Cisco Application Centric Infrastructure

Learning ACI: Bringing Up A Fabric

This ACI resource will guide you through the process for provisioning a fabric and bringing it online. Before you begin, ensure that everything is labeled correctly (leaf nodes to spine nodes, APIC controllers to leaf nodes, APIC out of band connectivity, etc).

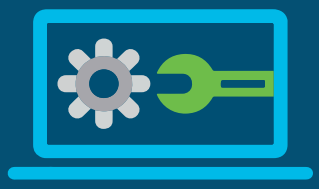

Before you start, review these parameters and have them handy during the setup process.

- *Controller names* are the names you want to give your APIC controllers. (Normally there will be three of these.)
- *TEP IP address pool* will be used by the fabric to automatically allocate Tunnel End Point (TEP) addresses. Every node (leaf and spine) will be allocated at least one TEP address. It's generally a good practice to dedicate a unique range for this.
- *Multicast address pool* is a pool of addresses used for multicast traffic through the fabric.
- *VLAN ID for the infrastructure network* is the VLAN ID used for internal communication in the fabric, but it also may be extended outside the fabric for communication to AVS virtual switches. Therefore, it's best to allocate an unused VLAN. (Be aware that this VLAN overlaps with the "reserved" range on some other Nexus platforms.)
- *Management IP address and gateway* is the out-of-band address you will use to browse to the APIC itself.

You will usually use the CIMC utility (essentially a UCS C series server) to access a KVM session, so you can set up the APIC. Once you have started the setup process, you'll see a simple script which will ask you for the parameters above.

```
 Cluster configuration ...
  Enter the fabric name [ACI Fabric1 #1]:
  Enter the number of controllers in the fabric (1-16) [3]:
  Enter the controller ID (1-3) [2]:
  Enter the controller name [apic2]:
  Enter address pool for TEP addresses [10.0.0.0/16]:
  Enter the VLAN ID for infra network (1-4094)[4093]:
  Enter address pool for BD multicast addresses (GIPO) [225.0.0.0/15]:

Out-of-band management configuration ...
  Enter the IP address for out-of-band management: 192.168.10.2/24
  Enter the IP address of the default gateway [None]: 192.168.10.254
  Enter the interface speed/duplex mode [auto]:

Administrator user configuration...
  Enable strong passwords? [Y]
  Enter the password for admin:
```

Once you have completed the script, the APIC initializes and after a short period of time, you should be able to browse to it using the out-of-band address you configured earlier. Now that you have a single APIC up and running, you can start to discover the physical switching nodes in your fabric.

In the APIC GUI, click on the **Fabric** tab and then the **Inventory** sub-menu. On the left hand side, click on **Fabric Membership**. If all is well, you should see a single leaf node listed in here; this is the physical switch that your APIC is connected to. Notice that there is no node ID or node name listed, and the IP address is set to 0.0.0.0. That means that this switch is not yet registered. In order to register the node and add it to the network, you simply need to assign a node ID and node name by double clicking on those fields. Once done, you should (after a minute or so) see an IP address appear for that node. Notice that this address is taken from the TEP range you specified during the setup process. This means that the node is now "registered."

Now that you have a single leaf node registered you can discover the rest of our network? With the first node registered, you should start to see additional switches (spine nodes) pop up on the Fabric Membership page. Register these in the same way, which in turn will lead to the remaining switches appearing and being available for registration. The end result will look similar to the below.

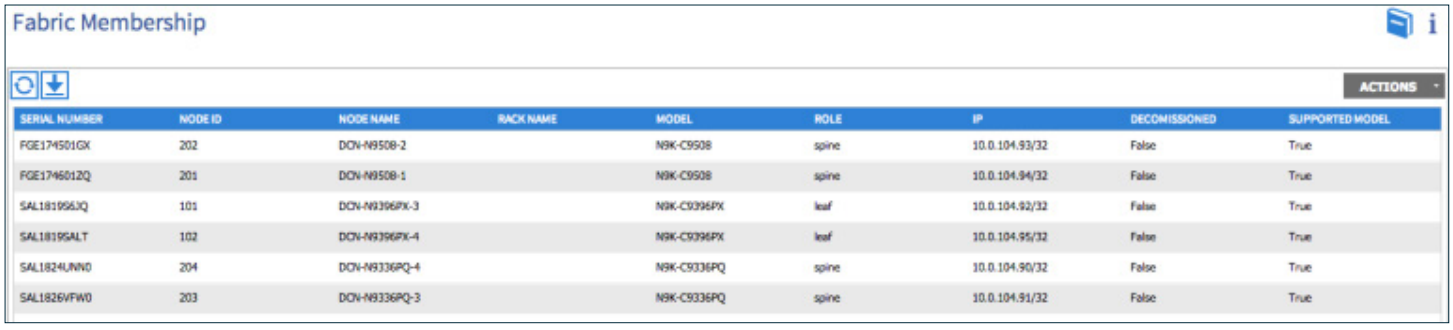

SERIAL NUMBER	NODE ID	NODE NAME	RACK NAME	MODEL	ROLE	IP	DECOMISSIONED	SUPPORTED MODEL
FGE174501GX	202	DCN-N9508-2		N9K-C9508	spine	10.0.104.93/32	False	True
FGE174601ZQ	201	DCN-N9508-1		N9K-C9508	spine	10.0.104.94/32	False	True
SAL1819S6JQ	101	DCN-N9396PX-3		N9K-C9396PX	leaf	10.0.104.92/32	False	True
SAL1819SALT	102	DCN-N9396PX-4		N9K-C9396PX	leaf	10.0.104.95/32	False	True
SAL1824UNN0	204	DCN-N9336PQ-4		N9K-C9336PQ	spine	10.0.104.90/32	False	True
SAL1826VFW0	203	DCN-N9336PQ-3		N9K-C9336PQ	spine	10.0.104.91/32	False	True

At this point, you have your full switching topology discovered, but you still only have one APIC. As your other APICs are very likely to be connected to different leaf switches, you needed to wait until those switches were discovered before you could initialize the remaining APICs and form the cluster. Now you can run the setup script on the other APICs, using the same basic information as before. Make sure you use different controller IDs, management IP address, etc.

Now that your remaining APICs are up and running, your fabric is ready to go. Go to the **Fabric | Inventory** section in the APIC GUI and click on **Topology** for a full graphical representation of your fabric.

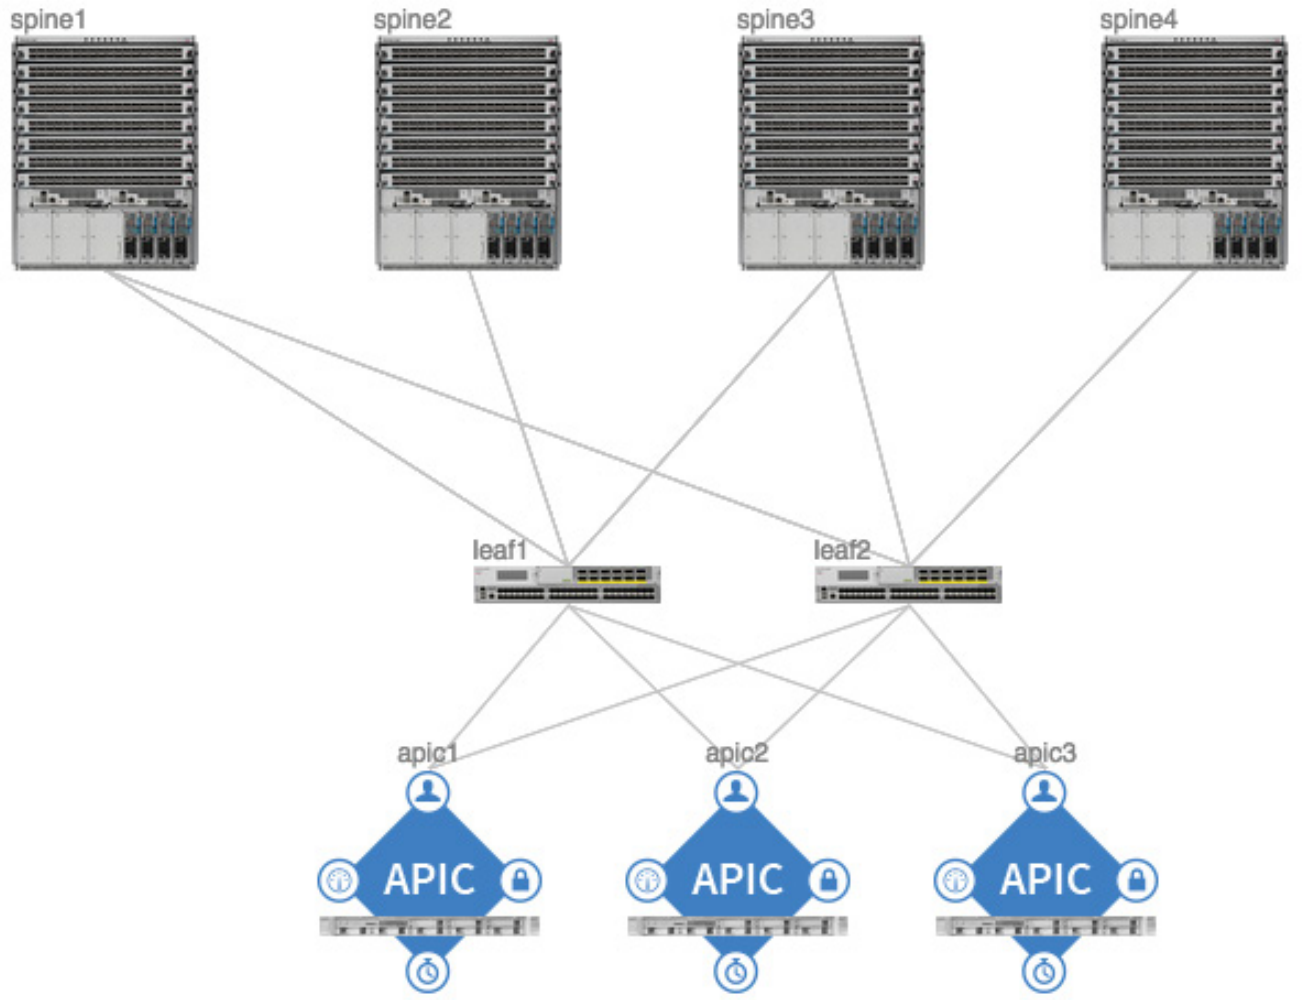

Before you go any further, you may wish to configure your switching nodes with management IP addresses so you can SSH to them directly. Do this by clicking on the **Tenants** tab and then the "Mgmt" tenant. On the left hand side you'll find **Node Management Addresses**, which will allow you to configure management IPs for every node in the fabric. Also under the "Mgmt" tenant, you'll find **Out-Of-Band Contracts** listed under the **Security Policies** menu. You'll need to configure at least one contract here to allow traffic into your OOB management interfaces. Finally, under **Node Management EPGs**, you can add the OOB contract to your OOB EPG.

What is really happening during discovery?

Now that you understand how to discover your fabric, it's time to understand exactly how this process is happening. We use LLDP to determine what is connected to each port. The process is as follows:

1. When the first leaf initializes, it sends LLDP messages on its active ports to determine which port the APIC is connected to.
2. The leaf receives an LLDP message back from the APIC indicating the node type.
3. The leaf node sends a DHCP Discover message to the APIC. The APIC responds with an IP address from a locally configured pool (the pool you configured during the APIC setup script).
4. The leaf node uses this IP address as the Tunnel End Point (TEP) address.
5. Once the first leaf node has been discovered and registered by the APIC, a similar process is followed for the remaining devices in the fabric.