v1.0

Cisco SD-Access Lab Workbook

Mason Reimert
masonreimert.com

Buy Me a Coffee

Cisco SD-Access Overview

Cisco SD-Access is a campus technology that allows you to build LISP overlays on top of a routed campus network. All links in an SD-Access network are routed, with no spanning tree. There are four primary node types in Cisco SD-Access:

Edge Nodes (HQ-EDGE-1, HQ-EDGE-2) – These are the switches that your clients connect to. These switches encapsulate the client traffic into VXLAN tunnels to transport the traffic across the overlay. These switches also act as the default gateways for every VLAN in the campus site.

Control Plane Nodes (HQ-BCP-1, HQ-BCP-2) – Control plane nodes are routers or switches that are responsible for maintain the LISP control plane for the fabric. They maintain mappings of EIDs (Endpoint IPs) to RLOCs (Router Locators). When a client sends traffic to a new destination the edge node queries the control plane node to ask what switch the destination sits behind.

Border Nodes (HQ-BCP-1, HQ-BCP-2) – Border nodes can be collocated on control plane nodes. A border node sits between the SD-Access fabric and external networks such as a datacenter or the internet. In our lab, the border node takes routes from the fusion router and sends them into the fabric, and vice versa.

Intermediate Nodes (HQ-IN-1, HQ-IN-2) – Intermediate nodes are nodes that pass traffic in the underlay but are not aware of the SD-Access fabric riding over top. These nodes can be any switch or router that is capable of participating in the routing protocol you are using for your underlay.

Workbook Lab Overview

This lab is a single site SD-Access fabric. It allows you to get a foundation for SD-Access concepts by using a virtual topology. If this is the first time you have worked with SD-Access I suggest starting by working through the answer key, then later completing the tasks without looking at the answer key. The lab consists of two edge nodes where clients can connect, two collocated border/control plane nodes, and two intermediate nodes. The fabric uplinks to a fusion router that provides a way for decapsulated traffic to exit the fabric and reach the datacenter where DNA Center and ISE reside.

This lab is best run on two servers. One server to host DNA Center and ISE, and another server to host CML. The reason behind using CML is that the virtualization in CML seems to provide better latency when virtualizing Catalyst 9kv switches compared to other network modeling software. You will need to install DNA Center and ISE before starting this lab. You should install them with the IP addresses from the table below, and on a vSwitch that you can bridge into CML. This process is very well documented including a video I made on another lab where I show how I install DNA Center and ISE and bridge them into the topology. In new versions of DNA Center, you also need to ensure you install the SD-Access app as it is not installed by default.

Server Requirements

By far, the most common question I get about Software Defined Access is what are the hardware requirements needed to lab. Unfortunately, I do not have the resources or bandwidth to make a comprehensive list of what will or won't work. But I can tell you what has worked for me. Firstly, I really do not recommend renting racks. I do not believe rack rentals work for a technology as complex as SDA. You need more time hands on with the fabric to build it, break it, and troubleshoot it. With that said, I know you can run all of this on one server. I choose to run across two servers with ISE and DNA on one, and CML on the other. With how heavy the c9kv nodes are, you really should have two servers.

My server specifications (each server):

- 2x Intel(R) Xeon(R) Silver 4114 CPU @ 2.20GHz
- 190GB of RAM
- 3TB of storage

You can definitely get by with less, as with this lab running, I hover around 75% RAM usage, 50% CPU usage, and 40% storage usage. The problem you will run into is astronomical boot times for DNA and ISE. You can learn more about my personal setup in this video.

Software Versions

Software	Version	Reasoning
Cisco Modeling Labs	>= 2.7	2.7 is the first version of CML with the c9kv included in the refplat
Cisco ISE	>= 3.1	3.1 is the version this was tested with in the lab, it may be possible to go older
Cisco DNA Center	>= 2.3.5.5	2.3.5.5 is when support for c9kv was included in the DNA device packs, anything less will not work

Prerequisites

DNA Center and ISE must be installed and connected to a port group/virtual switch that is bridged into CML, if you need help bridging into CML, I go over that here.
Download the lab YAML file from here and load it into CML.
You must have a way to bridge internet access into CML, or otherwise transfer the installer for the wpa_supplicant Linux software to all the workstation hosts. I made a separate video on that here. It is also covered in the respective lab task.
The SD-Access application needs to be installed into DNA Center. Starting in version 2.X SD-Access is not installed by default. This is an offline process and can be completed in about 30 minutes.
You need to swap out the external connectors in the lab topology with bridges that are valid in your environment. The bridge to DNAC/ISE needs to be connected to the L2 segment with DNAC/ISE attached. The bridge for internet needs to be connected via bridge or NAT to the internet.

Pointers

If you start with the lab YAML you do not need to address links or set hostnames, that is done for you.
Please ensure you have the correct version of the c9kv. Using the wrong version will get you very far into the lab but will be a waste of time when it does not adopt into DNA. The version I use is C9KV-UADP-8P / 17.10.20220531. Also, verify your serial numbers are unique across c9kvs upon first boot.
If this is your first time using bridges in CML, you should make sure your vSwitches in your hypervisor have promiscuous mode and forged transmits allowed to prevent bridge issues.

Buy Me a Coffee

Topology

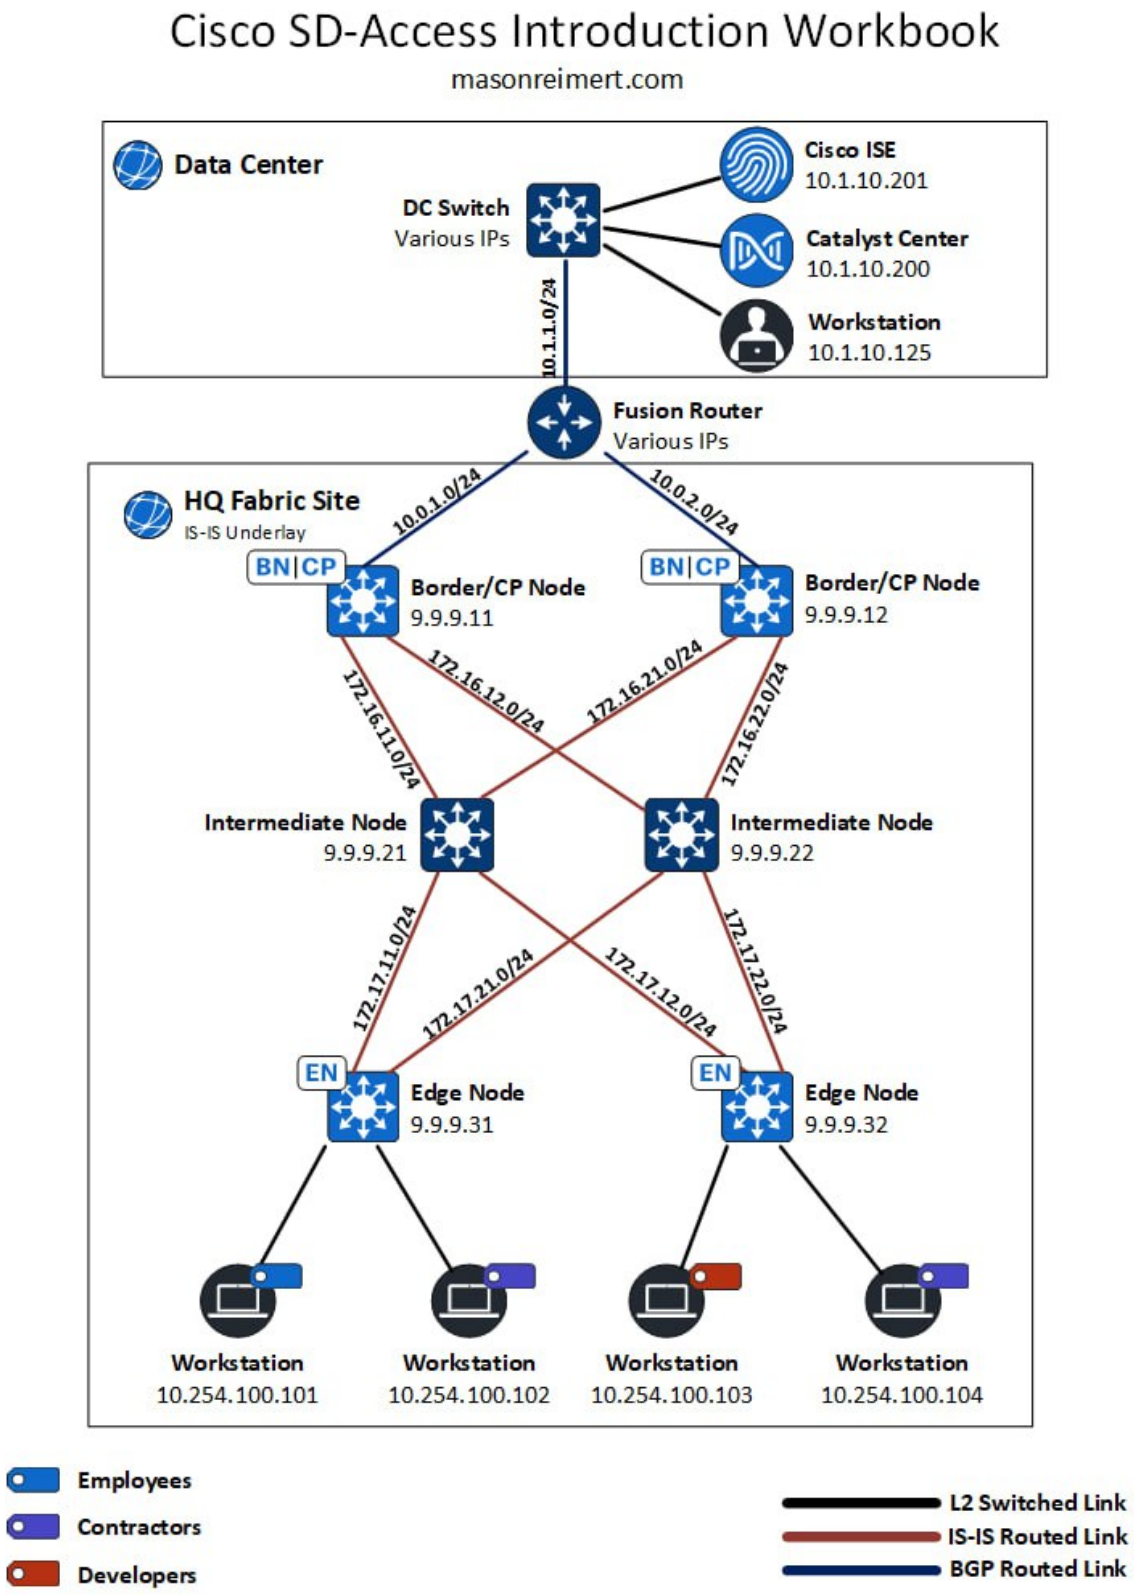

Lab Tasks

Task 1 – Datacenter Routing:
Establish a BGP peering between DC-SW-1 (AS 650001) and FUSION-1 (AS 650002) interface addresses and advertise the datacenter subnet to the fusion router.

Task 2 – Fusion Routing:
Configure BGP between the fusion router (AS 650002) and both border nodes (AS 650003) based on interface addresses. This task is for underlay reachability and only peers the underlay global VRF, not any overlay VNs.

Task 3 – Internal Underlay Routing:
Run IS-IS on all underlay links in the fabric (172.X.X.X) and loopback0 interfaces (9.9.9.X). At this point you should have underlay reachability to all lookback interfaces within the fabric domain. The IS-IS net statement should be XXX where XX is the last octet of the loopback0 address.

Task 4 – Redistribution
On the border nodes, redistribute IS-IS into BGP and vice versa. Also, advertise the loopbacks of the border nodes into BGP because they will not redistribute via IS-IS. You should now be able to ping from your loopback0 interfaces to ISE and Catalyst Center.

Task 5 – Fabric Node Bootstrapping
Perform the necessary tasks to enable device manageability from Catalyst Center on the fabric nodes (BCP/EDGE Nodes).

- Enable SNMP with RW community “cisco”
- Define the AAA method list for exec authorization
- Enable AAA and add user “dna” with password “ISEisC00L” with the highest privilege level
- Configure the VTY lines to allow SSH login with local authentication and the highest privilege level
- Enable NETCONF
- Set the device’s hostname respectively and the domain name to “ccie.local”

Task 6 – Discover Devices
Discover the four fabric enabled devices into Catalyst Center. You can either add the settings in the network hierarchy or input the device parameters upon each device add. Ensure to enable NETCONF manageability.
Hint: Access Catalyst Center via the VNC console of the admin workstation.

Task 7 – Network Hierarchy
Create an area named USA. Inside USA create a building named SDA-Campus.

Task 8 – Integrate Catalyst Center with ISE
Enable PXGrid on ISE, then proceed to create the integration between Catalyst Center and ISE. When complete, re-provision devices and ensure they show up under the network devices list in ISE.

Task 9 – Add AAA Settings to Building
Add ISE as the AAA server for client access under network settings in the hierarchy.

Task 10 – Provision Devices to Site
Provision the fabric nodes to the SDA-Campus site.

Task 11 – Create IP Address Pools
At the global level add the 10.254.0.0/16 address pool as “SDA-Campus”.
Then, reserve the following at the building level:

- 10.254.100.0/24 - “SDA-Campus-VLAN-100”
- 10.254.251.0/30 - “SDA-Campus-L3-Handoff-1”
- 10.254.252.0/30 - “SDA-Campus-L3-Handoff-2”

Task 12 – Create Fabric Site
Create a Fabric Site for SDA-Campus using closed authentication.

Task 13 – Create Transit Site
Create an IP-Based transit site named “TO-FUSION” with the Fusion Router’s ASN 650002.

Task 14 – Create VN
Create an SD-Access VN named CORP and assign it to your fabric site.

Task 15 – Provision Fabric Site
Create a fabric site using the four discovered devices with their respective roles from the topology diagram. Ensure to enable L3 handoff on both border nodes with local as 650003 and remote-as 650002. Use VLAN 101 to handoff BCP-1, and VLAN 102 to handoff BCP-2.
Keep the following in mind:

- Uncheck “Do not Import Internal Routes”

- Ensure that you add the transit site, external interface, and IP pool to the L3 Handoff on the Border Nodes
- This lab was tested using LISP/BGP not Pub/Sub, if you want to follow along 100% use LISP/BGP

Task 16– Configure IP Handoff

Configure sub interfaces on the fusion router for each L3 handoff VLAN created in the last task (101/102). Then, configure BGP peerings to both border nodes on the fusion router. Since we are intending to leak all routes from the GRT into the CORP VN, the peerings do not been to be VRF aware on the fusion side.

Task 17 – Create Anycast Gateways

Depending on which view you are using in Catalyst center either create “Anycast Gateways” or “VLANS” with the VLAN number of 100 using the previously reserved IP space “SDA-Campus-VLAN-100”.

Task 18 – Run TrustSec Migration

Migrate TrustSec administration to DNA center within the Group Based Policy menu.

Task 19 – Define Security Policy

Define the following group-based policy, then deploy the policy to ISE.

SRC DST	**Employees**	**Contractors**	**Developers**
Employees			
Contractors			
Developers			

Task 20 – Create Groups

Create the following user groups in ISE:

- SDA_Employees
- SDA_Contractors
- SDA_Developers

Task 21 – Create Users

Create the following users in ISE as local users with membership to the respective groups.

Username	Password	Group
bob	ISEisC00L	SDA_Employees
joe	ISEisC00L	SDA_Contractors
susan	ISEisC00L	SDA_Contractors
kim	ISEisC00L	SDA_Developers

Task 22 – Create ISE Authorization Result

Create an authorization result that returns an access-accept with VLAN 100.

Task 23 – Create ISE Policy Sets

Disable the default ISE policy sets. Create a new policy set that uses local authentication. The policy set should match on the condition of each identity group from above and return the appropriate SGT in the authorization response to the edge node.

Task 24 – Configure Supplicants

This task is not a Cisco task, but rather a nuance of using the Ubuntu node definition in CML. You need to install the wpa_supplicant software as it does not come with Ubuntu. The tricky part is internet is required for this, or you could download it from your computer and transfer it. Below is using the direct internet method. I made a video about this here, but below are the basic steps:

1. Connect the bridge connected to the “OOB-iNET” switch to a network with internet and DHCP
2. Login to the workstations (1-4) with cisco/cisco
3. Run “sudo apt update”
4. Run “sudo apt install wpasupplicant”

Next create a file at `/etc/wpa_supplicant/wpa_supplicant.conf` that includes:

```
country=US
ctrl_interface=/var/run/wpa_supplicant
update_config=1
ap_scan=0
network={
        key_mgmt=IEEE8021X
        eap=PEAP
        identity="bob"
        password="ISEisC00L"
        eapol_flags=0
}
```


This will need to be done on each workstation with the following users that we created in ISE:

Username	Workstation	Group
bob	Workstation 1	SDA_Employees
joe	Workstation 2	SDA_Contractors
kim	Workstation 3	SDA_Developers
susan	Workstation 4	SDA_Contractors

Then you need to shut down the internet facing interface, and bring up the interface facing the fabric. Be sure to replace X with the workstation number:

```
1. sudo ip link set ens3 up
2. sudo ip address add 10.254.100.10X/24 dev ens3
3. sudo ip route add 0.0.0.0/0 via 10.254.100.1
4. sudo ip link set ens2 down
```

Next you need to start the wpa_supplicant process and watch the authentication occur:

```
1. sudo wpa_supplicant -c /etc/wpa_supplicant/wpa_supplicant.conf -D wired -i ens3
```

Task 25 – Verify Fabric Operations and TrustSec Policy

Verify via pings that the intended security policy is enacted (ex: Contractors should not be able to ping Employees but Developers should):

SRC DST	**Employees**	**Contractors**	**Developers**
Employees			
Contractors			
Developers			

Also, verify that you can ping destinations external to the fabric from the workstations such as ISE and DNA. This checks functionality of the IP handoff.

Conclusion – Bonus Tasks

Congratulations! If you made it this far you have a functional fabric! The answer key will stop here, but here are some additional bonus tasks that are achievable using the current topology:

- Create a second VN named Guest that can only access one IP past the fusion router (simulating the internet)
- Imagine you have a host not capable of dot1x, use host onboarding to configure a switchport with a manual SGT and open authentication. Do not use the CLI!
- Redo the BGP peering for the IP handoff the extend the macro segmentation to the Fusion router and use route leaking to selectively inject one route from the GRT into the fabric.
- Configure the fabric to prefer border one, and ensure return traffic does the same.

I am always open to suggestions about how to make my work better. Feel free to contact me via LinkedIn or via email mreimert[at]mertandhouse[dot]com.

Answer Key / Explanations

Task 1

Configuration

On DC-SW-1:

```
1. router bgp 650001
2. network 10.1.10.0 mask 255.255.255.0
3.    neighbor 10.1.1.2 remote-as 650002
```

On FUSION-1:

```
1. router bgp 650002
2.    neighbor 10.1.1.1 remote-as 650001
```

Verification / Troubleshooting

- show ip bgp neighbors
- ping opposite neighbor addresses

Task 2

Configuration

On FUSION-1:

```
1. router bgp 650002
2.    neighbor 10.0.1.2 remote-as 650003
3.    neighbor 10.0.2.2 remote-as 650003
```

On HQ-BCP-1:

```
1. router bgp 650003
2.    neighbor 10.0.1.1 remote-as 650002
```

On HQ-BCP-2:

```
1. router bgp 650003
2.    neighbor 10.0.2.1 remote-as 650002
```

Verification / Troubleshooting

- show ip bgp neighbors
- ping opposite neighbor addresses

Task 3

Configuration

On HQ-BCP-1:

```
1. router isis
2.    net 49.0001.0000.0000.0001.00
3. interface range g1/0/2-3,lo0
```

```
4.    ip router isis
```

On HQ-BCP-2:

```
1. router isis
2.    net 49.0001.0000.0000.0002.00
3. interface range g1/0/2-3,lo0
4.    ip router isis
```

On HQ-IN-1:

```
1. router isis
2.    net 49.0001.0000.0000.0101.00
3. interface range g0/0-3,lo0
4.    ip router isis
```

On HQ-IN-2:

```
1. router isis
2.    net 49.0001.0000.0000.0102.00
3. interface range g0/0-3,lo0
4.    ip router isis
```

On HQ-EDGE-1:

```
1. router isis
2.    net 49.0001.0000.0000.0011.00
3. interface range g1/0/1-2,lo0
4.    ip router isis
```

On HQ-EDGE-1:

```
1. router isis
2.    net 49.0001.0000.0000.0012.00
3. interface range g1/0/1-2,lo0
4.    ip router isis
```

Verification / Troubleshooting

- show isis neighbors
- ping loopback addresses sourced from local loopback address

Task 4

Configuration

On HQ-BCP-1:

```
1. router bgp 650003
2.    redistribute isis level-1-2
3.    network 9.9.9.11 mask 255.255.255.255
4. router isis
5.    redistribute bgp 650003
```

On HQ-BCP-2:

```
1. router bgp 650003
2.    redistribute isis level-1-2
3.    network 9.9.9.12 mask 255.255.255.255
```

```
4. router isis
5.    redistribute bgp 650003
```

Verification / Troubleshooting

- show ip bgp on fusion <- verify 9.9.9.XX addresses are present
- show ip route isis on Edge/BCP nodes <- verify DNAC subnet is present (10.1.10.0/24)

Task 5

Configuration

On **all Fabric Nodes**:

```
 1. snmp-server community cisco rw
 2. aaa new-model
 3. aaa authorization exec default local
 4. username dna privilege 15 password ISEisC00L
 5. line vty 0 4
 6.  exec-timeout 0 0
 7.  privilege level 15
 8.  transport input ssh
 9. ip domain name ccie.local
10. netconf-yang
```

Verification / Troubleshooting

- Use SSH to hop from one router to another to ensure SSH is running properly
  - ssh -l dna 9.9.9.XX
- SSH to NETCONF port on router (830) to ensure router is accepting NETCONF sessions
- Ensure crypto keys exist for SSH

Task 6

Configuration

On **DNAC (Provision > Inventory > Add Device) , repeat for each fabric enabled device** (not intermediate nodes):

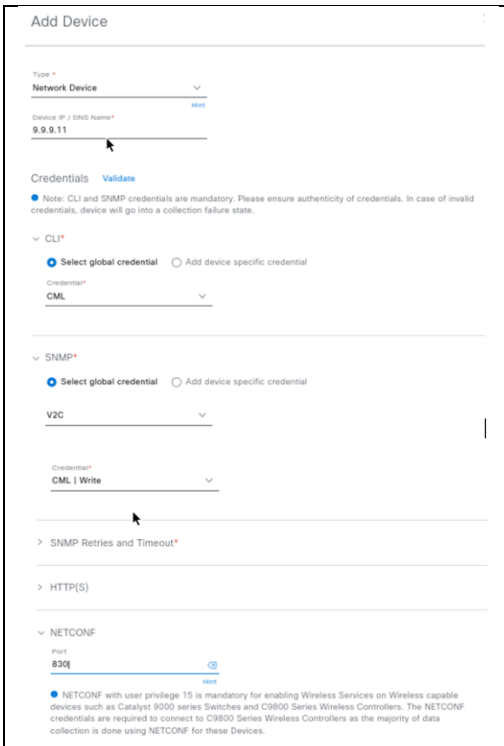

Verification / Troubleshooting

- Use SSH to hop from one router to another to ensure SSH is running properly
  - ssh -l dna 9.9.9.XX
- SSH to NETCONF port on router (830) to ensure router is accepting NETCONF sessions
- Ensure crypto keys exist for SSH
- Verify SNMP is running with the correct community (cisco)

Task 7

Configuration

On **DNAC (Design > Network Hierarchy > Add Area)** add USA:

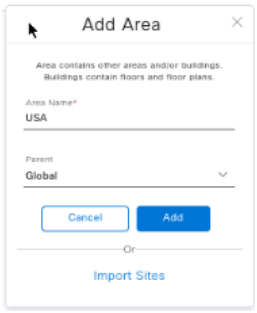

Under USA, add a building named SDA-Campus:

Verification / Troubleshooting

- Verify USA > SDA-Campus exists in the network hierarchy:

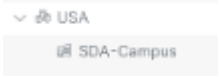

Task 8

Configuration

On **ISE (Administration > Deployment),** enable PXGrid:

On **DNAC (System > Settings > Authentication and Policy Servers)** add ISE):

Accept the certificate warning for ISE's self-signed HTTPS cert

Verification / Troubleshooting

- Verify reachability from DNAC to ISE by logging into DNAC command line as maglev user and pinging ISE
- Ensure PXGrid is enabled on ISE
- ISE web and console password must match for DNAC support

Task 9

Configuration

On **DNAC (Design > Network Settings > USA > SDA-Campus)** add Servers -> AAA):

Verification / Troubleshooting

- Ensure ISE is fully added to DNAC, verify in System 360 that the status is green

Task 10

Configuration

On **DNAC (Provision > Inventory)** select all devices, then Assign Device to Site:

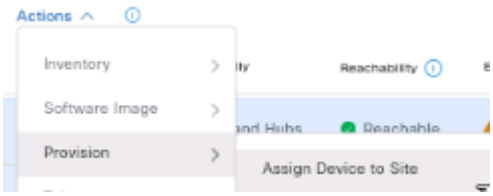

Assign all to the SDA-Campus site:

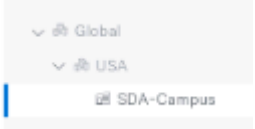

Proceed with suggested settings

Verification / Troubleshooting

- Ensure all devices are “Reachable” and “Managed”

Task 11

Configuration

On **DNAC (Design > Network Settings > IP Address Pools > Add)** :

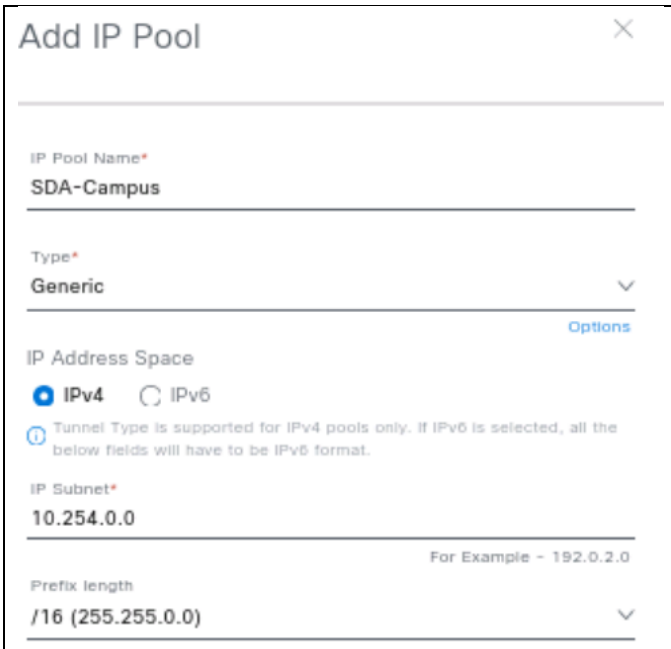

You do not need to define Gateway, DHCP, or DNS servers because those will be defined when the pool is reserved.

Then, at the building level:

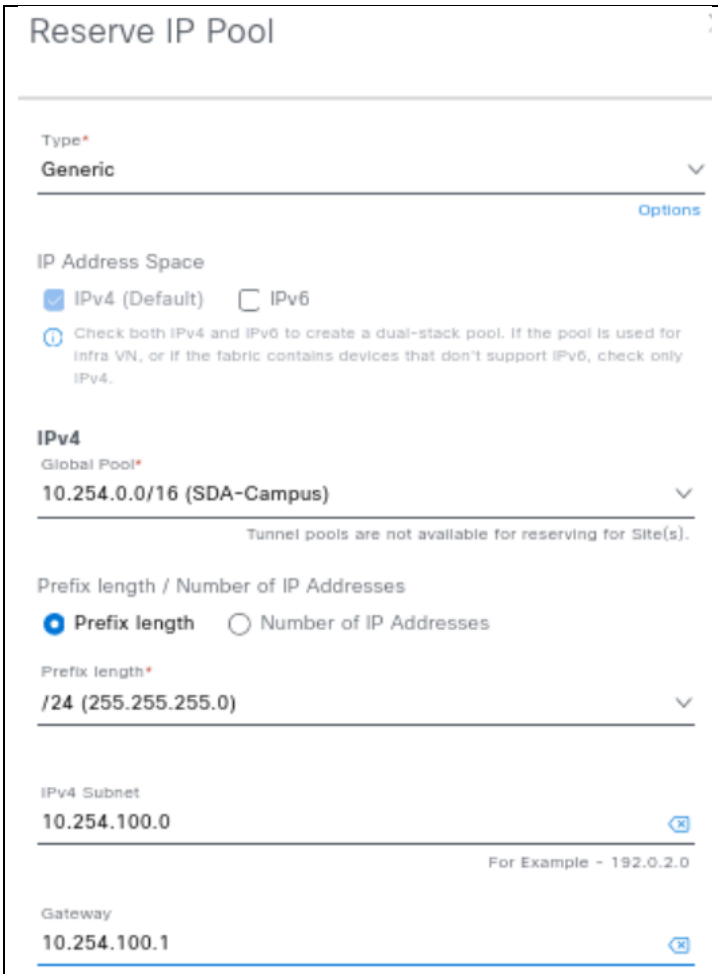

Repeat this for the L3 Handoff pools using a /30 subnet mask.

Verification / Troubleshooting

- Ensure you are selected under “Global” in the hierarchy

Task 12

Configuration

On **DNAC (Provision > Fabric Sites > Create Fabric Site)** using the following parameters :

1. SDA-Campus Building
2. Closed Authentication
3. Do not setup Fabric Zones

Verification / Troubleshooting

- Ensure AAA server is added under Network Settings

Task 13

Configuration

On **DNAC (Provision > Transits)** create a transit using the following parameters**:**

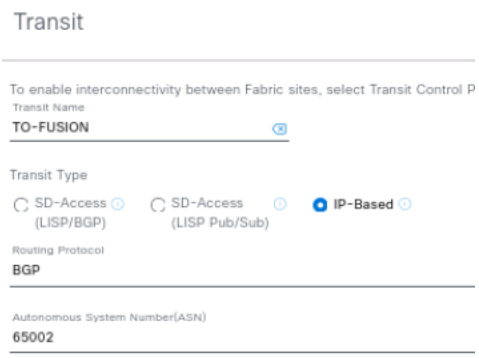

Verification / Troubleshooting

- Ensure the SD-Access Application is fully installed into DNA Center

Task 14

Configuration

On **DNAC (Provision > Virtual Networks)** create the following virtual network:

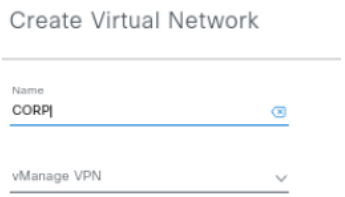

On **DNAC (Provision > Fabric Sites)** add the VN to your fabric site under “Host Onboarding”:

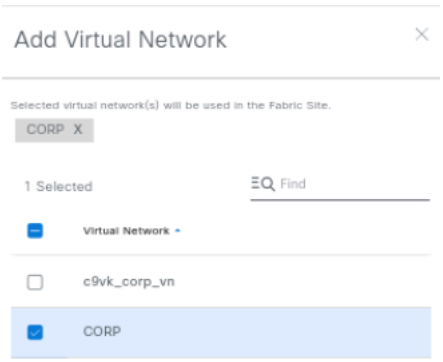

Verification / Troubleshooting

- Ensure the SD-Access Application is fully installed into DNA Center

Task 15

Configuration

On **DNAC (Provision > Fabric Sites)** enter the SDA-Campus site and make the following assignments:

1. Add the Edge Nodes (HQ-EDGE-1/HQ-EDGE-2) as Edge Nodes

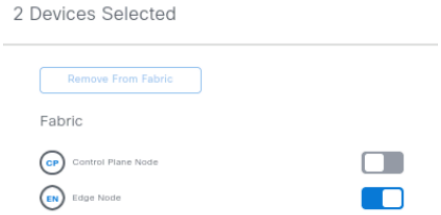

2. Add the Border and Control Plane nodes (HQ-BCP-1/HQ-BCP-2) as both Border and Control Plane Nodes. The border role is not assignable when using multiselect, you must select one node at a time to assign the border role. **This configuration must be completed on both border nodes!**
    a. Border node configuration

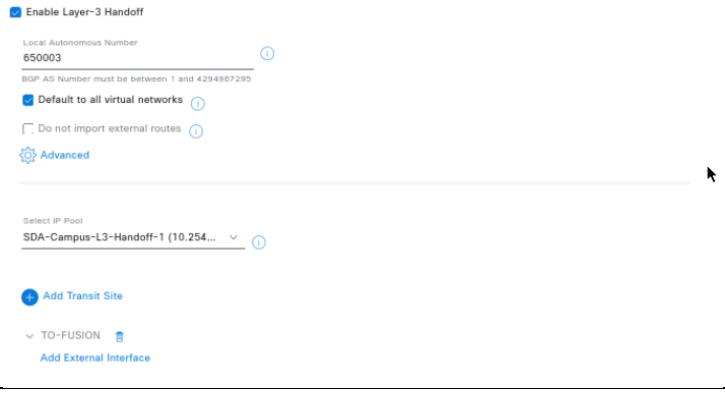

b. Transit site configuration

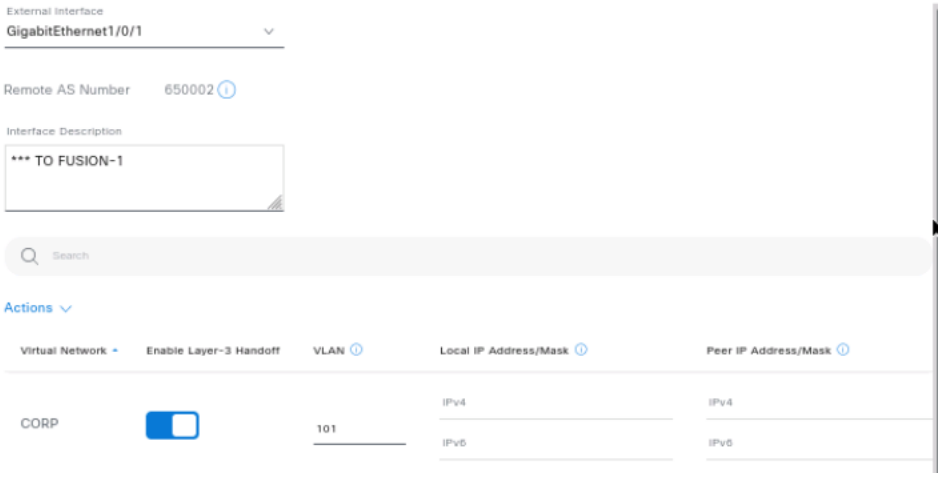

3. When assignments are complete your roles should mirror the following:

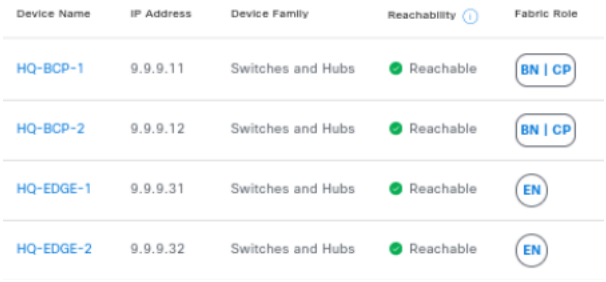

Verification / Troubleshooting

- Ensure you have the CORP VN created
- Ensure you have the IP Transit created
- All devices must be already provisioned to site (see task 10)
- IP Routing and Loopback0 interfaces must exist on each device that holds a fabric role
- This step has the most dependencies out of the entire lab, if a button is grayed out, or a task is failing most likely a you have skipped a previous step
- The links on the Border Nodes facing the fusion router must be switchports. They cannot be routed ports because DNA will not understand how to add VLANs for the L3 handoff. This means your BGP session will be sourced from an SVI in the native VLAN on that trunk. This is preconfigured if you started from the lab YAML.
- If you are getting an error that IP routing is not enabled and it is, I have some bad news for you. The image of c9kv you used is not correct and is not supported in DNAC. Even though the device will show as supported it will not be able to be added to a fabric.

Task 16

Configuration

On the fusion router add the following subinterfaces:


```
1. interface GigabitEthernet2.101
2.  encapsulation dot1Q 101
3.  ip address 10.254.251.2 255.255.255.252
4. interface GigabitEthernet3.102
5.  encapsulation dot1Q 102
6.  ip address 10.254.252.2 255.255.255.252
```

Then, add the following BGP peers under the existing BGP process:

```
1. router bgp 650002
2.    neighbor 10.254.251.1 remote-as 650003
3.    neighbor 10.254.252.1 remote-as 650003
```

Verification / Troubleshooting

- "show ip bgp summary" on the fusion router
- "show bgp vpnv4 unicast all summary" on each border
- "show ip route vrf CORP" on each border, ensure DNA subnet is present

Task 17

Configuration

On **DNAC (Provision > Fabric Sites > SDA-Campus > Host Onboarding > Virtual Networks)** click on the "CORP" VN and add the following gateway:

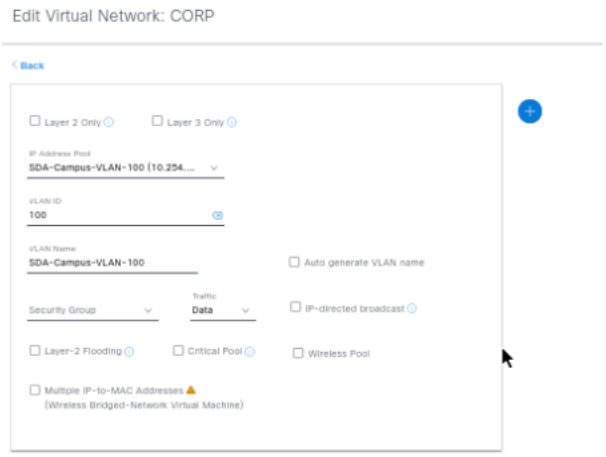

Verification / Troubleshooting

- If the VN is missing ensure the VN is added to the fabric site
- To check and make sure that the anycast gateways are deployed, run the following commands on your edge nodes: "show run interface vlan 100", the VN should also be

blue now that is is deployed and contains IP Pools.

```
HQ-EDGE-1(config)#do show run int vlan100 | beg interface
interface Vlan100
 description Configured from Cisco DNA-Center
 mac-address 0000.0c9f.f55d
 vrf forwarding CORP
 ip address 10.254.100.1 255.255.255.0
 no ip redirects
 ip route-cache same-interface
 no lisp mobility liveness test
 lisp mobility SDA-Campus-VLAN-100-IPV4
end
```

Task 18

Configuration

On **DNAC (Policy > Group Based Access Control)** start the policy matrix migration:

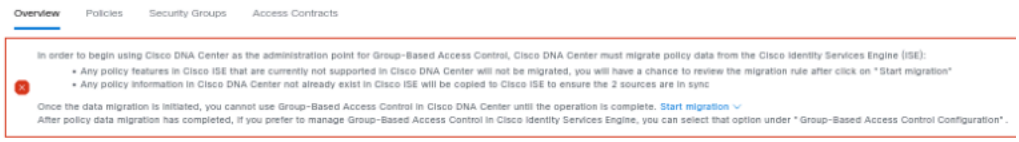

Verification / Troubleshooting

- Wait until the migration is complete until proceeding with changes

- If the migration fails, visit System360 in DNAC and ensure the connections to ISE are healthy

Task 19

Configuration

On **DNAC (Policy > Group Based Access Control > Policy)** define the policy desired in the workbook:

Then you must deploy the policy to ISE using the deploy dropdown.

Verification / Troubleshooting

- Ensure you pressed the deploy button and received a success, you can also view a read only matrix in the ISE UI to ensure changes propagated

Task 20

Configuration

On **ISE (Administration > Identity Management > Groups > User Identity Groups)** define the groups from the task.

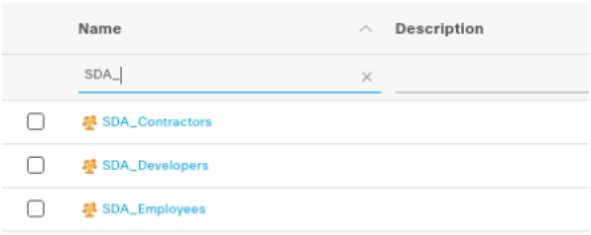

Task 21

Configuration

On **ISE (Administration > Identity Management > Identities)** define the users from the task.

Task 22

Configuration

On **ISE (Policy > Results) create an Authorization Result that returns VLAN 100:**

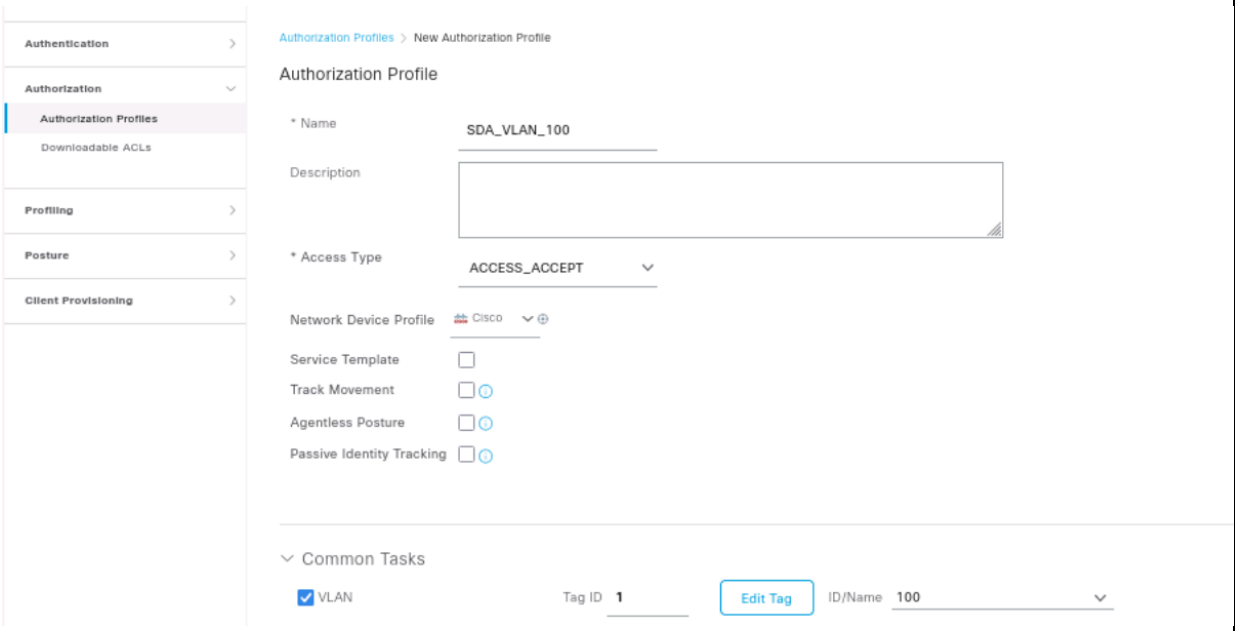

Task 23

Configuration

On **ISE (Policy > Policy Sets)** add a policy set above the default policy set:

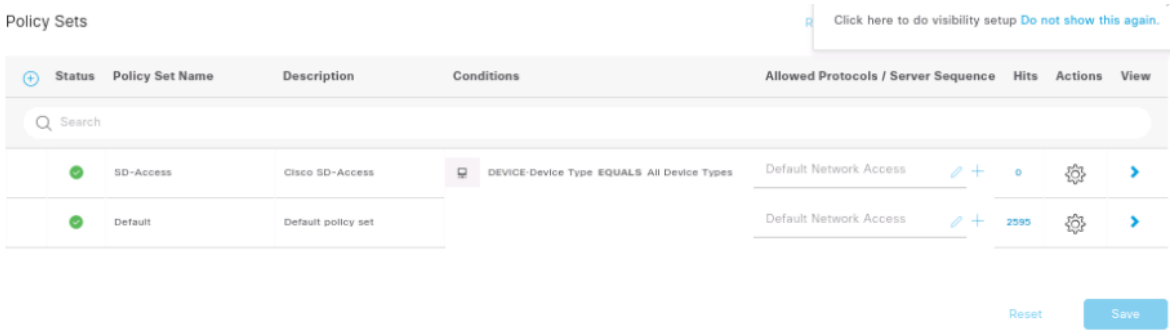

Most of the policy can be left default, but authorization policies need to be defined for each SGT we want to return to the switch based on what group the users are in. Create an entry for each group, and match on that group returning the SGT and the result that assigns VLAN 100:

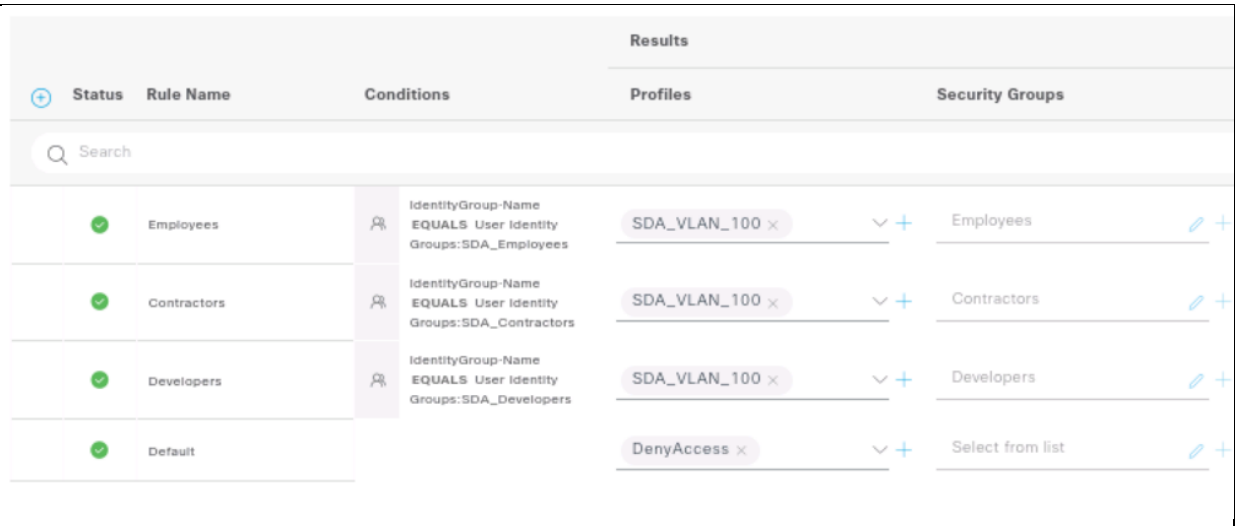

Verification / Troubleshooting

- Ensure the groups exist from the tasks above, the SGTs are default SGTs so they will always exist

Task 24

Configuration

You should get the following log message on the terminal of each workstation, for the configuration refer to the lab task or watch the video linked in the lab task:

EAP-MSCHAPV2: Authentication succeeded
ens3: CTRL-EVENT-EAP-SUCCESS EAP authentication completed successfully

Task 25

Verification

The workstations should all be able to ping their default gateway after authentication:

```
PING 10.254.100.1 (10.254.100.1) 56(84) bytes of data.
64 bytes from 10.254.100.1: icmp_seq=1 ttl=254 time=89.9 ms
64 bytes from 10.254.100.1: icmp_seq=2 ttl=254 time=73.0 ms
64 bytes from 10.254.100.1: icmp_seq=3 ttl=254 time=85.6 ms
```

The workstations should be able to ping each other selectively based on the defined TrustSec policy.

Verification / Troubleshooting

- Use “show access-session interface g1/0/7|g1/0/8” to verify the host is authenticated

- If reachability issues are across edge nodes ensure reachability between loopbacks of edge nodes and to the CP node.

