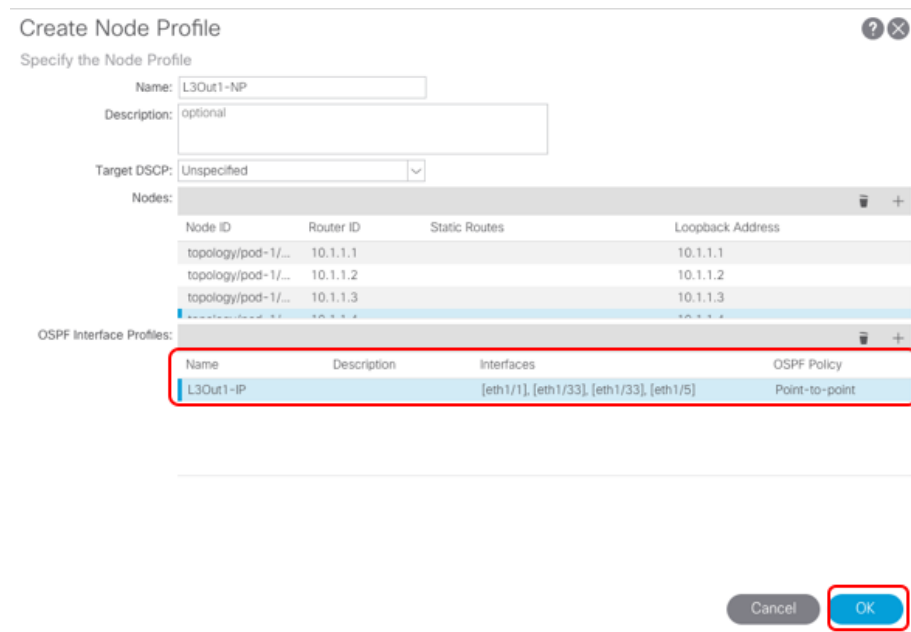
14. Verify Node Profile is added and click Next.

Figure 289. Verify Node Profile is added

15. Create External EPG Network by clicking + icon.

Figure 290. Create External EPG Network

16. Create Subnet by clicking + icon.

■ Name: L3Out1

Figure 291. Create External EPG Network

17. Create Subnet pop-up appears. Specify 0.0.0.0/0 that means all external route in this VRF is classified to the L3Out1 EPG.

■   IP Address: 0.0.0.0/0

■   Scope: External Subnets for the External EPG

Figure 292. Create External EPG Network

18. Verify External EPG Network is added and click Finish.

Figure 293. Finish to create L3Out1 and External EPG Network

19. Once you click Finish, you can see External EPG L3Out1 under the L3Out1.

Figure 294. Verify L3Out1 is created
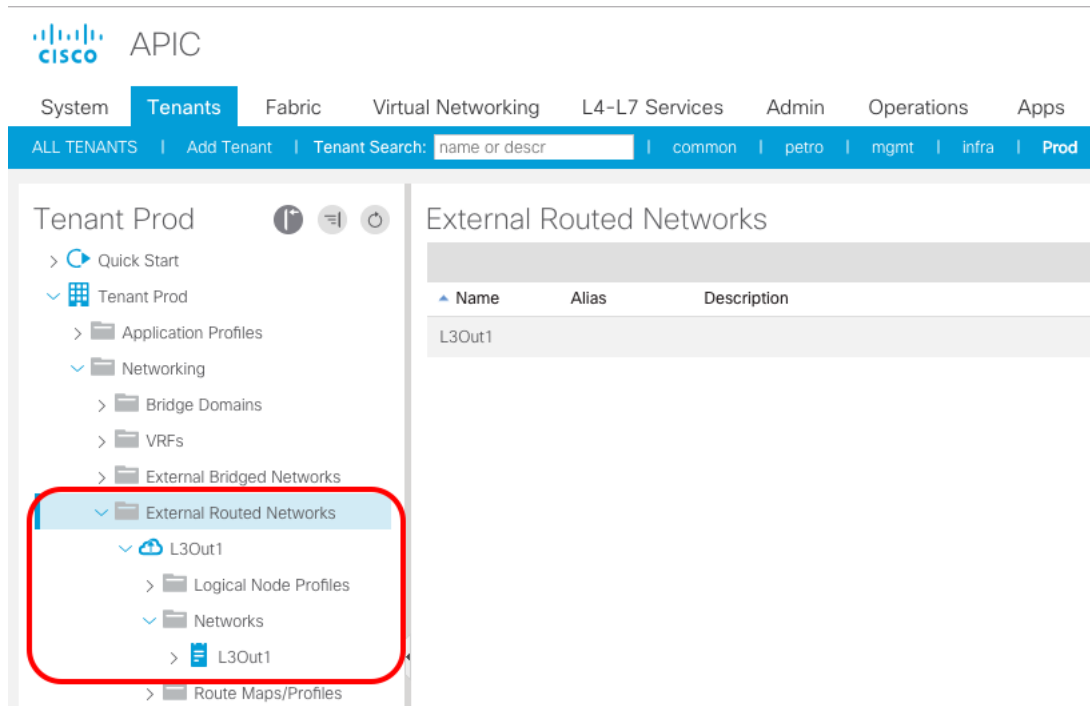
Now Border leaf nodes have OSPF interfaces. If external routers are also configured accordingly, border leaf nodes have OSPF neighbor established and external routes learned from external routers.

```
leaf101# show ip ospf interface vrf Prod:VRF1

 loopback3 is up, line protocol is up

    IP address 10.1.1.1/32, Process ID default VRF Prod:VRF1, area
backbone

    Enabled by interface configuration

    State LOOPBACK, Network type LOOPBACK, cost 1


 Ethernet1/33.23 is up, line protocol is up

    IP address 172.16.21.1/30, Process ID default VRF Prod:VRF1, area
backbone

    Enabled by interface configuration

    State P2P, Network type P2P, cost 4

    Index 93, Transmit delay 1 sec

    1 Neighbors, flooding to 1, adjacent with 1

    Timer intervals: Hello 10, Dead 40, Wait 40, Retransmit 5
```

```
      Hello timer due in 00:00:10

    No authentication

    Number of opaque link LSAs: 0, checksum sum 0


leaf101# show ip ospf neighbors vrf Prod:VRF1
 OSPF Process ID default VRF Prod:VRF1
 Total number of neighbors: 1
 Neighbor ID      Pri State            Up Time  Address         Interface
 10.1.2.1           1 FULL/ -          00:02:06 172.16.21.2     Eth1/33.23


leaf101# show ip route ospf vrf Prod:VRF1
IP Route Table for VRF "Prod:VRF1"
'*' denotes best ucast next-hop
'**' denotes best mcast next-hop
'[x/y]' denotes [preference/metric]
'%<string>' in via output denotes VRF <string>


0.0.0.0/0, ubest/mbest: 1/0
    *via 172.16.21.2, eth1/33.23, [110/1], 00:02:51, ospf-default, type-2
172.16.21.4/30, ubest/mbest: 1/0
    *via 172.16.21.2, eth1/33.23, [110/8], 00:02:51, ospf-default, intra
172.16.21.8/30, ubest/mbest: 1/0
    *via 172.16.21.2, eth1/33.23, [110/12], 00:02:51, ospf-default, intra
172.16.21.12/30, ubest/mbest: 1/0
    *via 172.16.21.2, eth1/33.23, [110/12], 00:02:51, ospf-default, intra
172.16.21.16/30, ubest/mbest: 1/0
    *via 172.16.21.2, eth1/33.23, [110/8], 00:02:51, ospf-default, intra
```

However, external routers haven't seen routes from ACI fabric yet. To advertise route from ACI fabric to external, contract and BD subnet scope need to be configured accordingly.

```
Router1# show ip route ospf vrf Prod:VRF1

IP Route Table for VRF "Prod:VRF1"

'*' denotes best ucast next-hop

'**' denotes best mcast next-hop

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

'%<string>' in via output denotes VRF <string>


10.1.1.1/32, ubest/mbest: 1/0

    *via 172.16.21.1, Eth1/1.1202, [110/5], 00:04:09, ospf-1, intra

10.1.1.2/32, ubest/mbest: 1/0

    *via 172.16.21.5, Eth1/2.1202, [110/5], 00:04:09, ospf-1, intra

10.1.1.3/32, ubest/mbest: 1/0

    *via 172.16.21.18, Eth1/17.1202, [110/9], 00:04:09, ospf-1, intra

10.1.1.4/32, ubest/mbest: 1/0

    *via 172.16.21.18, Eth1/17.1202, [110/9], 00:04:09, ospf-1, intra

172.16.21.8/30, ubest/mbest: 1/0

    *via 172.16.21.18, Eth1/17.1202, [110/8], 02:43:50, ospf-1, intra

172.16.21.12/30, ubest/mbest: 1/0

    *via 172.16.21.18, Eth1/17.1202, [110/8], 02:43:50, ospf-1, intra
```

Configure a BD to advertise BD subnet

By default, a BD subnet scope is "Private to VRF". To advertise the BD subnet to external, scope must be "Advertised Externally". In this example, as EPG-Web is going to have a contract with L3Out1 EPG, you need to set BD1 subnet scope to advertise BD1 subnet that is EPG-Web subnet.

Cisco Systems. All printed copies and duplicate soft copies are considered uncontrolled
and the original online version should be referred to for the latest version.

Page 183 of 215

20. The location is at Tenant > Networking > Bridge Domains > Bridge Domain > Subnets > Subnet

Figure 295. BD subnet scope setting: Advertised Externally
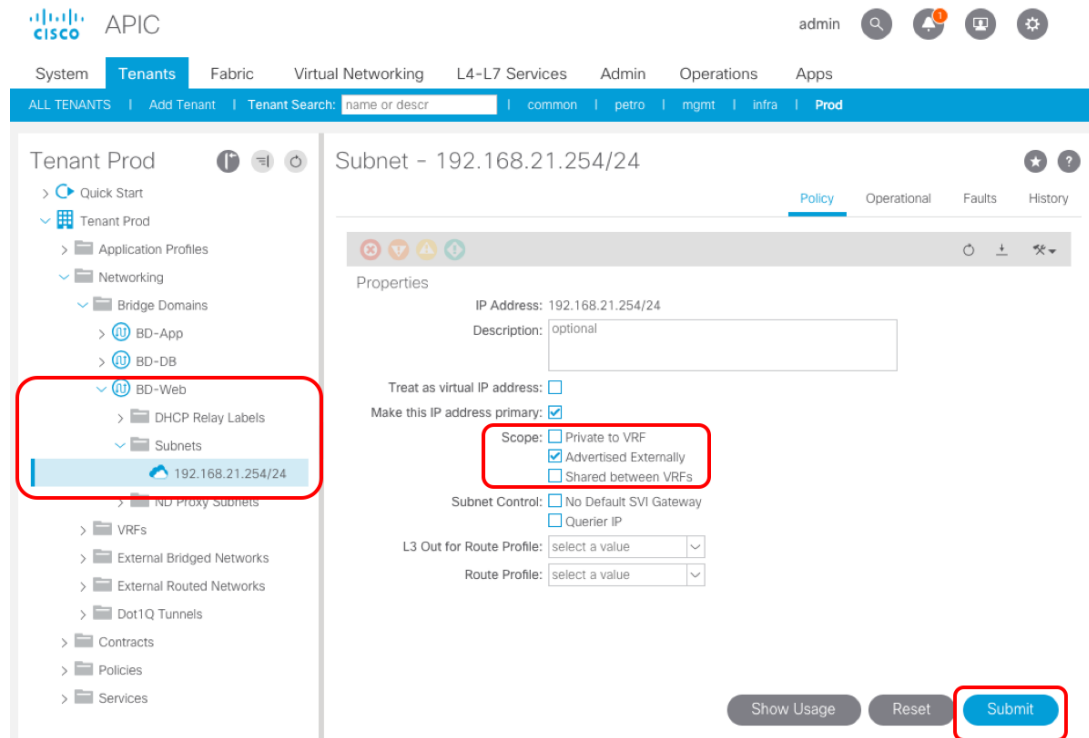
In addition to this, Associated L3 Outs need to be specified. So that ACI fabric administrator can manage which external network the subnet is advertised.

Cisco Systems. All printed copies and duplicate soft copies are considered uncontrolled
and the original online version should be referred to for the latest version.

Page 184 of 215

21. From Tenant > Networking > Bridge Domains > Policy > L3 Configurations, add Associated L3 Out by clicking + icon and select L3Out1.

Figure 296. BD subnet scope setting: Advertised Externally
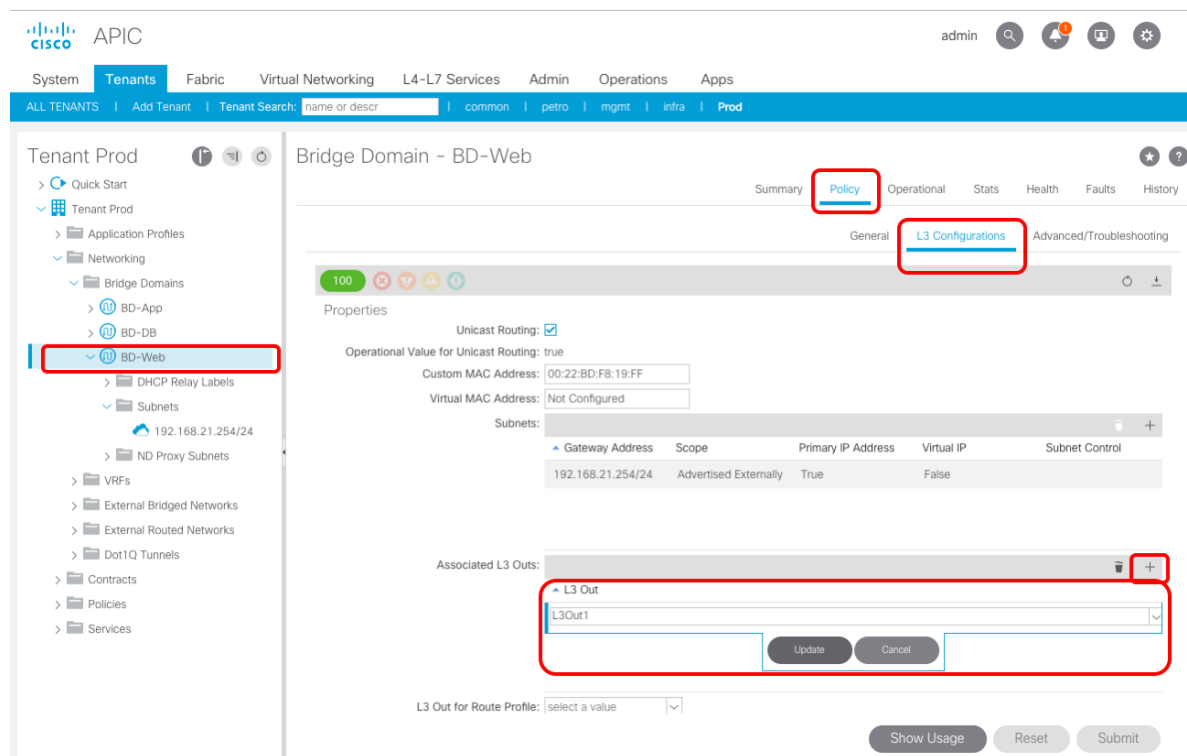
Then, border leaf starts advertising BD-Web subnet 192.168.21.0/24 to external router of L3Out1. Now you can see 192.168.21.0/24 on external routers.

```
Router1# show ip route ospf vrf Prod:VRF1

IP Route Table for VRF "Prod:VRF1"

'*' denotes best ucast next-hop

'**' denotes best mcast next-hop

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

'%<string>' in via output denotes VRF <string>


10.1.1.1/32, ubest/mbest: 1/0

    *via 172.16.21.1, Eth1/1.1202, [110/5], 00:09:56, ospf-1, intra

10.1.1.2/32, ubest/mbest: 1/0

    *via 172.16.21.5, Eth1/2.1202, [110/5], 00:09:56, ospf-1, intra

10.1.1.3/32, ubest/mbest: 1/0

    *via 172.16.21.18, Eth1/17.1202, [110/9], 00:09:56, ospf-1, intra
```

```
10.1.1.4/32, ubest/mbest: 1/0

    *via 172.16.21.18, Eth1/17.1202, [110/9], 00:09:56, ospf-1, intra

172.16.21.8/30, ubest/mbest: 1/0

    *via 172.16.21.18, Eth1/17.1202, [110/8], 02:49:37, ospf-1, intra

172.16.21.12/30, ubest/mbest: 1/0

    *via 172.16.21.18, Eth1/17.1202, [110/8], 02:49:37, ospf-1, intra

192.168.21.0/24, ubest/mbest: 2/0

    *via 172.16.21.1, Eth1/1.1202, [110/20], 00:01:39, ospf-1, type-2

    *via 172.16.21.5, Eth1/2.1202, [110/20], 00:01:39, ospf-1, type-2
```

Configure a contract between the L3out and an EPG

Finally, you are going to create a contract External-Web for between External EPG L3Out1 and EPG-Web. Otherwise external traffic can't reach to endpoints in EPG-Web as the traffic is not permitted by ACI contract even though routes are there. If you use network centric tenant that has vzAny contract, you may skip this step.

1. Navigate to Tenant > Contracts. Then, Create Contract pop-up appears. In this example, we are going to reuse a filter that was created in previous section.

   ■ Name: External-Web

   ■ Add subject by clicking + icon. Then, Create Contract pop-up appears.

   ■ Subject name: sbjct-permit-all

      ● Check Apply Both Directions and Reverse Filter Ports (default)

      ● Add a filter by clicking + icon

      ● Filter: ftr-permit-all

Cisco Systems. All printed copies and duplicate soft copies are considered uncontrolled
and the original online version should be referred to for the latest version.

Page 186 of 215

Figure 297. Create contract

Figure 298. Create contract subject

2. Verify subject is added and click Submit.

Figure 299. Create contract

3. Then, configure contract for EPG-Web as provider and for External EPG L3Out1 as consumer.

For EPG-Web, the location is at Tenant > Application Profiles > APP1 > Application EPGs > EPG-Web > Contracts. Then, Add Provided Contract pop-up appears.

Figure 300. Configure contract for EPG-Web

4. For External EPG L3Out1, the location is at Tenant > Networking> External Routed Networks > L3Out1 > Networks > L3Out1 > Policy > Contracts > Consumed Contracts.

Figure 301. Configure contract for External EPG L3Out1
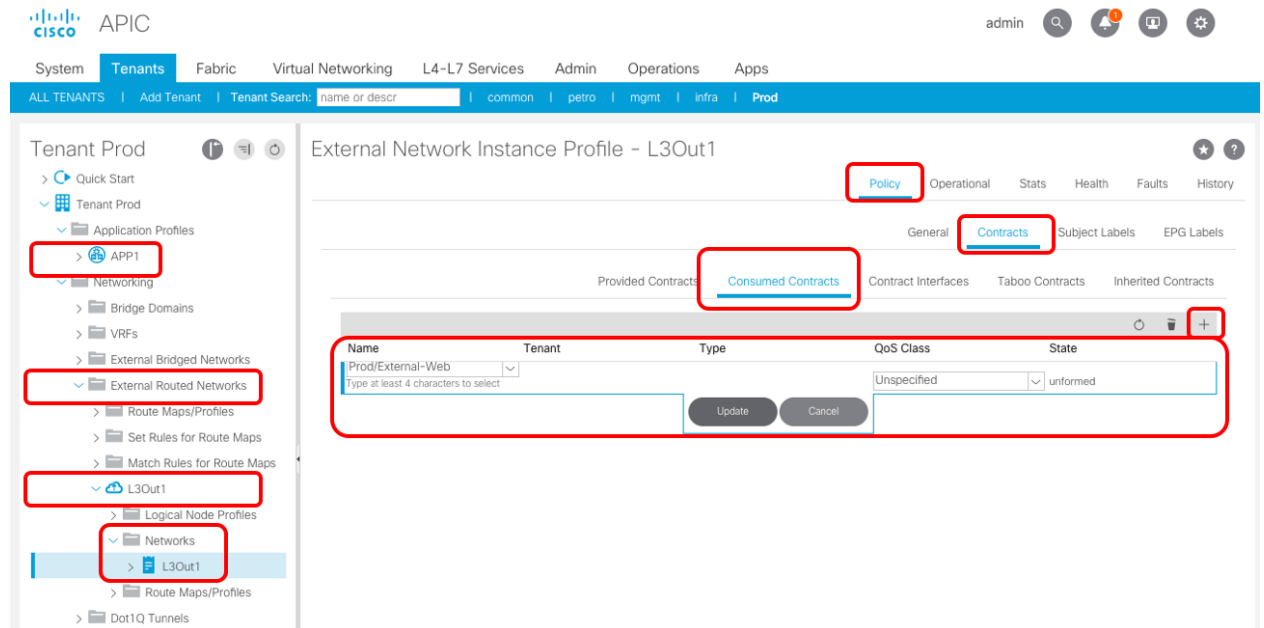
5. Once a contract between EPGs are configured, endpoint in EPG-Web can communicate with IP in L3Out1. For example, external client 172.16.10.1 can ping to 192.168.21.11 in EPG-Web.

Figure 302. External connectivity
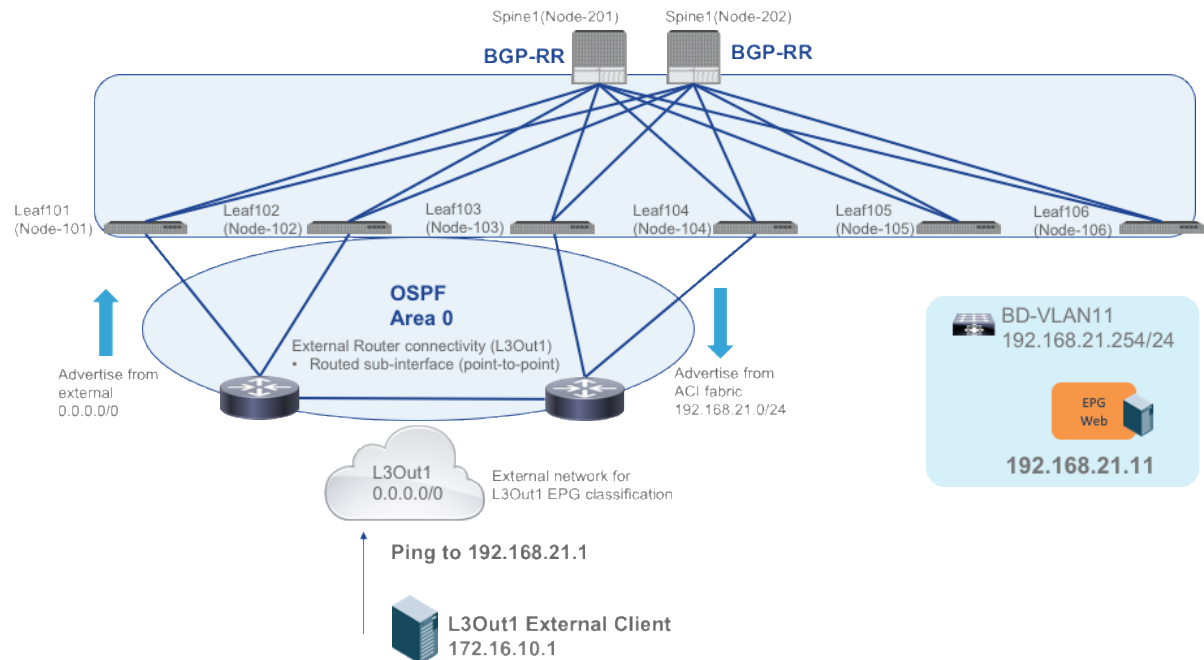
```
[root@rescue /]# ifconfig eth0
eth0      Link encap:Ethernet  HWaddr 00:50:56:B6:07:F1
          inet addr:172.16.10.1  Bcast:172.16.10.255  Mask:255.255.255.0
          inet6 addr: fe80::250:56ff:feb6:7f1/64 Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:306 errors:0 dropped:0 overruns:0 frame:0
          TX packets:5372 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:18854 (18.4 Kb)  TX bytes:1880766 (1.7 Mb)

[root@rescue /]# ping 192.168.21.11
PING 192.168.21.11 (192.168.21.11) 56(84) bytes of data.
64 bytes from 192.168.21.11: icmp_seq=1 ttl=61 time=0.431 ms
64 bytes from 192.168.21.11: icmp_seq=2 ttl=61 time=0.354 ms
64 bytes from 192.168.21.11: icmp_seq=3 ttl=61 time=0.382 ms
64 bytes from 192.168.21.11: icmp_seq=4 ttl=61 time=0.417 ms
64 bytes from 192.168.21.11: icmp_seq=5 ttl=61 time=0.334 ms
64 bytes from 192.168.21.11: icmp_seq=6 ttl=61 time=0.377 ms
```

If you remove the contract "External-Web" from either EPG-Web or L3Out1 EPG, ping stops working. If you add it back, ping starts working again.

Transit L3out Design

When multiple L3Outs are there, external routes learned from one L3Out can be advertised through another L3Out, making the Cisco ACI fabric as a transit network. Below is an example of transit L3out design that is a contract "Transit" between L3Out1 EPG as consumer and L3Out2 EPG as provider.

Figure 304. Transit L3out
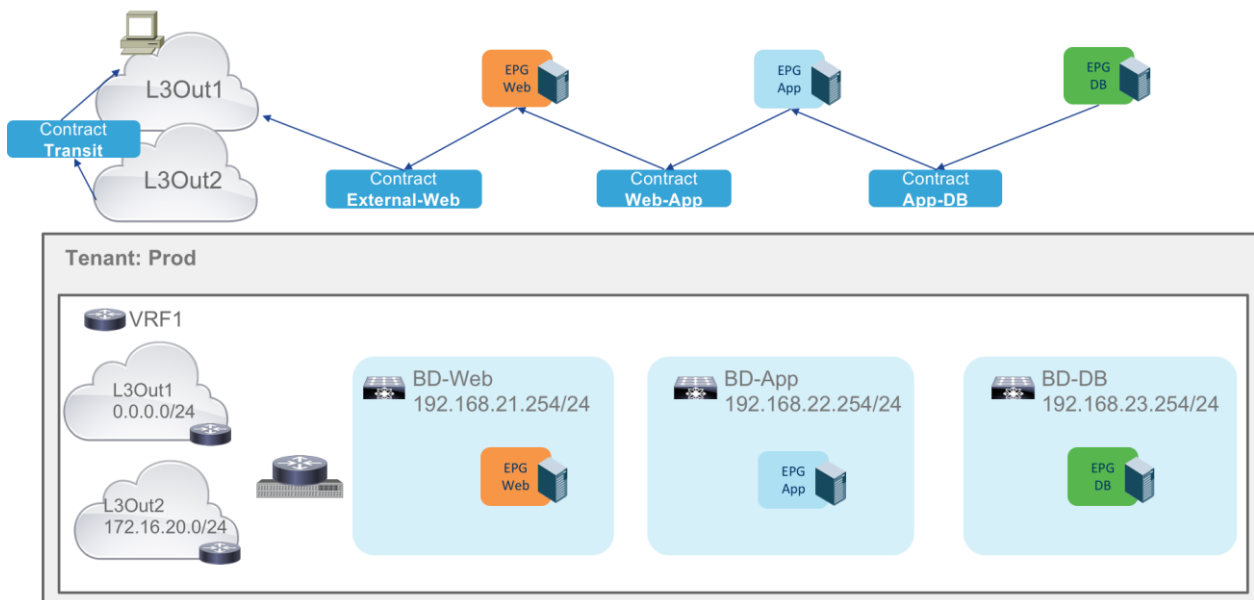
By default, Cisco ACI will not advertise routes learned from one L3out to another L3out. The Cisco ACI does not allow transit by default. Transit routing is controlled by creating export route control policies for a L3Out. It means you need to specify which route learned from a L3out should be exported to which L3out.

The following figure illustrates the example covered in this document where we use a dedicated pair of border leaf nodes and SVI on vPC with static route for external router for L3Out2. From L3Out1 border leaf nodes, ACI fabric advertises the subnet 172.16.20.0/24 from L3Out2 to external.

- Dedicated pair of border leaf nodes for L3Out2

- SVI over vPC

- Static route for 172.16.20.0/24 via L3Out2

- 172.16.20.0/24 as L3Out2 EPG subnet

- Export 172.16.20.0/24 to L3Out1 (OSPFv2)

Figure 305. External connectivity
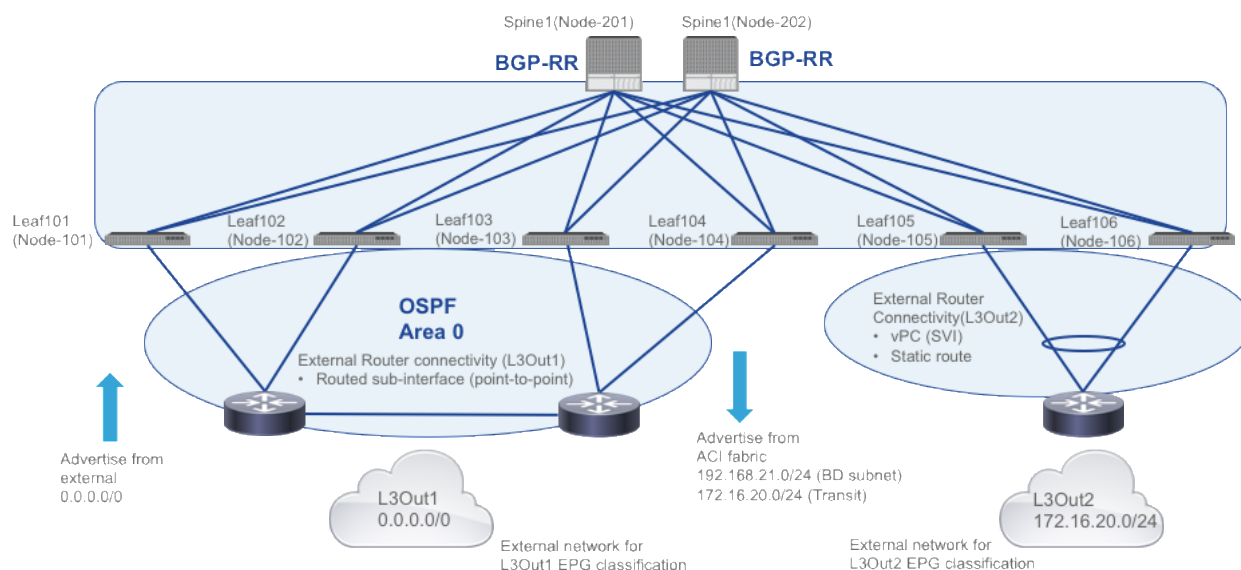
Note: Not all transit routing combinations are currently supported in ACI. For information about the currently supported transit routing combinations, see the Cisco APIC and Transit Routing document at the following URL: http://www.cisco.com/c/en/us/support/cloud-systems-management/application-policy-infrastructure-controller-apic/tsd-products-support-series-home.html

Transit L3out Configuration step

The Transit L3out configuration in ACI includes the following steps:

1. Create a second L3out

   - Create Routed Outside

   - Create Node Profile

   - Create Interface Profile

   - Create External EPG Network

2. Configure a contract between the L3outs

   The rest part of this session detailed the steps for a configuration of a contract between L3Out1 and L3Out2.

   The following table and figure show the IP addressing that we are going to use in configuration steps.

Table 6. IP addressing

Device, Interface	IP address/subnet mask
Leaf5 router-id	10.1.1.5
Leaf5 Primary (Interface vlan 1222)	172.16.22.252/24
Leaf5 Secondary (Interface vlan 1222)	172.16.22.254/24
Leaf6 router-id	10.1.1.6
Leaf6 Primary (Interface vlan 1222)	172.16.22.253/24
Leaf6 Secondary (Interface vlan 1222)	172.16.22.254/24
External router IP (Interface vlan 1222)	172.16.22.1 (next hop for static route)

Figure 306. IP addressing
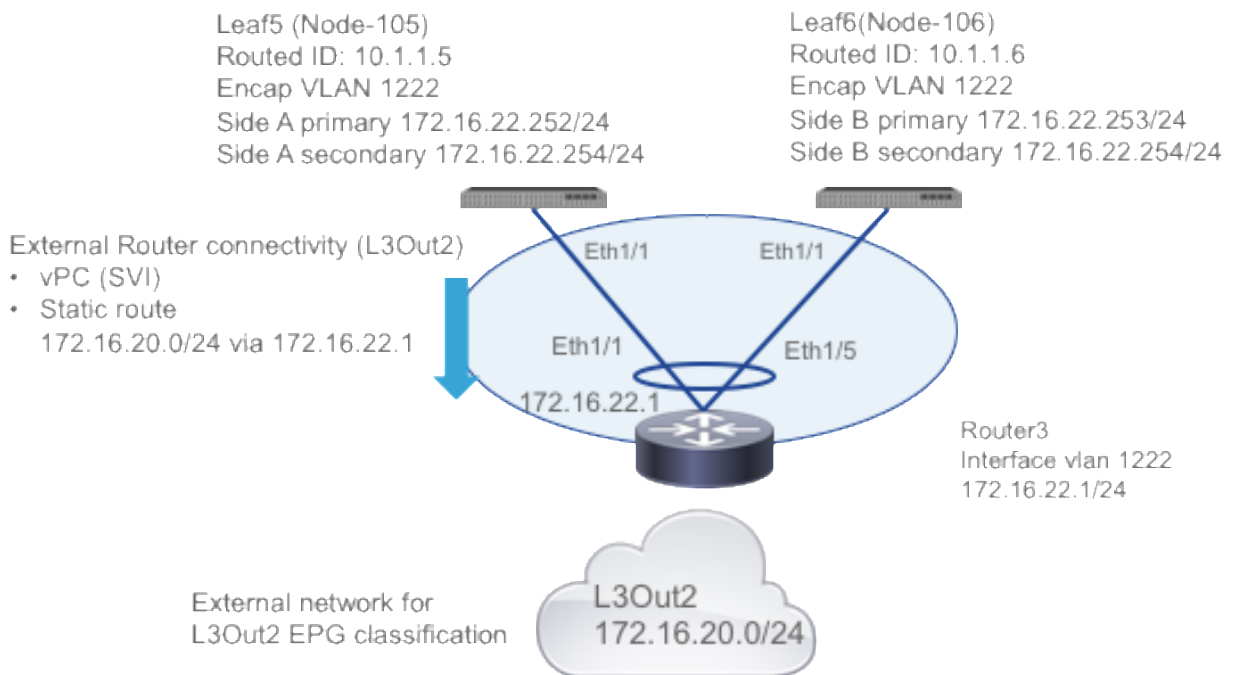
Create a second L3out

Similar to L3Out1 configuration, you are going to create L3Out2 with SVI on vPC interface and static route configuration.

1. Navigate to Tenant > Networking > External Routed Networks > Create Routed Outside

Figure 307. Create Routed outside for L3Out2
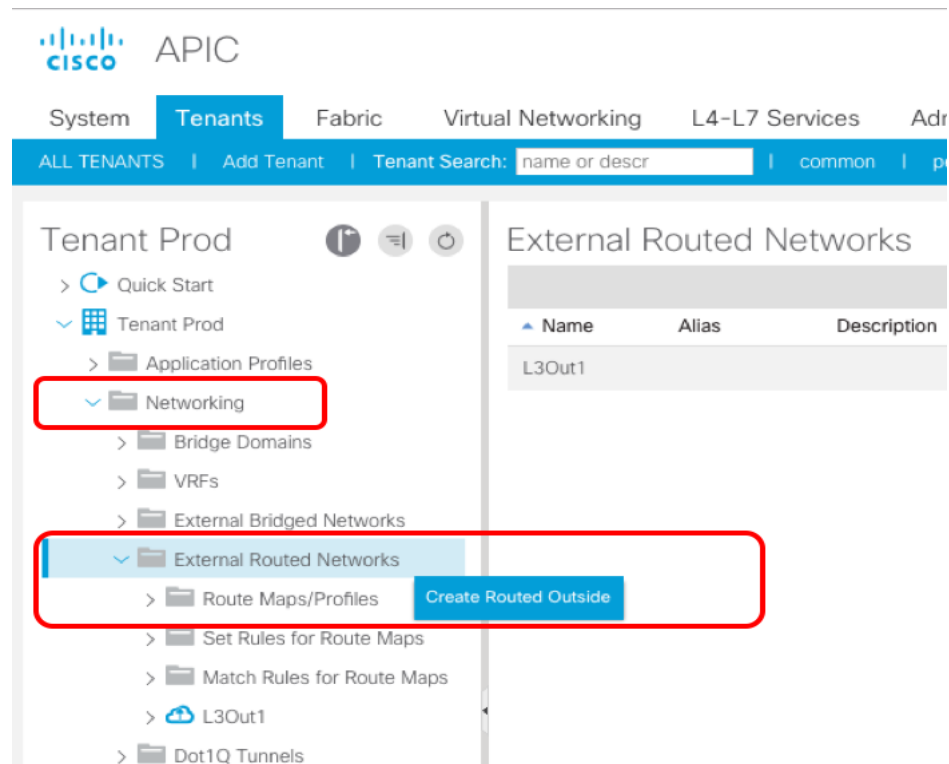
2. Create Routed Outside wizard appears. Select VRF and Routing Protocol. In this example, External Routed Domain needs to be specified as we are going to use SVI on vPC interface.

   ■ Name: L3Out2

   ■ VRF: VRF1

   ■ External Routed Domain: RoutedDomain2

   ■ Add Node and Interfaces Protocol Profiles by clicking + icon on the bottom.

Figure 308. Select VRF and External Routed Domain

3. Create Node Profile pop-up appears. Then, add Nodes.

 ■ Name: L3Out2-NP

 ■ Add Nodes by clicking + icon

Figure 309. Create Node Profile

4. Select Node pop-up appears. Add Node Leaf105 and add its static route. Then, repeat same steps for Leaf106.

 ■ Name: Leaf105(Node-105)

 • Router ID: 10.1.1.5

- Prefix: 172.16.20.0/24

- Next Hop IP: 172.16.22.1

- Preference: 1

■ Name: Leaf6(Node-106)

- Router ID: 10.1.1.6

- Prefix: 172.16.20.0/24

- Next Hop IP: 172.16.22.1

- Preference: 1

Figure 310. Add Node

Cisco Systems. All printed copies and duplicate soft copies are considered uncontrolled
and the original online version should be referred to for the latest version.

Page 195 of 215

5. Create Static Route pop-up appears.

Figure 311. Create static route

Figure 312. Add Node

6. Verify Nodes are added and add Interfaces Profiles by clicking + icon.

Figure 313. Create Interface Profile

7. Create Interface Profile wizard appears. In this example, you are going to skip protocol profile setting by unchecking Config Protocol Profiles.

   ■ Name: L3Out2-IP

   ■ Uncheck Config Protocol Profiles

Figure 314. Create Interface Profile

8. Specify the SVI interfaces for L3Out2.

   ■ Select SVI

■ Add SVI Interfaces by clicking + icon.

Figure 315. Specify the Interfaces
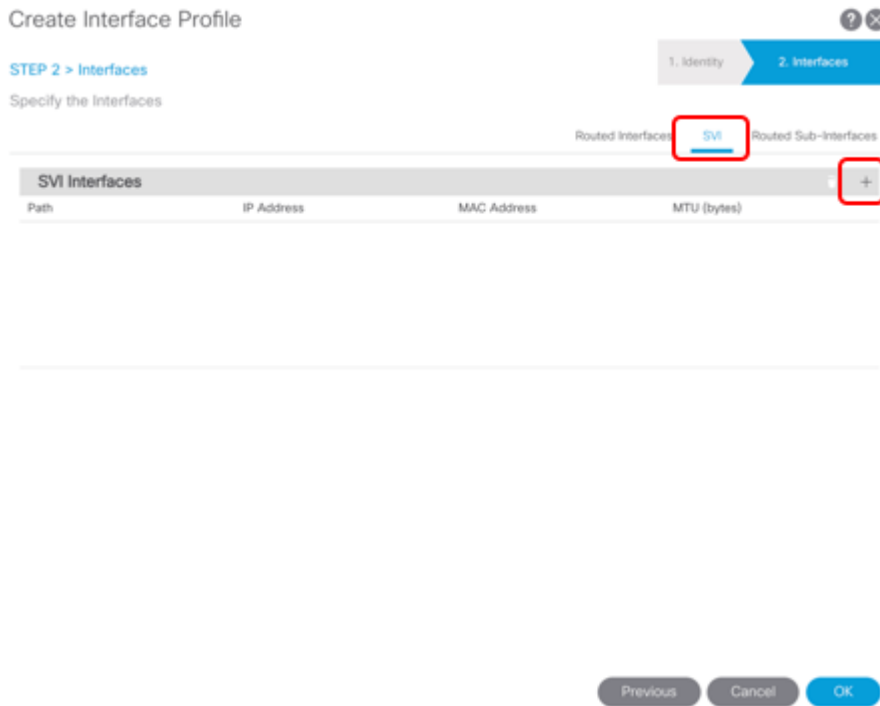
9. Select SVI pop-up appears. You are going to configure vPC interface between Leaf5 and Leaf6.

■ Path Type: Virtual Port Channel

■ Path: Router3_PolGrp

■ Encap: VLAN 1222

■ Side A IPv4 Primary: 172.16.22.252/24

■ Side A IPv4 Secondary: 172.16.22.254/24

■ Side B IPv4 Primary: 172.16.22.253/24

■ Side B IPv4 Secondary: 172.16.22.254/24

Note: Secondary IP address is common IP for both Side A and Side B, which is the floating IP for the vPC pair leaf nodes. External router uses the secondary IP as next-hop IP toward the ACI fabric.

Cisco Systems. All printed copies and duplicate soft copies are considered uncontrolled
and the original online version should be referred to for the latest version.

Page 198 of 215

10. Verify SVI is added and click OK.

Figure 316. Add SVI

11. Verify Interface Profile is added and click OK.

Figure 317. Create Node Profile

12. Verify Node Profile is added and click Next.

Figure 318. Verify Node Profile is added

13. Create External EPG Network by clicking + icon.

Figure 319. Create External EPG Network
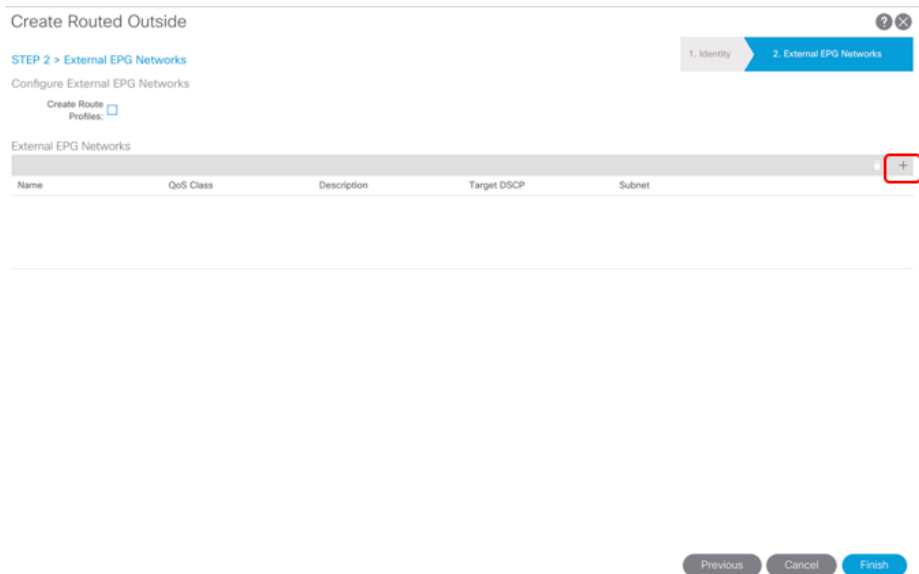
14. Create External EPG Network pop-up appears. Create Subnet by clicking + icon.

   ■ Name: L3Out2

Figure 320. Create External EPG Network
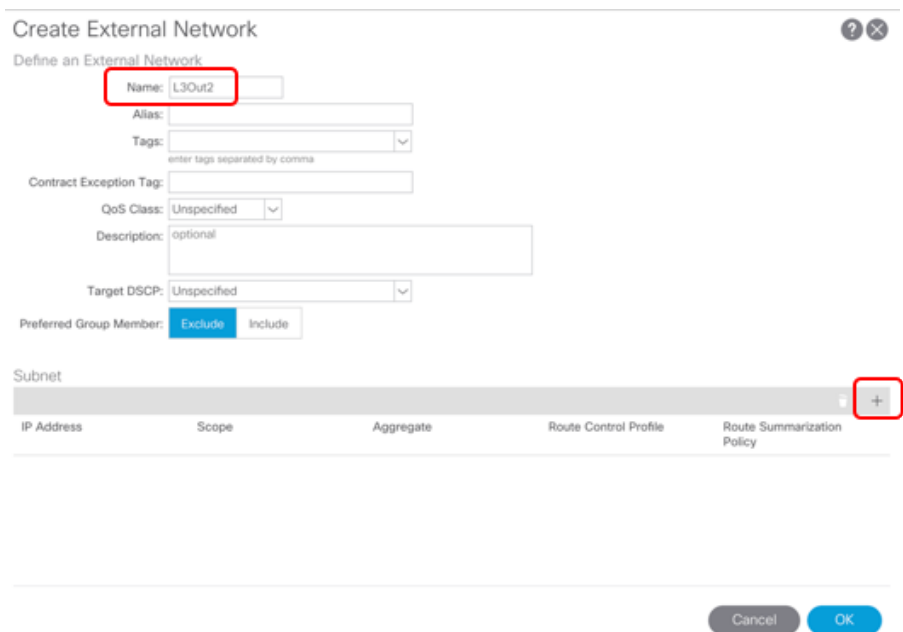
15. Create Subnet pop-up appears. Similar to L3out1 EPG subnet with " External Subnets for the External EPG" configuration, you are going to define L3out2 EPG subnet. Specify 172.16.20.0/24 that means 172.16.20.0/24 subnet in this VRF is classified to the L3Out2 EPG.

   ■ IP Address: 172.16.20.0/24

   ■ Scope: External Subnets for the External EPG

Create External EPG Network for L3Out2

Create Subnet

Specify the Subnet

IP Address: 172.16.20.0/24

scope: ☐ Export Route Control Subnet
☐ Import Route Control Subnet
☑ External Subnets for the External EPG
☐ Shared Route Control Subnet
☐ Shared Security Import Subnet

aggregate: ☐ Aggregate Export
☐ Aggregate Import
☐ Aggregate Shared Routes

Route Control Profile:

Name	Direction

Cancel   OK

Note: Either "Shared Route Control Subnet" or "Shared Security Import Subnet" is NOT required for this use case as these are for inter-VRF route-leaking that is not covered in this document.

16. Verify 172.16.20.0/24 is added and click OK.

Figure 321. Finish to define External EPG Network subnet

Create External Network

Define an External Network

Name: L3Out2
Alias:
Tags:
      enter tags separated by comma
Contract Exception Tag:
QoS Class: Unspecified
Description: optional

Target DSCP: Unspecified
Preferred Group Member: Exclude Include

Subnet

IP Address	Scope	Aggregate	Route Control Profile	Route Summarization Policy
172.16.20.0/24	External Subnets for the External EPG			

Cancel   OK

Cisco Systems. All printed copies and duplicate soft copies are considered uncontrolled and the original online version should be referred to for the latest version.

Page 202 of 215

17. Verify External EPG Network L3Out2 is added and click Finish.

Figure 322. Finish to create L3Out2 and External EPG Network
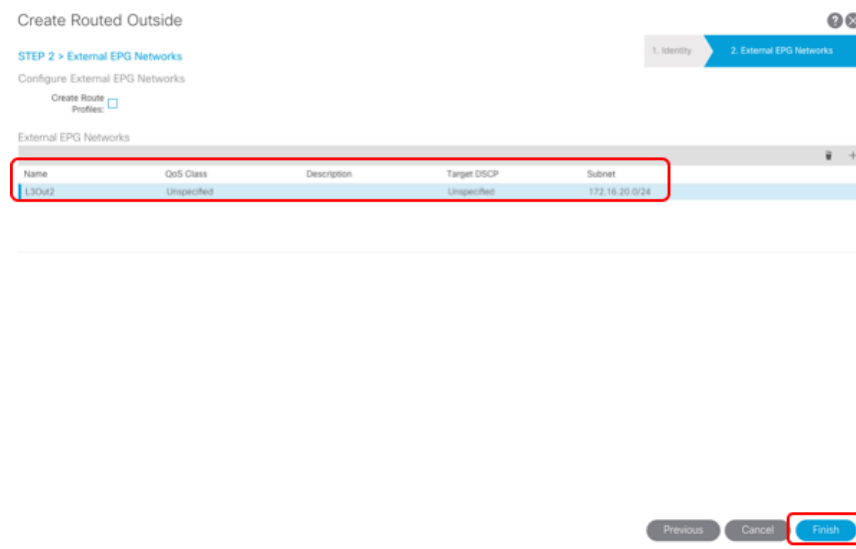
18. Once you click Finish, you can see External EPG L3Out2 under the L3Out2.

Figure 323. Verify L3Out2 is created
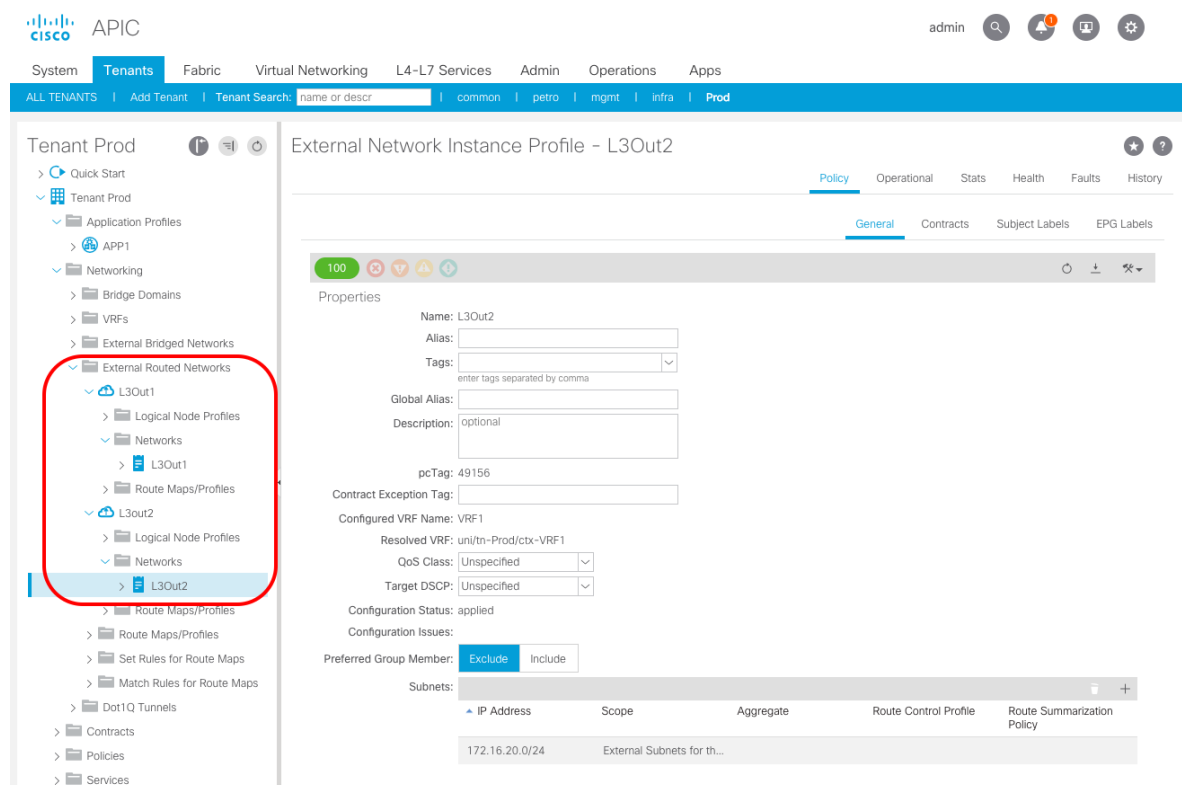
Border leaf nodes (Leaf105 and Leaf106) have SVI interfaces and the static route created.

```
leaf105# show ip interface vrf Prod:VRF1

IP Interface Status for VRF "Prod:VRF1"
```

Cisco Systems. All printed copies and duplicate soft copies are considered uncontrolled
and the original online version should be referred to for the latest version.

Page 203 of 215

```
vlan16, Interface status: protocol-up/link-up/admin-up, iod: 97, mode:
external

  IP address: 172.16.22.252, IP subnet: 172.16.22.0/24

  IP address: 172.16.22.254, IP subnet: 172.16.22.0/24 secondary

  IP broadcast address: 255.255.255.255

  IP primary address route-preference: 1, tag: 0

lo3, Interface status: protocol-up/link-up/admin-up, iod: 98, mode:
unspecified

  IP address: 10.1.1.5, IP subnet: 10.1.1.5/32

  IP broadcast address: 255.255.255.255

IP primary address route-preference: 1, tag: 0



leaf105# show ip route vrf Prod:VRF1

IP Route Table for VRF "Prod:VRF1"

'*' denotes best ucast next-hop

'**' denotes best mcast next-hop

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

'%<string>' in via output denotes VRF <string>



0.0.0.0/0, ubest/mbest: 4/0

    *via 10.0.224.69%overlay-1, [200/1], 00:02:10, bgp-65551, internal,
tag 65551

    *via 10.0.224.70%overlay-1, [200/1], 00:02:10, bgp-65551, internal,
tag 65551

    *via 10.0.224.64%overlay-1, [200/1], 00:02:10, bgp-65551, internal,
tag 65551

    *via 10.0.224.67%overlay-1, [200/1], 00:02:10, bgp-65551, internal,
tag 65551

10.1.1.1/32, ubest/mbest: 1/0

    *via 10.0.224.64%overlay-1, [1/0], 00:02:10, bgp-65551, internal, tag
65551
```

```
10.1.1.2/32, ubest/mbest: 1/0

    *via 10.0.224.70%overlay-1, [1/0], 00:02:10, bgp-65551, internal, tag
65551

10.1.1.3/32, ubest/mbest: 1/0

    *via 10.0.224.67%overlay-1, [1/0], 00:02:10, bgp-65551, internal, tag
65551

10.1.1.4/32, ubest/mbest: 1/0

    *via 10.0.224.69%overlay-1, [1/0], 00:02:10, bgp-65551, internal, tag
65551

10.1.1.5/32, ubest/mbest: 2/0, attached, direct

    *via 10.1.1.5, lo3, [1/0], 00:02:11, local, local

    *via 10.1.1.5, lo3, [1/0], 00:02:11, direct

10.1.1.6/32, ubest/mbest: 1/0

    *via 10.0.224.71%overlay-1, [1/0], 00:02:10, bgp-65551, internal, tag
65551

172.16.20.0/24, ubest/mbest: 1/0

    *via 172.16.22.1, vlan16, [1/0], 00:02:11, static

172.16.21.0/30, ubest/mbest: 1/0

    *via 10.0.224.64%overlay-1, [200/0], 00:02:10, bgp-65551, internal,
tag 65551

172.16.21.4/30, ubest/mbest: 1/0

    *via 10.0.224.70%overlay-1, [200/0], 00:02:10, bgp-65551, internal,
tag 65551

172.16.21.8/30, ubest/mbest: 1/0

    *via 10.0.224.67%overlay-1, [200/0], 00:02:10, bgp-65551, internal,
tag 65551

172.16.21.12/30, ubest/mbest: 1/0

    *via 10.0.224.69%overlay-1, [200/0], 00:02:10, bgp-65551, internal,
tag 65551

172.16.21.16/30, ubest/mbest: 4/0

    *via 10.0.224.69%overlay-1, [200/8], 00:02:10, bgp-65551, internal,
tag 65551
```

```
    *via 10.0.224.70%overlay-1, [200/8], 00:02:10, bgp-65551, internal,
tag 65551

    *via 10.0.224.64%overlay-1, [200/8], 00:02:10, bgp-65551, internal,
tag 65551

    *via 10.0.224.67%overlay-1, [200/8], 00:02:10, bgp-65551, internal,
tag 65551

172.16.22.0/24, ubest/mbest: 2/0, attached, direct

    *via 172.16.22.252, vlan16, [1/0], 00:02:11, direct

    *via 172.16.22.254, vlan16, [1/0], 00:02:11, direct

172.16.22.252/32, ubest/mbest: 1/0, attached

    *via 172.16.22.252, vlan16, [1/0], 00:02:11, local, local

172.16.22.254/32, ubest/mbest: 1/0, attached

    *via 172.16.22.254, vlan16, [1/0], 00:02:11, local, local
```

Other leaf nodes (Leaf101, 102, 103 and 104) learn external route 172.16.20.0/24 via BGP in ACI fabric.

```
leaf101# show ip route vrf Prod:VRF1

IP Route Table for VRF "Prod:VRF1"

'*' denotes best ucast next-hop

'**' denotes best mcast next-hop

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

'%<string>' in via output denotes VRF <string>


0.0.0.0/0, ubest/mbest: 1/0

    *via 172.16.21.2, eth1/33.23, [110/1], 00:47:10, ospf-default, type-2

10.1.1.1/32, ubest/mbest: 2/0, attached, direct

    *via 10.1.1.1, lo3, [1/0], 00:47:22, local, local

    *via 10.1.1.1, lo3, [1/0], 00:47:22, direct

10.1.1.2/32, ubest/mbest: 1/0

    *via 10.0.224.70%overlay-1, [1/0], 00:47:22, bgp-65551, internal, tag
65551
```

```
10.1.1.3/32, ubest/mbest: 1/0
    *via 10.0.224.67%overlay-1, [1/0], 00:47:23, bgp-65551, internal, tag
65551

10.1.1.4/32, ubest/mbest: 1/0
    *via 10.0.224.69%overlay-1, [1/0], 00:47:23, bgp-65551, internal, tag
65551

10.1.1.5/32, ubest/mbest: 1/0
    *via 10.0.224.68%overlay-1, [1/0], 00:03:43, bgp-65551, internal, tag
65551

10.1.1.6/32, ubest/mbest: 1/0
    *via 10.0.224.71%overlay-1, [1/0], 00:03:43, bgp-65551, internal, tag
65551

172.16.20.0/24, ubest/mbest: 2/0
    *via 10.0.224.68%overlay-1, [1/0], 00:03:43, bgp-65551, internal, tag
65551

    *via 10.0.224.71%overlay-1, [1/0], 00:03:43, bgp-65551, internal, tag
65551

172.16.21.0/30, ubest/mbest: 1/0, attached, direct
    *via 172.16.21.1, eth1/33.23, [1/0], 00:47:22, direct

172.16.21.1/32, ubest/mbest: 1/0, attached
    *via 172.16.21.1, eth1/33.23, [1/0], 00:47:22, local, local

172.16.21.4/30, ubest/mbest: 1/0
    *via 172.16.21.2, eth1/33.23, [110/8], 00:47:11, ospf-default, intra

172.16.21.8/30, ubest/mbest: 1/0
    *via 172.16.21.2, eth1/33.23, [110/12], 00:47:11, ospf-default, intra

172.16.21.12/30, ubest/mbest: 1/0
    *via 172.16.21.2, eth1/33.23, [110/12], 00:47:11, ospf-default, intra

172.16.21.16/30, ubest/mbest: 1/0
    *via 172.16.21.2, eth1/33.23, [110/8], 00:47:11, ospf-default, intra

172.16.22.0/24, ubest/mbest: 2/0
```

Cisco Systems. All printed copies and duplicate soft copies are considered uncontrolled
and the original online version should be referred to for the latest version.

Page 207 of 215

```
    *via 10.0.224.68%overlay-1, [200/0], 00:03:43, bgp-65551, internal,
tag 65551

    *via 10.0.224.71%overlay-1, [200/0], 00:03:43, bgp-65551, internal,
tag 65551

192.168.21.0/24, ubest/mbest: 1/0, attached, direct, pervasive

    *via 10.0.160.66%overlay-1, [1/0], 00:39:53, static

192.168.21.254/32, ubest/mbest: 1/0, attached, pervasive

    *via 192.168.21.254, vlan19, [1/0], 03:37:56, local, local
```

Border leaf node 101, 102, 103 or 104 doesn't advertise transit route 172.16.20.0/24 to external. Thus, external Router1 or 2 doesn't learn 172.16.20.0/24 at this point.

```
Router1# show ip route ospf vrf Prod:VRF1

IP Route Table for VRF "Prod:VRF1"

'*' denotes best ucast next-hop

'**' denotes best mcast next-hop

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

'%<string>' in via output denotes VRF <string>


10.1.1.1/32, ubest/mbest: 1/0

    *via 172.16.21.1, Eth1/1.1202, [110/5], 00:50:11, ospf-1, intra

10.1.1.2/32, ubest/mbest: 1/0

    *via 172.16.21.5, Eth1/2.1202, [110/5], 00:50:11, ospf-1, intra

10.1.1.3/32, ubest/mbest: 1/0

    *via 172.16.21.18, Eth1/17.1202, [110/9], 00:50:11, ospf-1, intra

10.1.1.4/32, ubest/mbest: 1/0

    *via 172.16.21.18, Eth1/17.1202, [110/9], 00:50:11, ospf-1, intra

172.16.21.8/30, ubest/mbest: 1/0

    *via 172.16.21.18, Eth1/17.1202, [110/8], 03:29:52, ospf-1, intra

172.16.21.12/30, ubest/mbest: 1/0

    *via 172.16.21.18, Eth1/17.1202, [110/8], 03:29:52, ospf-1, intra

192.168.21.0/24, ubest/mbest: 2/0
```

```
 *via 172.16.21.1, Eth1/1.1202, [110/20], 00:41:54, ospf-1, type-2

 *via 172.16.21.5, Eth1/2.1202, [110/20], 00:41:54, ospf-1, type-2
```

Configure Export Route Control Subnet

To advertise transit route, Export Route control subnet needs to be configured. You are going to configure Export Route control subnet in L3Out1 EPG to advertise 172.16.20.0/24.

1. From Tenant > Networking > External Routed Networks > L3Out1 > Networks > L3Out1, add Subnet by clicking + icon.

Figure 324. Add Export Route Control Subnet in L3Out1 EPG

2. Create Subnet pop-up appears. Specify 172.16.20.0/24 with "Export Route Control Subnet" checked and "External Subnets for the External EPG" unchecked.

    ■ IP Address: 172.16.20.0/24

    ■ Scope: Export Route Control Subnet

**Cisco Systems. All printed copies and duplicate soft copies are considered uncontrolled
and the original online version should be referred to for the latest version.**

Page 209 of 215

3. Verify 172.16.20.0/24 with scope Export Route control Subnet is added.

Figure 326. Verify Subnets

Border leaf nodes 101, 102, 103 and 104 start advertising transit route 172.16.20.0/24. Thus, external Router1 and 2 learn 172.16.20.0/24.

```
Router1# show ip route ospf vrf Prod:VRF1

IP Route Table for VRF "Prod:VRF1"

'*' denotes best ucast next-hop

'**' denotes best mcast next-hop
```

```
'[x/y]' denotes [preference/metric]

'%<string>' in via output denotes VRF <string>


10.1.1.1/32, ubest/mbest: 1/0

    *via 172.16.21.1, Eth1/1.1202, [110/5], 00:54:24, ospf-1, intra

10.1.1.2/32, ubest/mbest: 1/0

    *via 172.16.21.5, Eth1/2.1202, [110/5], 00:54:24, ospf-1, intra

10.1.1.3/32, ubest/mbest: 1/0

    *via 172.16.21.18, Eth1/17.1202, [110/9], 00:54:24, ospf-1, intra

10.1.1.4/32, ubest/mbest: 1/0

    *via 172.16.21.18, Eth1/17.1202, [110/9], 00:54:24, ospf-1, intra

172.16.20.0/24, ubest/mbest: 2/0

    *via 172.16.21.1, Eth1/1.1202, [110/1], 00:00:57, ospf-1, type-2, tag
4294967295

    *via 172.16.21.5, Eth1/2.1202, [110/1], 00:00:57, ospf-1, type-2, tag
4294967295

172.16.21.8/30, ubest/mbest: 1/0

    *via 172.16.21.18, Eth1/17.1202, [110/8], 03:34:05, ospf-1, intra

172.16.21.12/30, ubest/mbest: 1/0

    *via 172.16.21.18, Eth1/17.1202, [110/8], 03:34:05, ospf-1, intra

192.168.21.0/24, ubest/mbest: 2/0

    *via 172.16.21.1, Eth1/1.1202, [110/20], 00:46:07, ospf-1, type-2

    *via 172.16.21.5, Eth1/2.1202, [110/20], 00:46:07, ospf-1, type-2
```
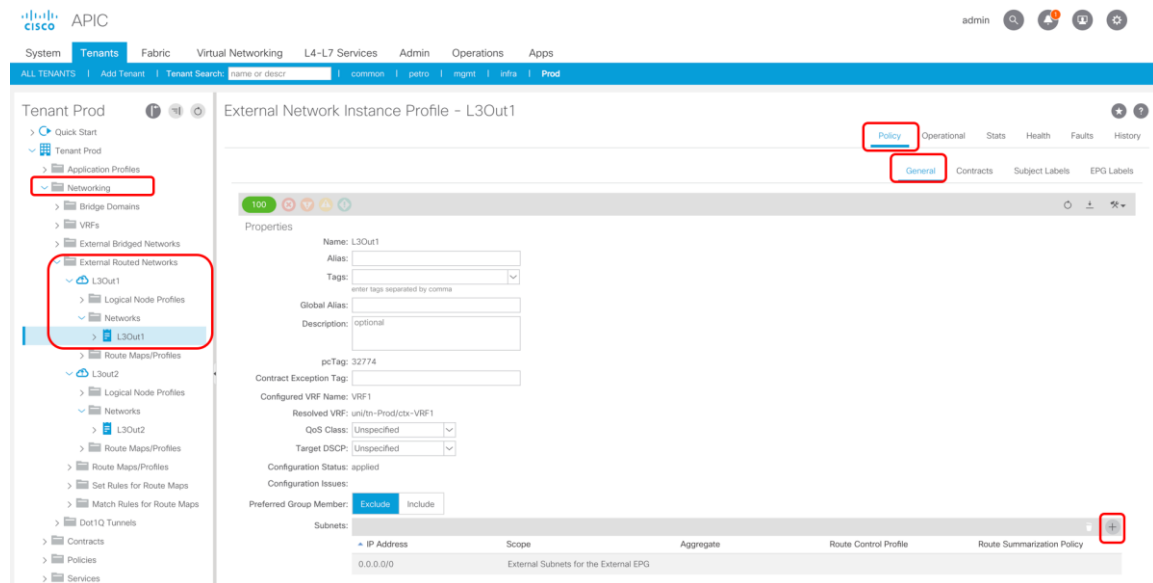
Note: In this example, we don't add Export Route Control Subnet in L3Out2 EPG as L3Out2 uses static route. If you need to advertise a transit route to L3Out2, you need to add Export Route Control Subnet in L3Out2 EPG too.

Configure a contract between the L3outs

Then, you are going to create a contract External for between External EPG L3Out1 and External EPG L3Out2.

1. From Tenant > Contracts, Create Contract.

    In this example, we are going to reuse a filter that was created in previous section.

Cisco Systems. All printed copies and duplicate soft copies are considered uncontrolled
and the original online version should be referred to for the latest version.

Page 211 of 215

- Name: Transit

- Add subject by clicking + icon. Then, Create Contract Subject pop-up appears.

- Subject name: sbjct-permit-all

  - Check Apply Both Directions and Reverse Filter Ports (default)

- **Add a filter by clicking + icon**
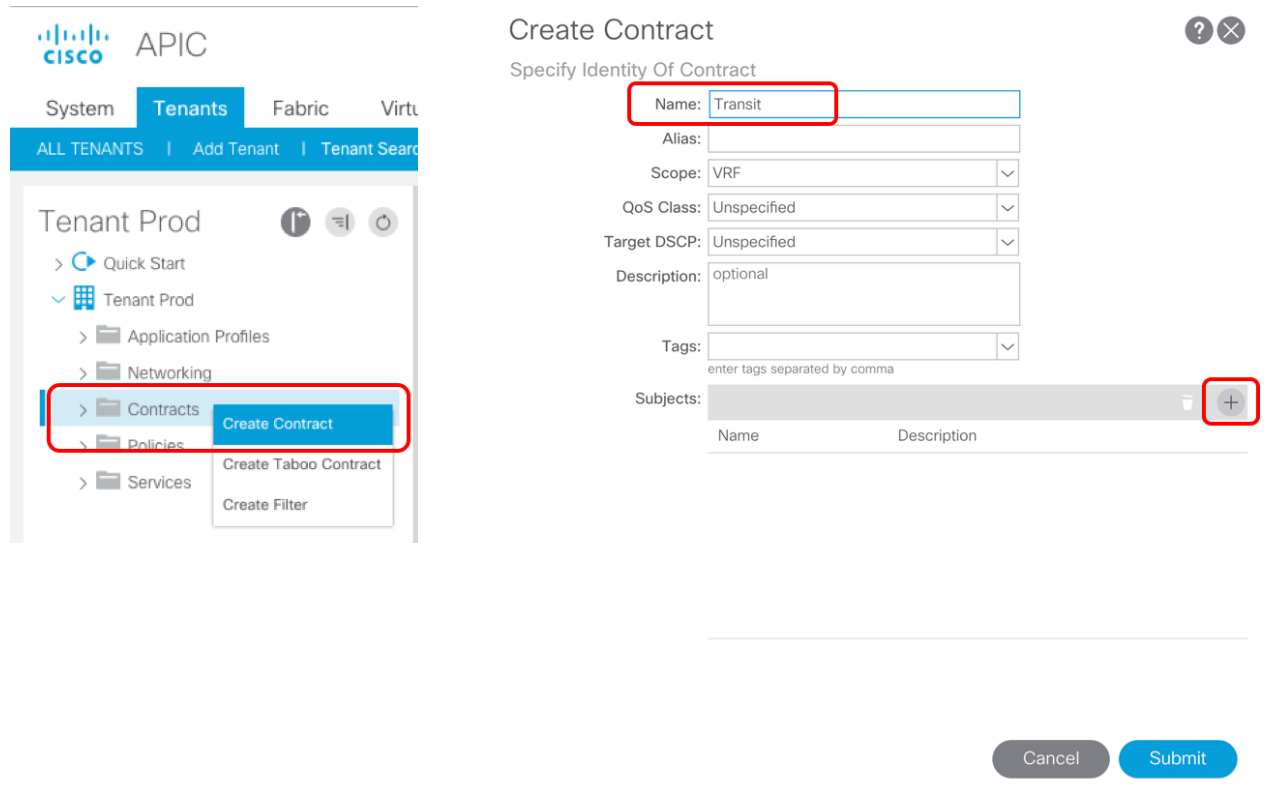
  - Filter: ftr-permit-all

Figure 327. Create contract

Cisco Systems. All printed copies and duplicate soft copies are considered uncontrolled
and the original online version should be referred to for the latest version.

Page 212 of 215

Figure 328. Create contract subject

2. Verify the subject is added and click Submit.

Figure 329. Create contract
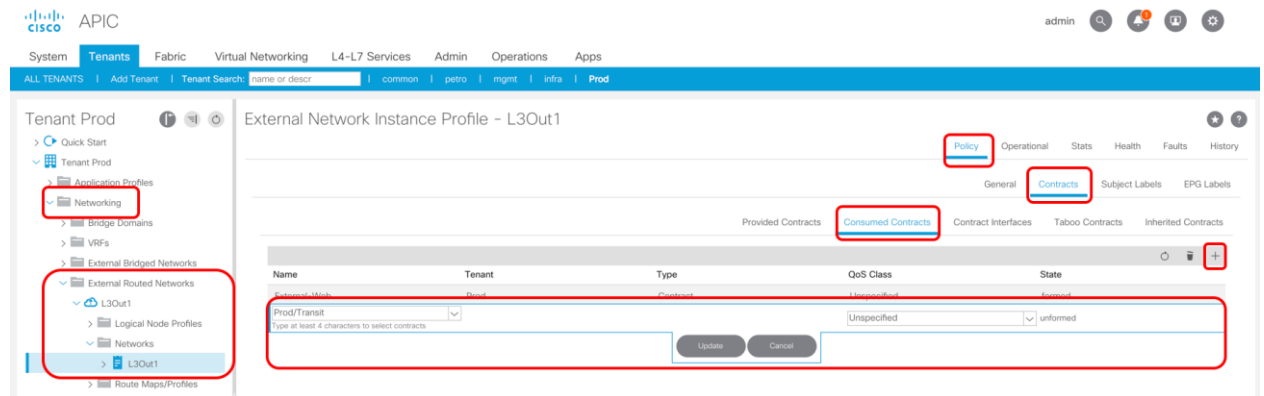

3. Then, configure contract for External EPG L3Out1 as consumer and for External EPG L3Out2 as provider.

For External EPG L3Out1, the location is at Tenant > Networking> External Routed Networks > L3Out1 > Networks > L3Out1 > Policy > Contracts > Consumed Contracts.

Figure 330. Configure contract for External EPG L3Out1

For External EPG L3Out2, the location is at Tenant > Networking> External Routed Networks > L3Out2 > Networks > L3Out2 > Policy > Contracts > Provided Contracts.

Figure 331. Configure contract for External EPG L3Out2
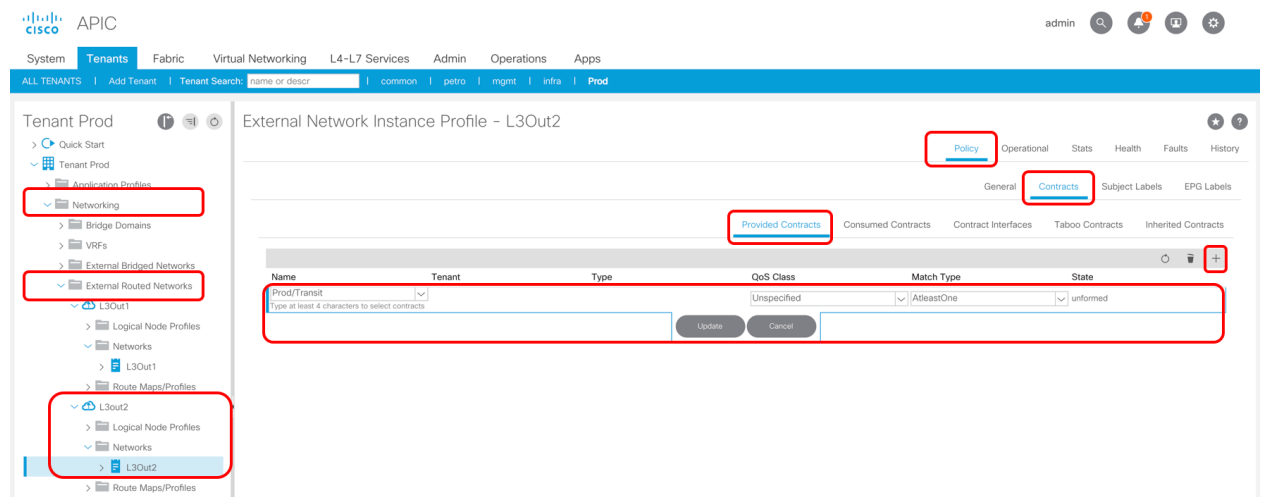
Once a contract between L3Out EPGs are configured, IP in L3Out1 can communicate with IP in L3Out2. An external client 172.16.10.1 behind L3Out1 can ping to an external client 172.16.20.1 behind L3Out2.

Cisco Systems. All printed copies and duplicate soft copies are considered uncontrolled
and the original online version should be referred to for the latest version.

Page 214 of 215
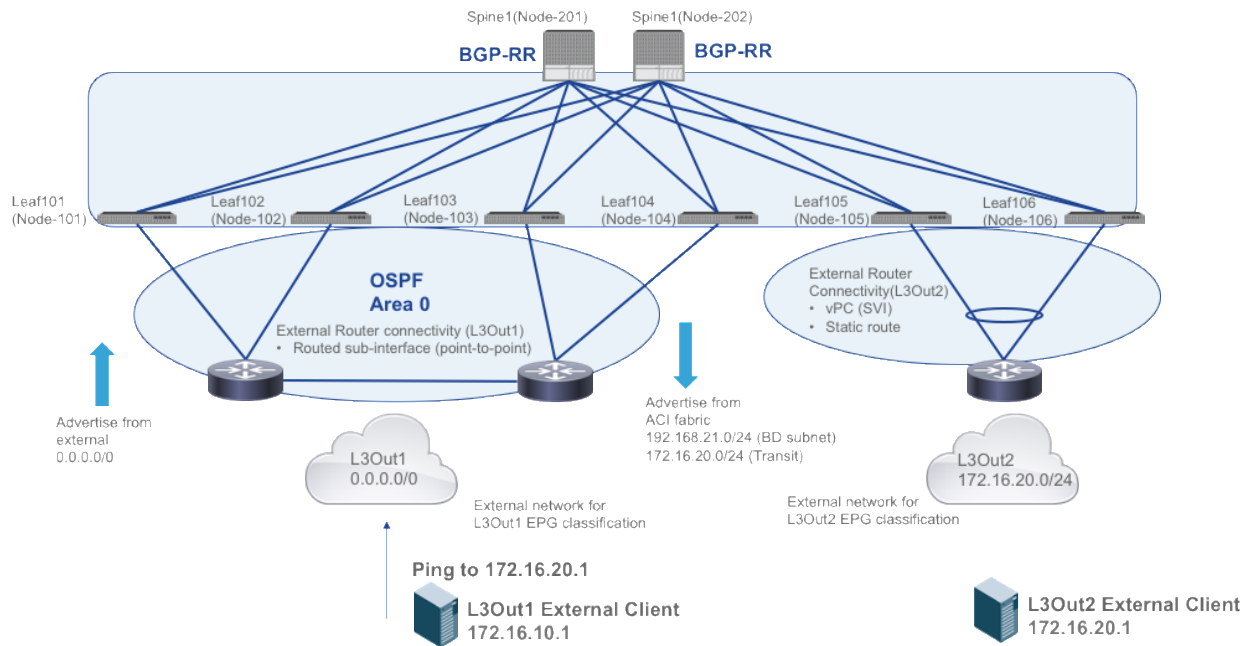
Figure 332. External connectivity

Figure 333. Ping from an external client behind L3Out1 can ping to an external client behind L3Out2

```
[root@rescue /]# ifconfig eth0
eth0      Link encap:Ethernet  HWaddr 00:50:56:B6:07:F1
          inet addr:172.16.10.1  Bcast:172.16.10.255  Mask:255.255.255.0
          inet6 addr: fe80::250:56ff:feb6:7f1/64 Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:341 errors:0 dropped:0 overruns:0 frame:0
          TX packets:5576 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:21980 (21.4 Kb)  TX bytes:1916462 (1.8 Mb)

[root@rescue /]# ping 172.16.20.1
PING 172.16.20.1 (172.16.20.1) 56(84) bytes of data.
64 bytes from 172.16.20.1: icmp_seq=1 ttl=251 time=0.503 ms
64 bytes from 172.16.20.1: icmp_seq=2 ttl=251 time=0.519 ms
64 bytes from 172.16.20.1: icmp_seq=3 ttl=251 time=0.541 ms
64 bytes from 172.16.20.1: icmp_seq=4 ttl=251 time=0.437 ms
64 bytes from 172.16.20.1: icmp_seq=5 ttl=251 time=0.418 ms
64 bytes from 172.16.20.1: icmp_seq=6 ttl=251 time=0.582 ms
```

If you remove the contract "Transit" from either L3Out1 EPG or L3Out2 EPG, ping stops working. If you add it back, ping starts working again.