

Cisco Software-Defined Access



Contents

Cisco SD-Access	3
Evolution of Networking Toward Fabric	3
Cisco SD-Access Overview	4
Migration Considerations.....	7
Network Considerations	8
Policy Considerations	9
Existing Network Design	10
Migration.....	11
Reference Network Design for Migration	11
Use new IP Subnets Optimized for SD-Access.....	12
Connecting the First Fabric Border/Control Plane and Fabric Edge Switch	12
Incremental Approach	13
Option 2: Map IP Subnets in SD-Access to Outside VLAN	25
Connecting the SD-Access Fabric to a Traditional Network.....	26
Service Interworking with SD-Access.....	29
Summary	30

Over the past few years, the digitization wave has gained more and more momentum throughout all industry sectors. Digital disruption now has the potential to overrun incumbents and redefine markets faster than perhaps any force in history.

Cisco SD-Access

Cisco® Digital Network Architecture (Cisco DNA™) is designed to meet the requirements of the rapid increase in digitization. The design philosophy behind the Cisco SD-Access architecture centers on policy-based automation, whereby network infrastructure is provisioned with secure user and device segmentation independent of media connectivity (wired or wireless users).

Automation and simplicity result in increased productivity, which enables IT staff to innovate early and be an industry leader in transforming to a digital enterprise, increasing operational effectiveness. A consistent segmentation framework that aligns with business policy regardless of transport medium (wired or wireless) is a key requirement for achieving effectiveness at the core. The following sections explain the tenets of Cisco SD-Access, including the terminology, components, and scale.

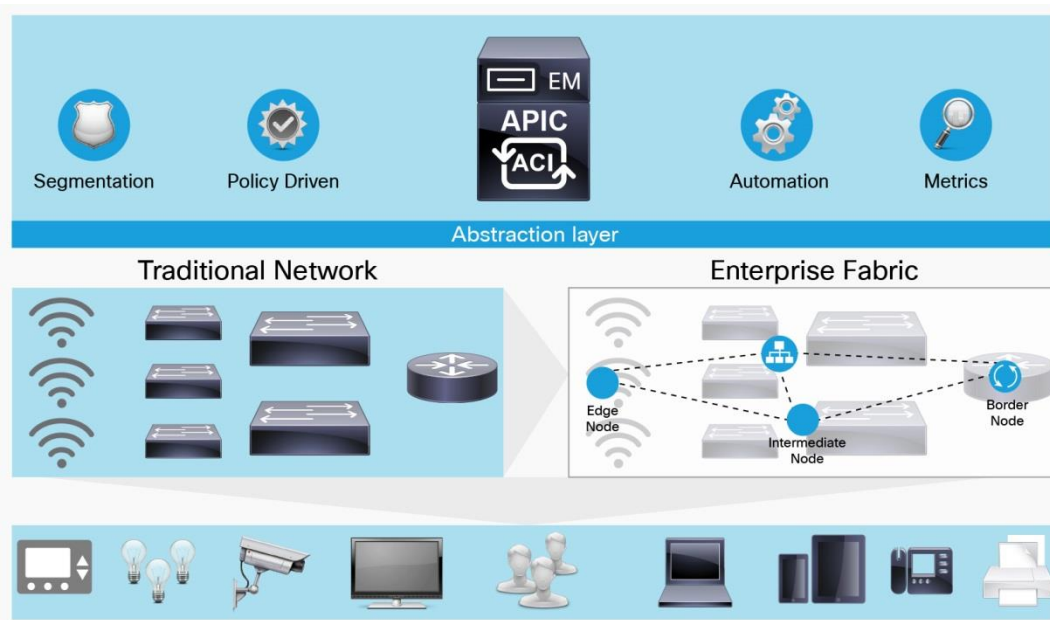
Cisco SD-Access offers the following technology benefits:

- **Simplified operations:** The ability to stretch IP subnets simplifies the overall network design. In addition, because all configuration is done on the fabric edge and border nodes, no hop-by-hop configuration is required.
- **Automation:** The deployment can be done through tools provided by the Cisco Application Policy Infrastructure Controller Enterprise Module (APIC-EM) or Cloud Managed Services Platform in the future, eliminating the need for a command-line interface (CLI) or CLI configuration templates.
- **Agility:** Network operations can now be done in a more agile manner to meet business requirements by reducing the number of manual configuration steps.
- **Security:** Embedded security and segmentation via neighborhoods (Virtual Routing and Forwarding (VRF)/virtual network) and user groups (Security Group Tag [SGT]/segment ID). It also provides macro-segmentation with neighborhoods, and micro-segmentation within a neighborhood via groups.
- **Consistent policies for wired and wireless:** Extends segmentation, visibility, and policy of the wired network to wireless. Distributed wireless termination helps scale the wireless network throughput while providing a centralized management and troubleshooting location.
- **Support for business analytics:** Provides analytics and other telemetry information aggregated into a single platform to make business decisions and enable planning for growth or diversification.

Evolution of Networking Toward Fabric

Cisco DNA builds an enterprise fabric that can be programmed to provide end-to-end connectivity between users, devices, branches, and data centers. The enterprise fabric is built on the traditional network using standards-based overlay technologies that are managed by the APIC-EM, as illustrated in Figure 1.

Figure 1. Cisco SD-Access Architecture



Today the IT group maintains a wide range of assets, including critical databases, vital company employee and customer information, classified commercial information, shared drives, and email and web servers, among many other elements. Any IT organization today will benefit from the following:

- **Identity-based segmentation and policy:** SD-Access decouples security policy definition from VLAN and IP address to enable rapid policy updates.
- **Automated network fabric:** SD-Access provides automation across wired and wireless, enabling IT to optimize resource utilization as well as traffic flows, in addition to being able to move away from device-by-device automation techniques to workflow-based simplified management. This results in the ability to maintain consistency at scale.
- **Insights and telemetry:** SD-Access leverages insights and analytics into user and application behavior for proactive issue identification and resolution.
- **Policy convergence** between wired and wireless.
- **Flexible authentication** options for users, devices, and things, including 802.1X, Active Directory, and static authentication.
- **Better positioning** for increased cloud usage via WAN and Internet; acceleration and optimization for cloud.

Cisco SD-Access Overview

The Cisco SD-Access fabric builds on a robust network with the design principles described in traditional networks. This fabric provides any-switch-to-any-switch connectivity via standards-based stateless tunnels. LISP and virtual extensible LAN (VXLAN) are the underlying technologies used to build the SD-Access fabric; however, these technologies are completely abstracted from the end user.

Note: As this fabric is built on top of a traditional network, it is normally referred to as an overlay network, and the traditional network is referred to as an underlay network.

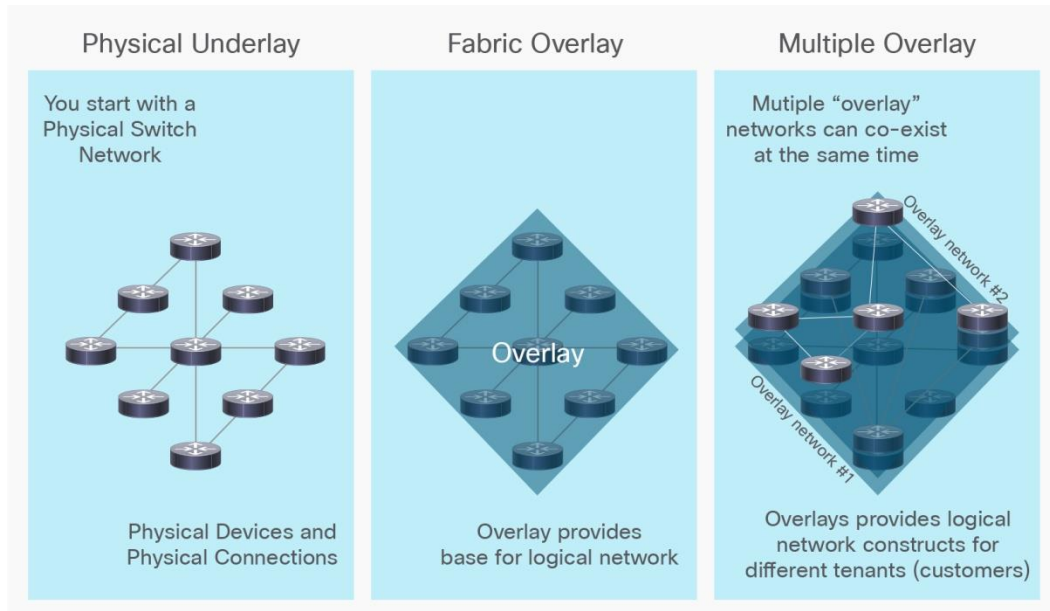
The networking approach used to build the Cisco SD-Access fabric consists of a prescriptive physical underlay and a programmable overlay with segmentation constructs such as virtual networks (VNs) and groups. This new approach enables enterprise networks to transition from the traditional VLAN-centric architecture to a new user group-centric architecture.

Table 1 defines the components used in the Cisco SD-Access fabric.

Table 1. Cisco SD-Access Terminology


Name	Description
Fabric edge node	Edge nodes are responsible for authenticating (static, 802.1X, Active Directory) and registering endpoints (laptops, phones, wireless devices, etc.) with the control plane node.
Fabric intermediate node	These devices are part of Layer 3 that interconnects edge nodes and border nodes. These nodes are labeled in red in Figure 3. They are responsible for routing the traffic within the fabric.
Fabric border node	These nodes are gateways between the fabric domain and networks external to the fabric. They connect traditional Layer 3 networks or different fabric domains to the SD-Access fabric and also are responsible for translating context information from the SD-Access fabric domain to another domain if it exists.
Control plane node	A database that tracks all endpoints in the fabric domain and maps them to respective fabric nodes. The control plane node enables the decoupling of the user or device from its identity (IP address).
Fabric domain	A collection of network entities consisting of fabric edge nodes, intermediate nodes, and fabric border nodes along with its own control plane node. It can be limited to a geographical location, or one can have multiple fabric domains within a location, depending upon scale and performance.
Virtual network	A high-level segmentation concept to separate different users or devices connecting to SD-Access and mapping to a VRF instance in traditional networks.
Application Policy Infrastructure Controller Enterprise Module (APIC-EM)	The controller that network operators can use to orchestrate the enterprise network space.
Cisco DNA Center	An application front-end user interface that provides access to all the applications that help automate/orchestrate the network.
Cisco Identity Services Engine (ISE)	The platform where policy and segmentation (SGT) definitions are programmed using Cisco DNA Center.
Physical underlay	The access switches, distribution switches, core routers, and WAN and data center interconnects. The end-to-end connectivity between the network elements in the underlay is provided by standard means, including routing protocols. Refer to Figure 2.
Fabric overlay	The fabric forms the overlay network on top of the underlay network. In simple terms, it provides any-to-any switch connectivity via stateless tunnels. The addresses of endpoints are carried in the overlay network. The fabric provides unicast and multicast communications between endpoints and supports anycast gateway services for mobility.
Endpoint addresses	Addresses assigned to the users and devices that connect to the downlink Ethernet ports of the access switches.

Figure 2. Logical Topologies Over the Fabric




Cisco SD-Access comprises the roles shown in Figure 3. They can be colocated on the same device or placed on different devices, depending upon the scale requirements.

Figure 3. Roles of Devices in Cisco SD-Access Fabric



The fabric edge node provides Layer 3 gateway services for IP subnets. The gateway IPv4/IPv6 addresses and gateway MAC are the same across all fabric edge nodes (anycast gateway). This enables the Cisco SD-Access fabric to extend a single subnet across multiple sites. These are shown by the FE icon in Figure 3. SD-Access can also support topologies in which a traditional network with Spanning Tree can connect to the SD-Access fabric, as shown in Figure 4.

Figure 4. SD-Access Fabric Connecting to a Traditional Spanning Tree Network



The control plane (CP) node enables the decoupling of the user or device from its identity (IP address). The SD-Access fabric supports dual active-active CP nodes. The fabric edge nodes will register endpoint entries to both CP nodes. There are two deployment models for the CP node in the fabric:

- The CP node is co-located on a fabric border node.
- The CP node is separate from the fabric border node.

The following are recommendations for selecting the platform to host the CP node when co-located with the fabric border node:

- For small fabric domains (up to 5000 endpoint entries), the platform of choice can be a network element such as the Cisco Catalyst® 3800 Series Switches.
- For medium-sized fabric domains (5000 to 20,000 endpoint entries), the platform of choice can be a network element such as the Cisco Catalyst 6800 Series or 9000 family of switches.
- For large fabric domains (20,000 to 40,000 endpoint entries), the platform of choice can be a network element such as the Cisco ASR 1000-X or 1000-HX router or the 4400 Series Integrated Services Routers (ISRs).

The deployment model in which the CP node is separate from the fabric border node is recommended for large-scale enterprise designs with more than 80,000 endpoints. The platform of choice for the CP node is a Cisco Cloud Services Router (CSR) 1000v.

Migration Considerations

The following are considerations to take into account before beginning the migration of the existing network to Cisco SD-Access. They are categorized as follows:

- Network considerations: Maximum transmission unit (MTU), network topology, IP addressing for underlay and overlay, and location of shared services.
- Policy considerations: Existing policy definition and enforcement points, virtual network, and Security Group Tags (SGTs).
- Hardware platform considerations: Switches, routers, wireless controllers, and access points that support SD-Access.
- Software platform considerations: Cisco DNA Center, ISE, Network Data Platform (NDP).

- Scale of deployment considerations: Scale of hardware platforms with respect to the role they play in the SD-Access architecture.
- Existing network design: Layer 2 access or routed access.

Network Considerations

Maximum Transmission Unit (MTU)

MTU is defined as the largest network protocol data unit that can be transmitted in a single transaction. The higher the MTU, the more efficient the network. The VXLAN encapsulation adds 50 bytes to the original packet. This can cause the MTU to go above 1500 bytes for certain applications. For example, wireless is deployed with SD-Access, where the additional Control and Provisioning of Wireless Access Points (CAPWAP) overhead needs to be considered. In general, increasing the MTU to 9100 bytes on interfaces across all switches and routers in the fabric domain (underlay and overlay) is recommended to cover most cases and to prevent fragmentation.

Network Topology

SD-Access fabric supports traditional hierarchical networks as well as arbitrarily designed networks such as ring topology or daisy-chained topologies. **Note:** A network designed with the Cisco Validated Design guidelines will have fewer considerations (steps) when migrating compared to arbitrarily designed networks that are inherently complex. Since fabric underlay topologies are based on a routed access design, if the existing network is routed access, it lends itself to easier migration to SD-Access.

IP Addressing for Underlay and Overlay

Existing campus networks are flat and do not have any concept of underlay and overlay. The IP address schema is flat, with no distinction between intranetwork prefixes and endpoint network prefixes. SD-Access, by its very nature, contains overlay and underlay to differentiate between the two spaces. It is recommended that two distinct IP ranges be selected, one for the endpoint network prefixes (overlay) and one for the intranetwork prefixes (underlay). The advantages are twofold. First, it enables summarization of the IP space when advertising in routing considerations. Second, troubleshooting is easier, since one has a clear understanding of which IP space one is looking at. For example, the overlay could be a 10.0.0.0/8 space, and the underlay range could be a 172.16.0.0/16 space.

Location of Shared Services

Shared services in the network include services such as Dynamic Host Configuration Protocol (DHCP), DNS, IP address management, Network Time Protocol (NTP), NetFlow collector, syslog, network management systems (NMS), and others. Traditionally, these services lie outside the campus or branch network in a data center. Some network designs do have some or all of these services in the campus or branch, connected to either a core or a distribution layer. Additionally, the shared services are normally in the global routing table (GRT), although in some deployments they might lie in a separate VRF context. It is essential that network devices and endpoints have access to basic services such as DHCP and DNS in order to connect to the network and forward traffic. The steps for migrating to SD-Access differ depending upon the physical location as well as the presence in either GRT or VRF of the shared services in the existing network.

Application of Features at the Distribution Layer

In a Layer 2 access design, in most cases features such as IP access control lists (ACLs), NetFlow, quality-of-service (QoS) classification, and marking and policing are configured at the distribution layer switches. Since SD-Access is a fabric solution, the incoming packets from the endpoints are encapsulated in the fabric data plane by the fabric edge, making the distribution layer switches act as intermediate nodes that switch IP packets back and forth between fabric edge (access layer) and upstream switches in the network. Due to the encapsulation at the fabric edge itself, the IP classification that the features were based on at the distribution layer is not available; hence the consideration of moving these features to the access layer switches in the network.

Routing between VRF and Underlay to External Network

The routing locator (RLOC) addresses (typically Loopback0) and underlay physical connectivity address space are in the GRT. The endpoint IP space will typically be in VRFs if not the default VRF. The network devices will still be reachable by the infrastructure and network management stations via the RLOC space in the GRT.

Policy Considerations

A mind shift is needed when SGT enforcement is considered, because the enforcement is based not on static IP ACLs but rather on dynamic downloaded security group (SG) ACLs, which are more secure. Implementation of 802.1X further strengthens the onboarding of endpoints onto the network, since network connections are now authenticated and/or profiled and placed in the right area in the network. How the users and things on the network should be isolated from each other is another consideration that the network administrator should work on with the security administrator of the network. SD-Access provides dual levels of segmentation within the network. With the deployment of VRF or VNs providing the classic path isolation among endpoints and SG ACL enforcement providing differentiated access control within the VN, it is imperative that network and security administrators work together to form a segmentation and access control policy that will be applied consistently in the network.

Hardware Platform Considerations

Table 2. Cisco SD-Access Hardware Requirements

Name	Description
Switching	Cisco Catalyst 3850 and 3650 Series Switches Cisco Catalyst 4500E Series Switches with Sup8-E supervisors Cisco Catalyst 9300, 9400, and 9500 Series Switches Cisco Catalyst 6500 and 6800 Series Switches with Sup2T and Sup6T with 6800 series line cards Cisco Catalyst 6840-X and 6880-X Switches Cisco Nexus 7700 Switches (with M3 series line cards)
Routing	Cisco 4400 Series ISRs Cisco 4300 Series ISRs ASR 1000-X and 1000-HX Routers
Wireless	Cisco 3504, 5520, and 8540 Wireless Controllers 802.11ac Wave 2 access points and later (Wave 1 access points with restrictions)

Software Platform Considerations

Based upon deployment size, scalability, and redundancy, the software functions can be run on top of individual virtual machines (VMs) or dedicated appliances for Cisco DNA Center and ISE.

Table 3 lists the current software requirements to implement the Cisco SD-Access architecture.

Table 3. Cisco SD-Access Software Requirements

Name	Description
Software components	Cisco DNA Center 1.0 ISE 2.3

Cisco SD-Access Scaling Considerations

Cisco SD-Access fabric scaling depends on the number of hosts and devices in a single site or across multiple sites. The Cisco DNA Center UI will support 500 network devices as fabric nodes (including fabric edge, fabric border, and fabric control plane nodes and wireless LAN controllers; excluding access points) and 20,000 endpoints per fabric domain. A total of 10 fabric domains is supported with Cisco DNA Center 1.0. Geographical locations that are in close proximity from a latency and performance standpoint can be controlled by a single APIC-EM instance. It is recommended that the APIC-EM be colocated near other software control plane entities such as Cisco Prime® Infrastructure, NDP, and ISE administrative nodes to reduce latency for communications between them. That way there is one variable, which is latency due to the WAN infrastructure (links and speeds), and not a combination of both. The locations might be in the data center or the main campus site, depending upon customer implementation.

It is recommended to run two control plane nodes per fabric domain for redundancy. For a given fabric domain, the choice of platform will depend upon the number of host entries to be managed by the control plane node. Hosting a single control plane node instance on a switch platform with active and standby supervisors or stack members provides an additional level of redundancy within a system. Hosting the other instance on another switch provides an additional layer of redundancy across systems. In the latter case, both control planes are active-active and all registrations are sent to both control plane nodes independently. There is no synchronization of the database across two control plane nodes.

Existing Network Design


There are generally two types of networks – branch and campus. Depending on their size, they can be two or more tiers. Sometimes the topology is constrained by physical factors such as space, which could lead to a daisy-chained or ring topology. A network topology designed using guidelines from the Cisco Validated Design is recommended. Whatever the topology of the physical design, it can be migrated to SD-Access – it is just that additional time and consideration should be given for deviations from the Cisco Validated Design. There are two ways of approaching migration as well:

- A parallel install involves erecting a new network parallel to the existing network. For a parallel install, prerequisites such as space, power, cooling, and cabling must be available in order to satisfy parallel live connections. This approach makes the migration easier, with the ability to revert to the existing network should the need arise, but it is also the most expensive in the sense that there are two networks running live, consuming power and space over a period of time while the migration is successfully completed.
- Migrating one switch at a time allows incremental migration of certain areas of the network rather than migration en masse. It is not dependent on factors such as space, power, and cabling, compared to a parallel install, and is more reasonable from a cost perspective. Network administrators have the option of experiencing how the new network will provide services and connections using a lot fewer network devices than in a parallel install. This approach also provides an opportunity to build a small SD-Access deployment over the existing network and test user scenarios in that small deployment.

Migration

Reference Network Design for Migration

Figure 5. Reference Network Topology



Consider the campus network shown in Figure 5. It is a multitiered network in which endpoints connect into the network at the access layer. This layer is aggregated into aggregation layer switches, typically at the main distribution frame (MDF). Typically, depending upon density, a building can be thought of as an aggregation block. The access network is a Layer 2 access design, and VLANs are spanned across the access layer, within an aggregation layer. Considering the above network, there are six buildings that connect back to the core – three buildings connecting to a single core block. The core block is typically a pair of switches for redundancy. In this example, the core blocks connect to a super-core block. In most networks, though, the aggregation layer connects to a single core block – the super core is collapsed into the core block. At the top in this example are the various blocks of the network that connect into infrastructure elements, namely network management stations, user repositories (Active Directory, ISE, or equivalent authentication, authorization, and accounting [AAA] RADIUS servers), DHCP servers, DNS servers, NTP servers, and NetFlow collectors, among others. This is called the shared services block. There is a WAN block that connects into the branch WAN network. A separate data center block connects to the data center over the WAN. Finally, at top right is the Internet block, which connects to the Internet. Typically, these days the services block, WAN block, and Internet block are absorbed into the data center block. Traffic from all the remote locations (campus and branch locations) is routed through the data center – since most of the traffic patterns are north-south (any location to the data center). This is a popular model, since it allows you to centralize all the services that are to be applied to the traffic going to the data center as well as the Internet. However, in this example, consider these blocks to be situated at the campus head-end network.

Use new IP Subnets Optimized for SD-Access

Consider the network shown in Figure 5 as an example for starting migration to Cisco SD-Access.

Choosing a new subnet to begin migration relieves the challenge in migrating existing networks – they can function as is, in the current configuration – but allows the network operations team to build out the SD-Access solution over the top, without disrupting any transactions in the existing network. In short, a fabric overlay is created by inserting a fabric edge and fabric border switch and creating the fabric over the existing network, using it as an underlay. Corresponding activities, such as provisioning a new DHCP scope, updating any firewall rules, and others, have to be done in advance. The new subnet approach also has the advantage of enabling the design of new subnets optimized for SD-Access, taking into account existing and future needs, compared to carrying on the legacy of networks designed in the past.

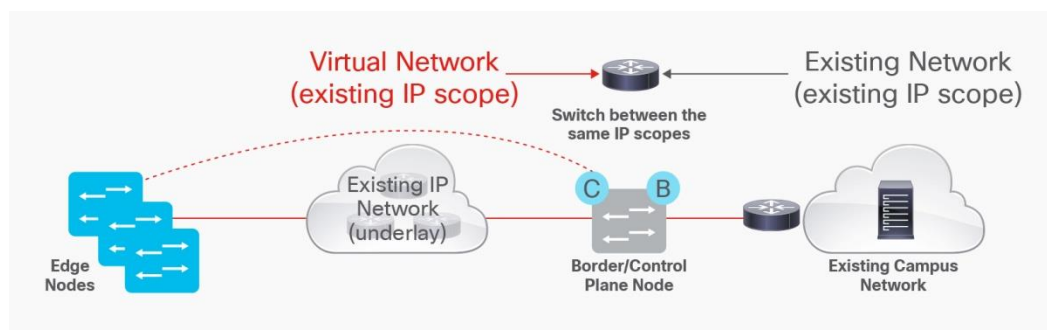
Connecting the First Fabric Border/Control Plane and Fabric Edge Switch

Connect the first fabric edge switch to the underlay. Configure the system MTU to be 9100 bytes on this switch. Ensure that the uplinks are configured as routed links up to the aggregation layer. Intermediate System to Intermediate System (IS-IS) is the recommended choice of Interior Gateway Protocol (IGP), but any existing IGP such as Open Shortest Path First (OSPF) or Enhanced Interior GRP (EIGRP) will also suffice. Cisco DNA Center's Plug and Play app can be used to automate the integration of this new access switch into the existing network.

The first fabric border switch that will also run a control plane for the SD-Access fabric will be inserted in the core layer, or the super-core layer in this case. There are two options to insert the fabric border.

- If the core switches support the fabric border functionality, one of the pairs can be reconfigured for the fabric border function.
- If the core switches do not support fabric border functionality, a new switch/router that will support the function as well as the required scalability can be connected off of the existing border pair. This will not affect the existing core layer in the network.

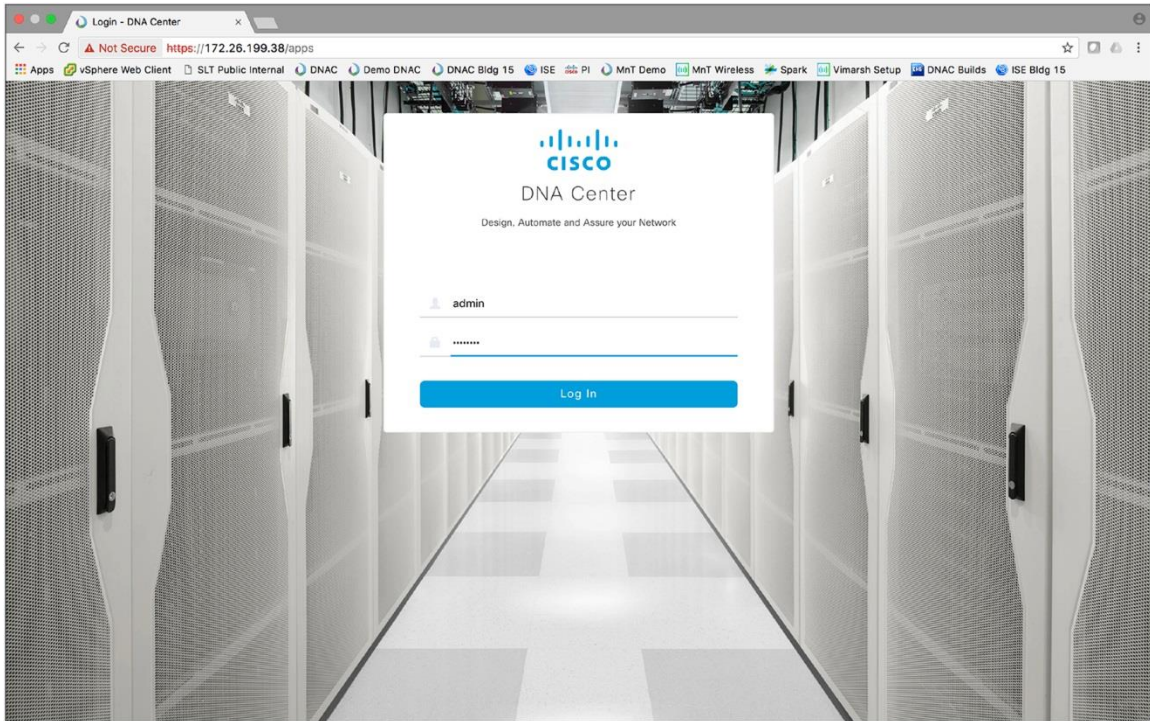
Using new IP subnets in the Cisco SD-Access fabric



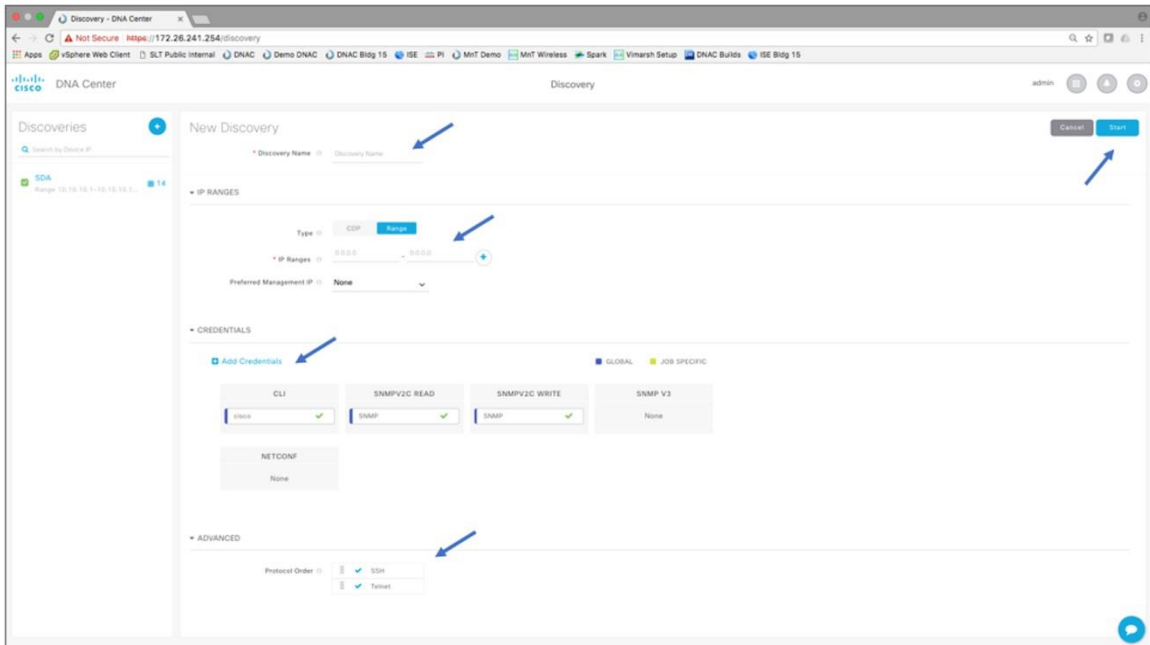
The idea is to build the Cisco SD-Access overlay network over the top of the existing network, which forms the underlay. The overlay network consists of multiple VNs that use new IP subnet space that is optimized for the SD-Access solution, which simplifies the IP design in the network as well. Deploy a border/control plane node that routes between the fabric network and the existing external networks.

Incremental Approach

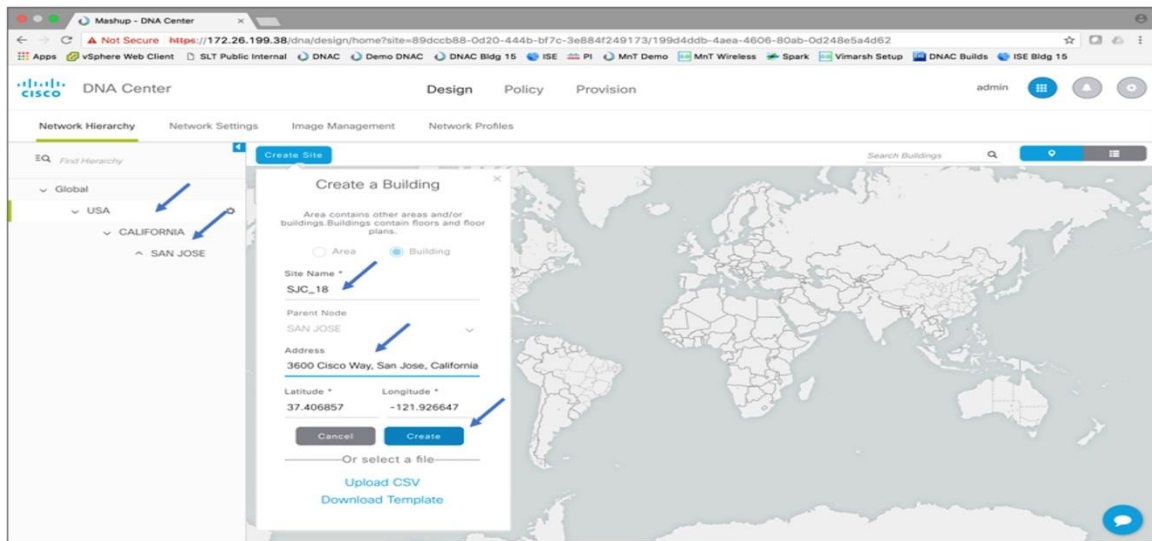
Step 1. Install and log in to Cisco DNA Center, as shown below.



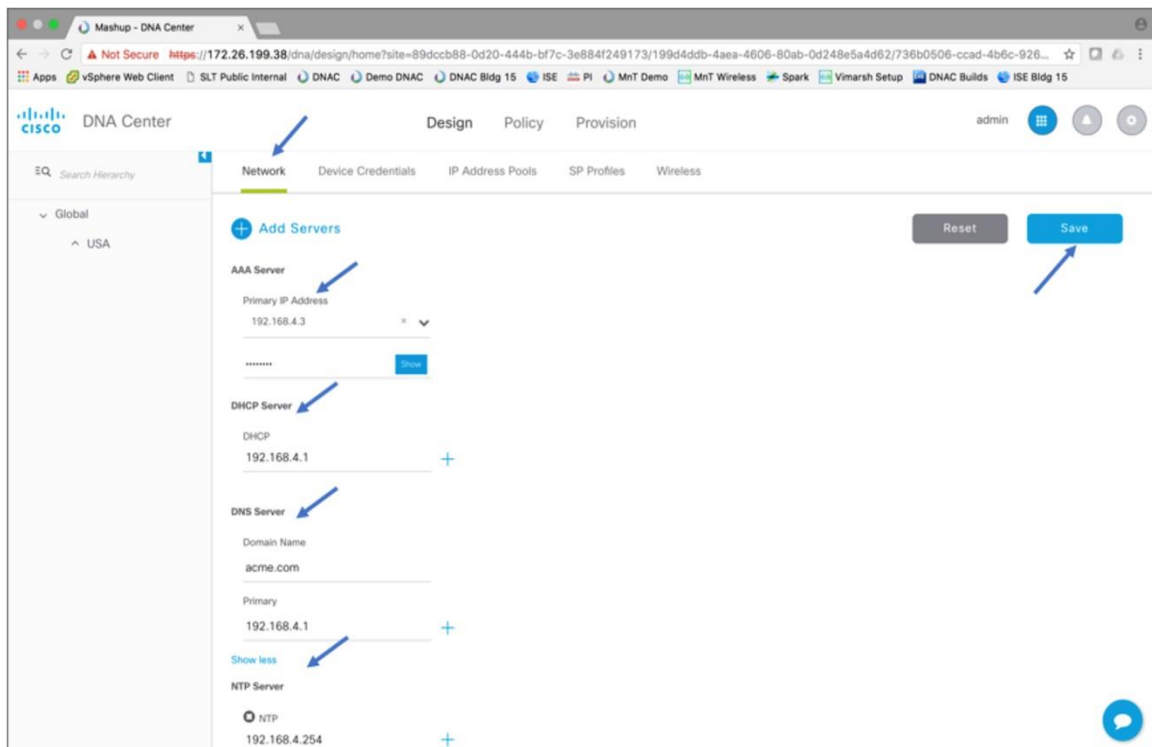
Step 2. Discover the existing network using network discovery within Cisco DNA Center, as shown below.

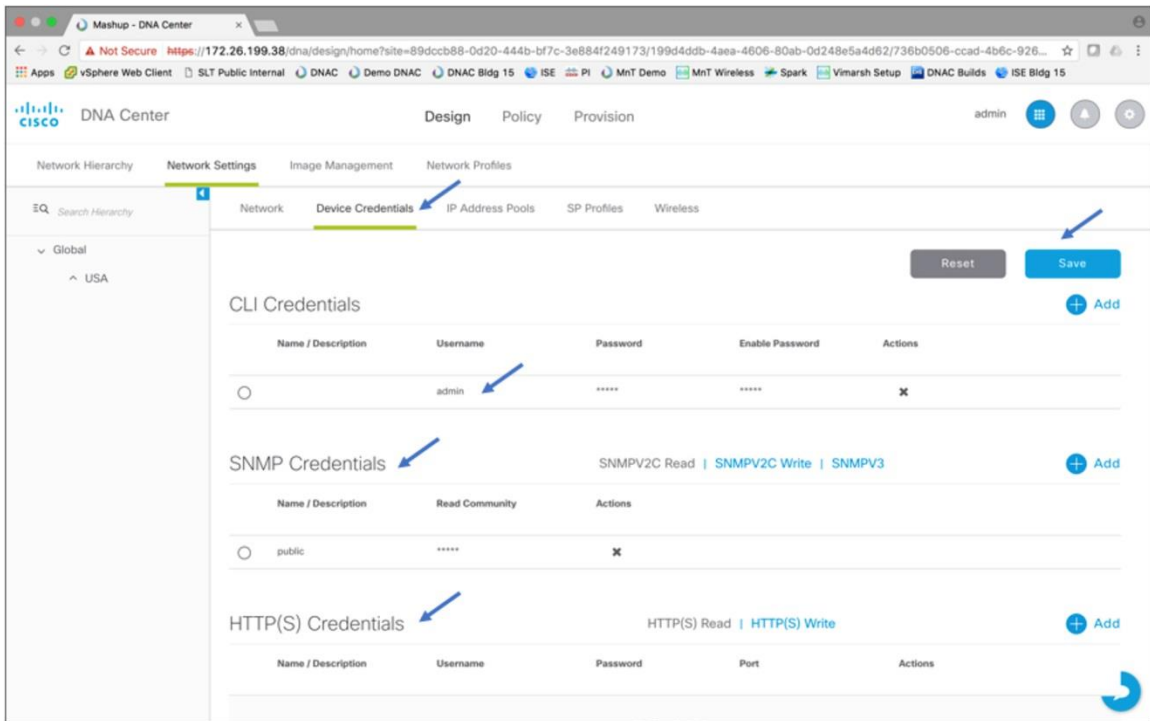


Step 3. Start with the Design tab, organizing network locations into a hierarchy. Create multiple sites as needed to be associated with a fabric domain, as shown with SJC_18 below.

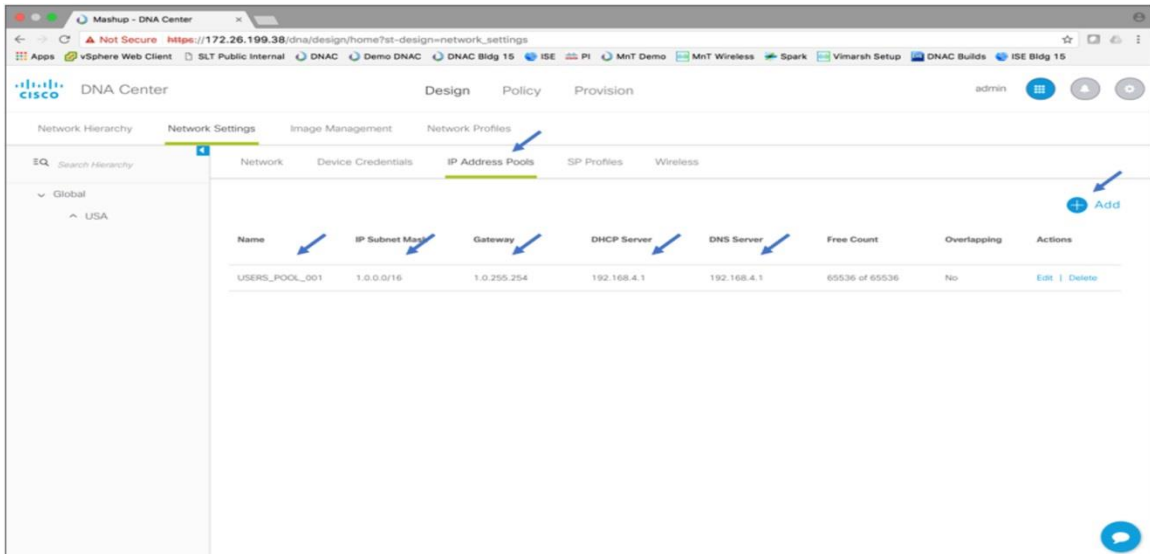


Step 4. Configure basic network settings, as shown below.





Step 5. Add a new IP subnet to the SD-Access fabric network, as shown below. This will be a new subnet prefix, and a /16 or /20 can be used to optimize the subnet size. Associate this IP pool with the right VN in Cisco DNA Center. This will set up the dynamic prefix registration for specific addresses from this prefix as well as set it up for advertisement to the external network.



Mashup - DNA Center

Not Secure https://172.26.199.38/dna/design/home?st-design=network_settings

Apps vSphere Web Client SLT Public Internal DNAC Demo DNAC DNAC Bldg 15 ISE PI Mnt Demo Mnt Wireless Spark Vimarsh Setup DNAC Builds ISE Bldg 15

admin

ADD IP POOL

IP Pool Name
AP

IP Subnet*
1.1.0.0

CIDR Prefix*
/16

Gateway IP Address*
1.1.255.254

DHCP Server(s)
192.168.4.1

DNS Server(s)
192.168.4.1

☐ Overlapping

Cancel Save

Mashup - DNA Center

Not Secure https://172.26.199.38/dna/design/home?st-design=network_settings

Apps vSphere Web Client SLT Public Internal DNAC Demo DNAC DNAC Bldg 15 ISE PI Mnt Demo Mnt Wireless Spark Vimarsh Setup DNAC Builds ISE Bldg 15

admin

Design Policy Provision

Network Hierarchy Network Settings Image Management Network Profiles

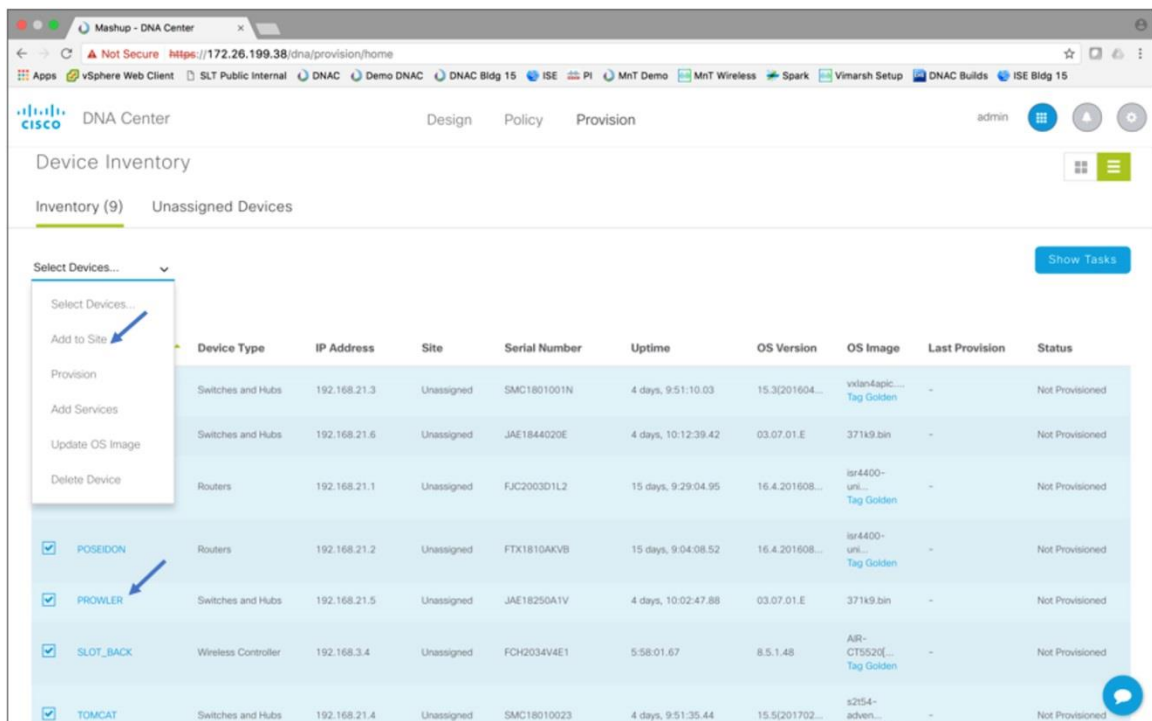
EQ Search Hierarchy

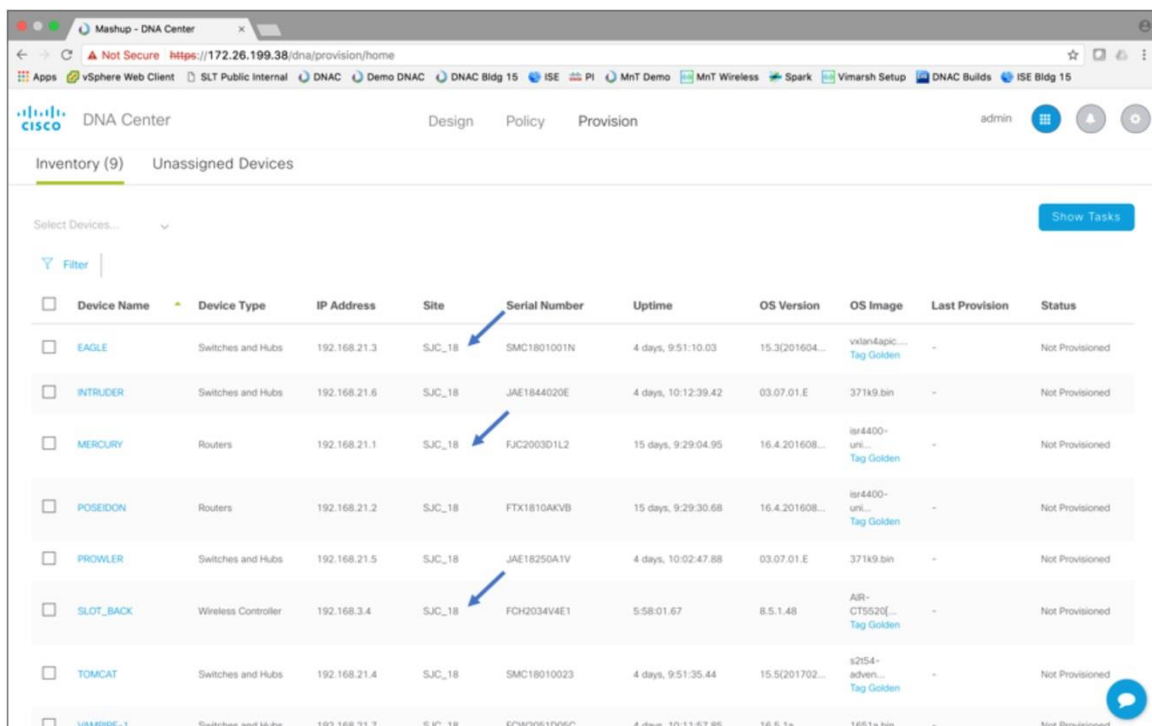
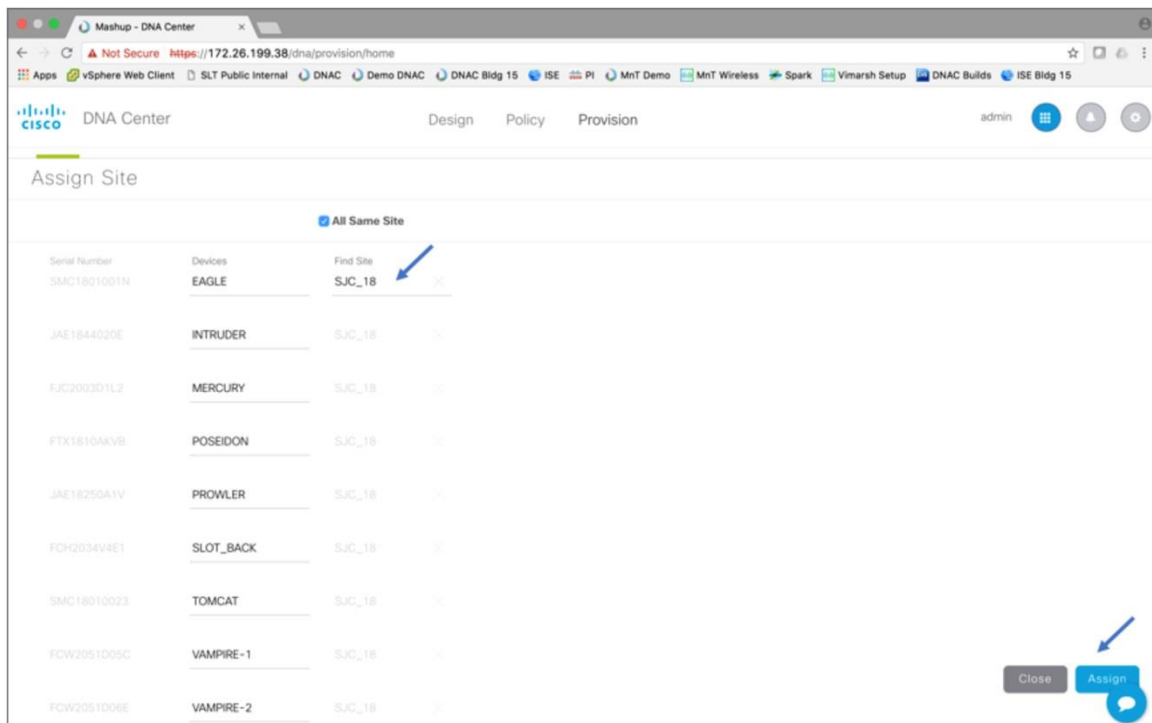
Global
USA

Network Device Credentials IP Address Pools SP Profiles Wireless

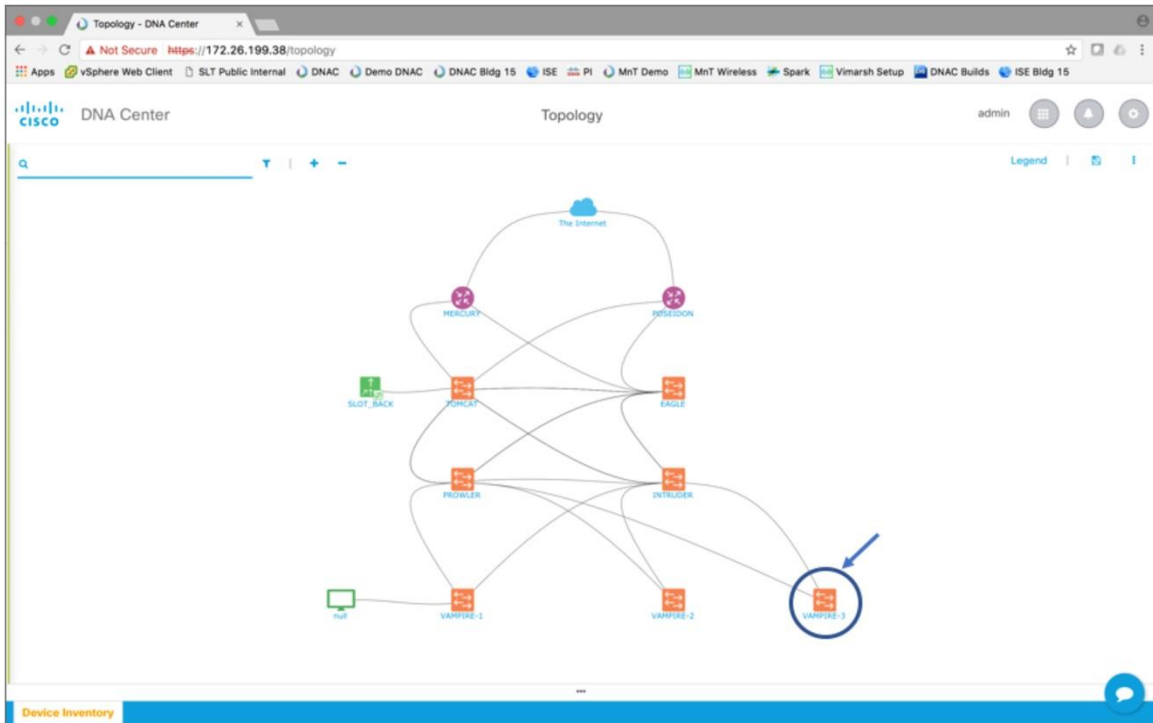
+ Add

Name	IP Subnet Mask	Gateway	DHCP Server	DNS Server	Free Count	Overlapping	Actions
AP	1.1.0.0/16	1.1.255.254	192.168.4.1	192.168.4.1	65536 of 65536	No	Edit Delete
USERS_POOL_001	1.0.0.0/16	1.0.255.254	192.168.4.1	192.168.4.1	65536 of 65536	No	Edit Delete





Step 8. Add a new switch to the network, and assign and provision it to the existing site (SJC_18 in the example below).



Device Inventory

Inventory (7) Unassigned Devices

Select Devices...

- Select Devices...
- Add to Site
- Provision
- Add Services
- Update OS Image
- Delete Device

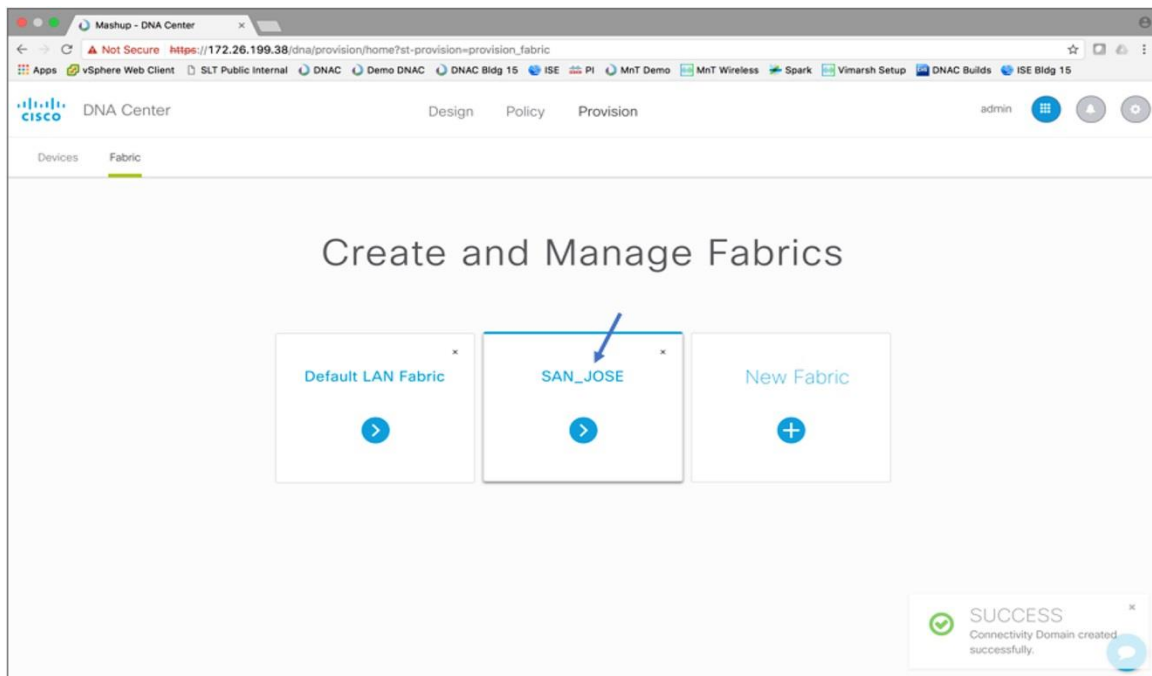
	Device Type	IP Address	Site	Serial Number	Uptime	OS Version	OS Image	Last Provision	Status
<input type="checkbox"/>	Switches and Hubs	192.168.21.3	SJC_18	SMC1801001N	4 days, 19:16:40.94	15.3(201604...	vlan4apic... Tag Golden	-	Not Provisioned
<input type="checkbox"/>	Switches and Hubs	192.168.21.6	SJC_18	JAE1844020E	4 days, 19:25:20.29	03.07.01.E	371k9.bin	-	Not Provisioned
<input type="checkbox"/>	Switches and Hubs	192.168.21.5	SJC_18	JAE18250A1V	4 days, 19:15:15.78	03.07.01.E	371k9.bin	-	Not Provisioned
<input type="checkbox"/>	TOMCAT	192.168.21.4	SJC_18	SMC18010023	4 days, 19:13:44.25	15.5(201702...	vlan4apic... Tag Golden	-	Not Provisioned
<input type="checkbox"/>	VAMPRE-1	192.168.21.7	SJC_18	FCW2051D05C	4 days, 19:02:32.65	16.5.1a	165.1a.bin	-	Not Provisioned
<input type="checkbox"/>	VAMPRE-2	192.168.21.8	SJC_18	FCW2051D06E	4 days, 19:02:22.38	16.5.1a	165.1a.bin	-	Not Provisioned
<input checked="" type="checkbox"/>	VAMPRE-3	192.168.21.9	SJC_18	FOC1744U0UW	0:47:10.60	16.5.1a	165.1a.bin	-	Not Provisioned

Show 10 entries

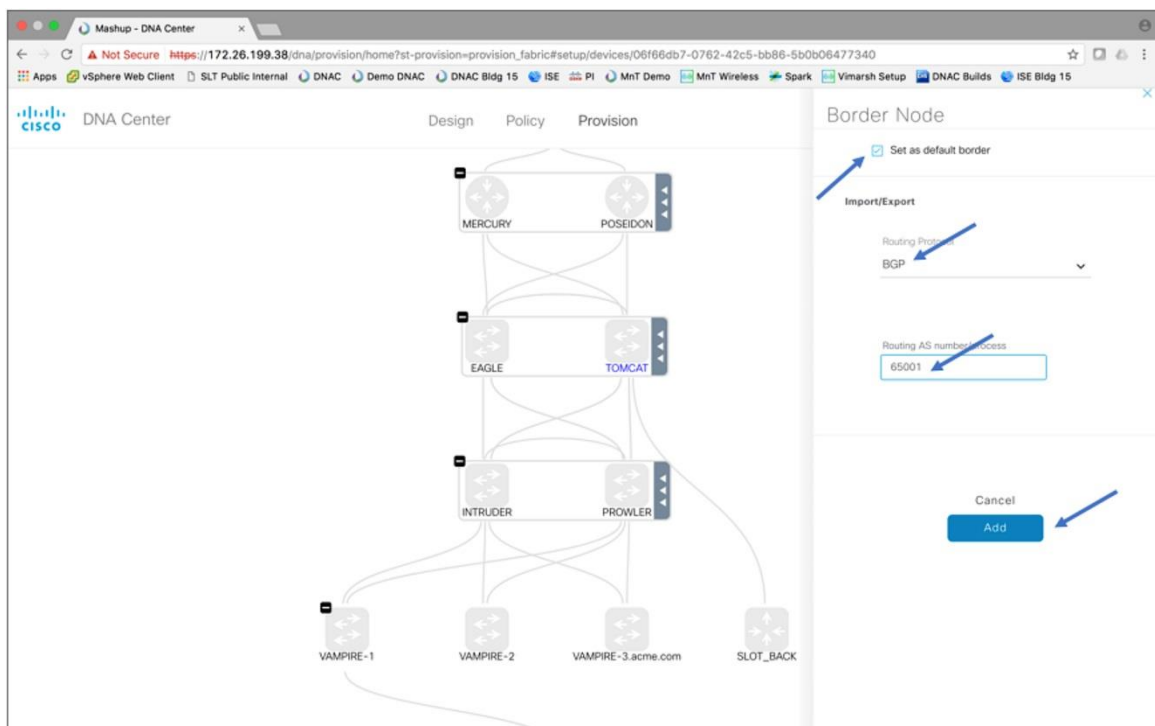
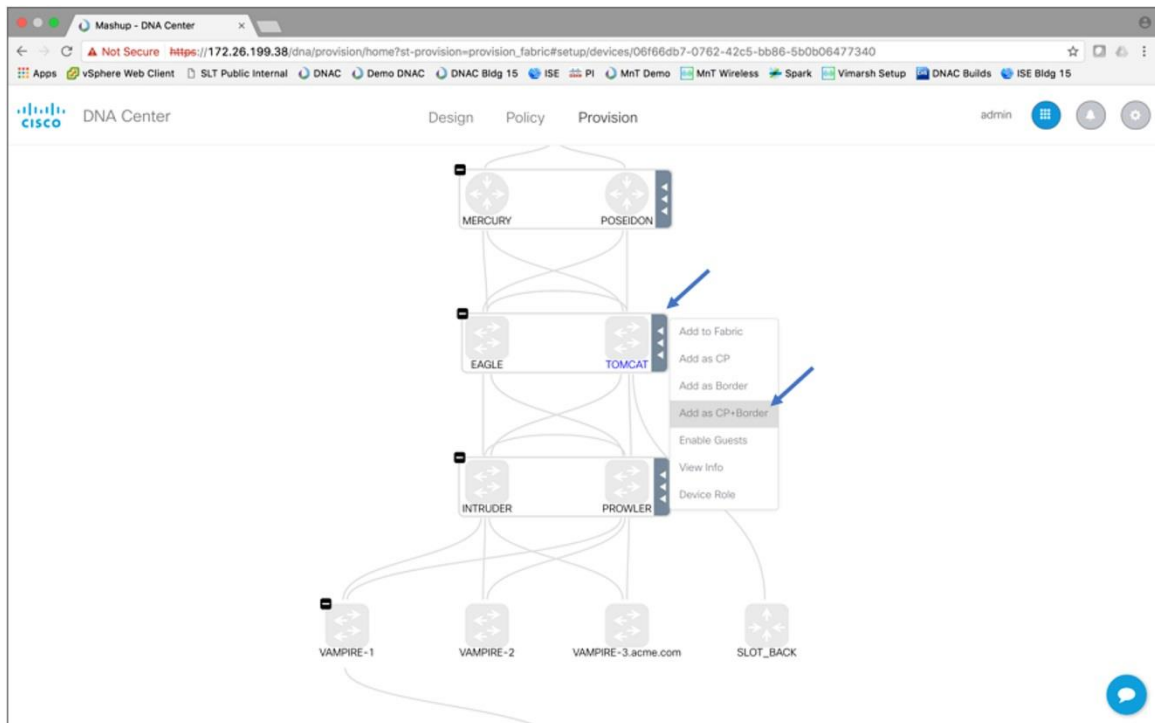
Showing 1 to 7 of 7 entries

Device Name	Device Type	IP Address	Site	Serial Number	Uptime	OS Version	OS Image	Last Provision	Status
EAGLE	Switches and Hubs	192.168.21.3	SJC_18	SMC1801001N	0:15:56.30	15.4(1.1.1)...	s20614.bin	-	Not Provisioned
INTRUDER	Switches and Hubs	192.168.21.6	SJC_18	JAE1844020E	4 days, 20:15:35.38	03.07.01.E	371K9.bin	-	Not Provisioned
PROWLER	Switches and Hubs	192.168.21.5	SJC_18	JAE18250A1V	4 days, 20:05:29.77	03.07.01.E	371K9.bin	-	Not Provisioned
TOMCAT	Switches and Hubs	192.168.21.4	SJC_18	SMC18010023	0:08:06.00	15.4(1.1.1)...	s20614.bin	Jun 17 2017 13:42:50	SUCCESS
VAMPIRE-1	Switches and Hubs	192.168.21.7	SJC_18	FCW2051D05C	4 days, 20:18:22.17	16.5.1a	1651a.bin	-	Not Provisioned
VAMPIRE-2	Switches and Hubs	192.168.21.8	SJC_18	FCW2051D06E	4 days, 20:18:09.47	16.5.1a	1651a.bin	-	Not Provisioned
VAMPIRE-3.acme.com	Switches and Hubs	192.168.21.9	SJC_18	POC174400UW	1:37:58.88	16.5.1a	1651a.bin	Jun 17 2017 12:46:10	SUCCESS

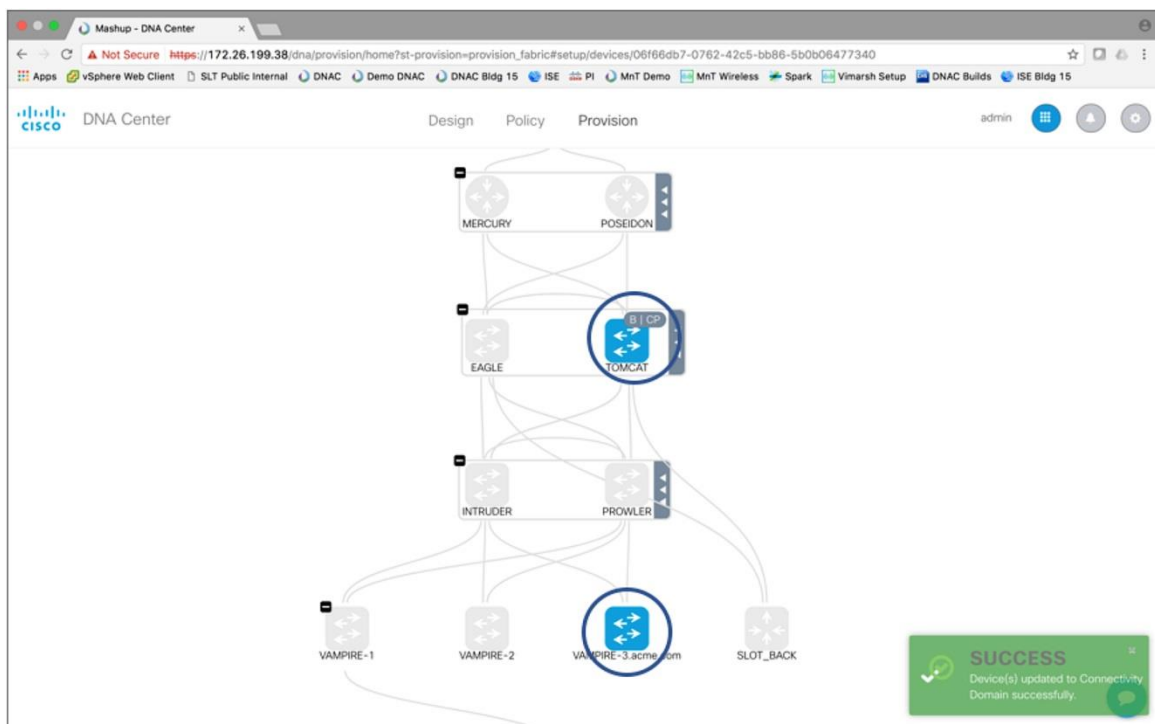
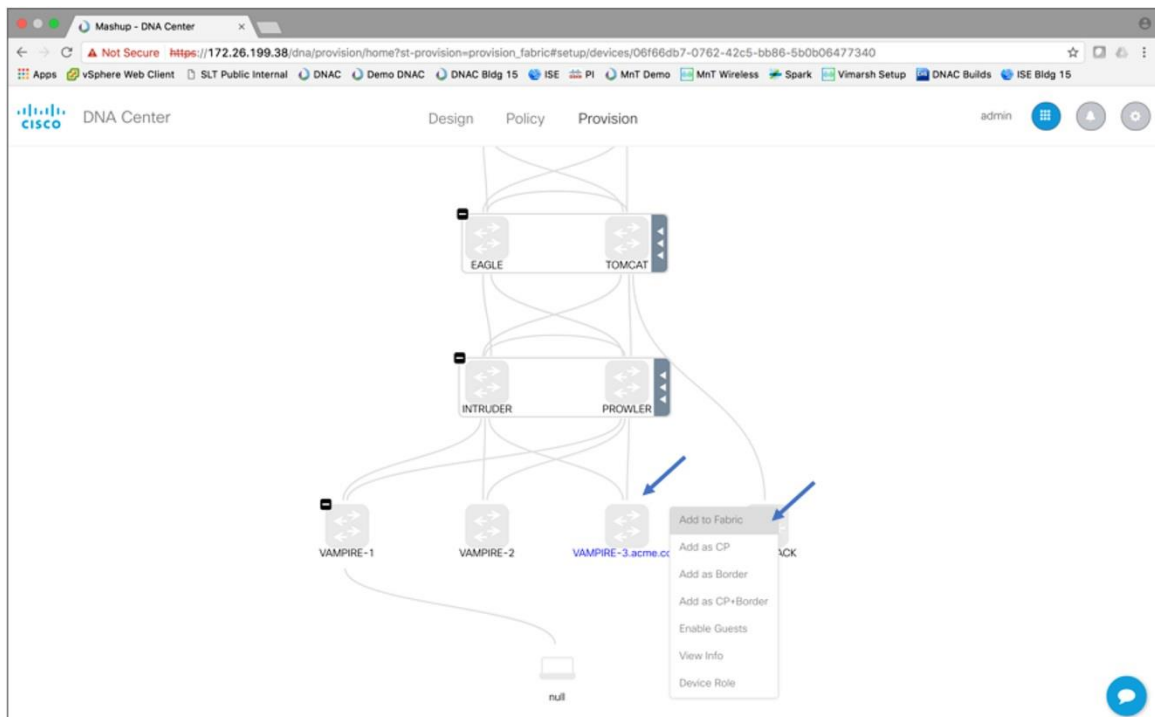
Step 9. Create a new fabric domain, as shown below.



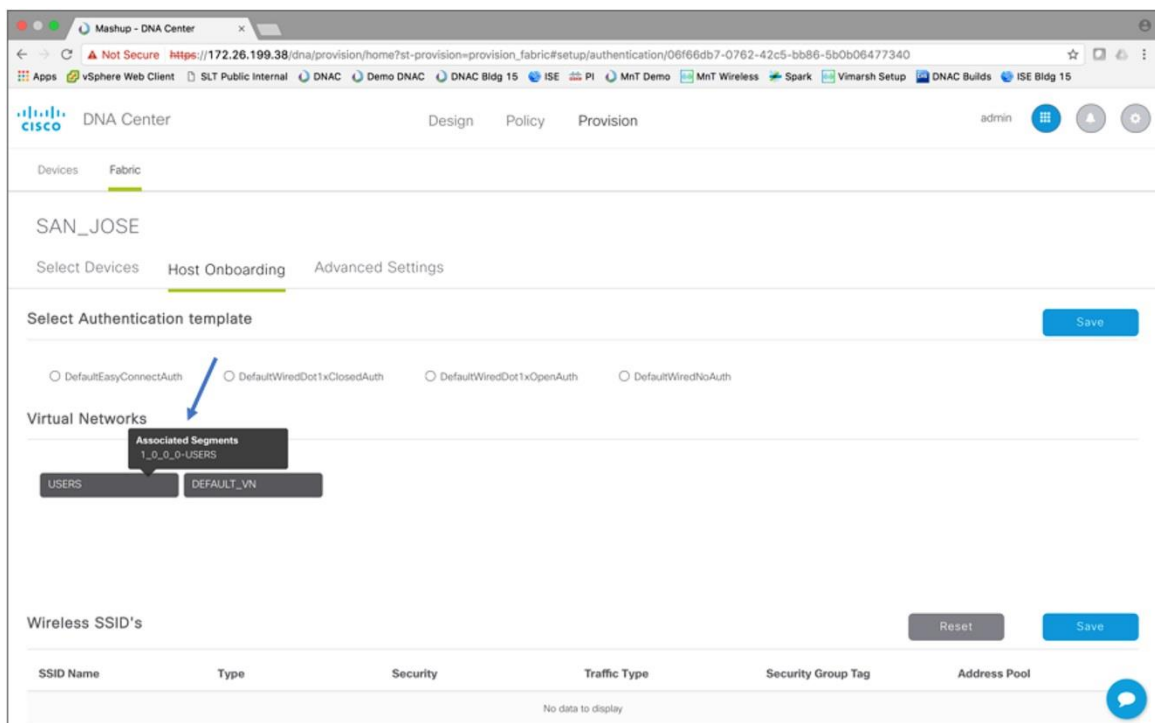
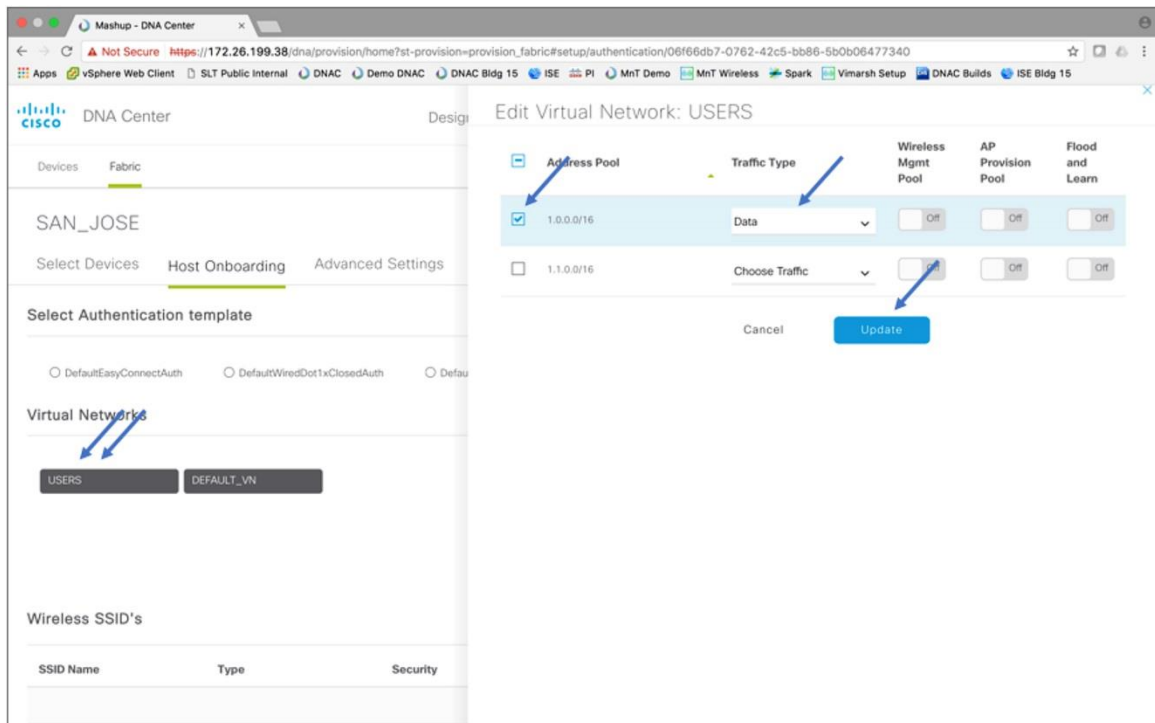
Step 10. Using the Provision tab, add a border node and control plane node to the fabric domain. Configure reachability at the border node for advertising the fabric prefixes out to the external network.



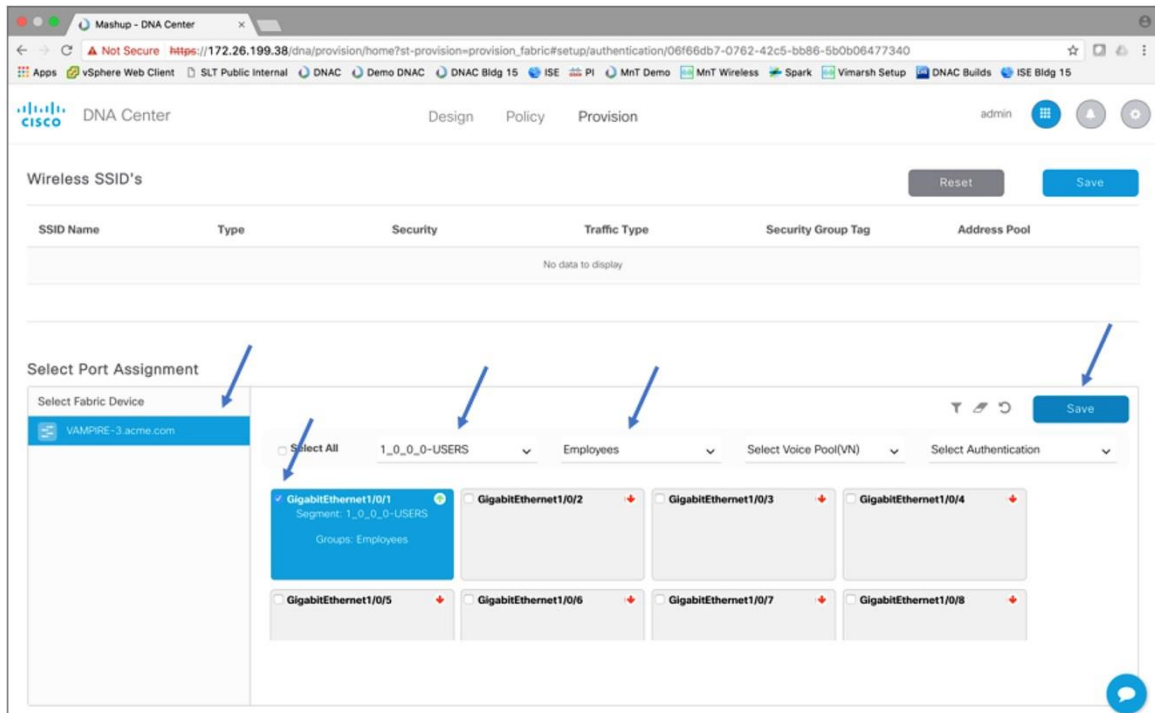
Step 11. Using the Provision tab, provision the newly added switch as a fabric edge node, as shown below.



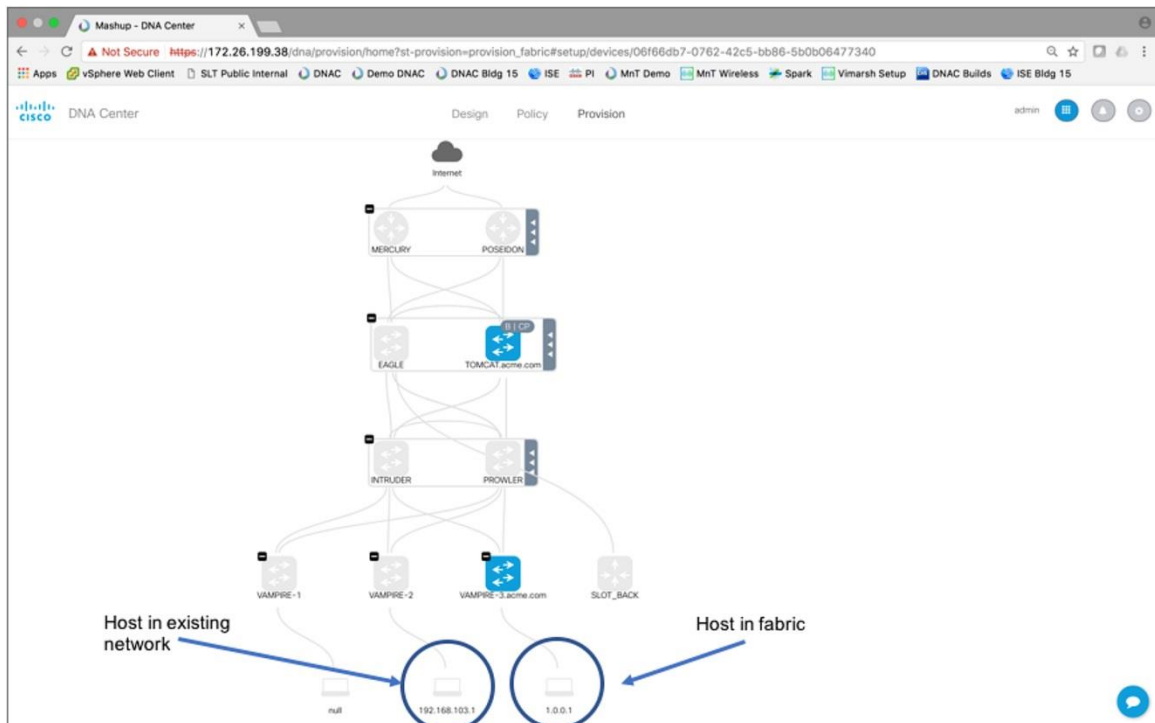
Step 12. Using the Onboarding subtab, provision the newly created IP subnet into the appropriate virtual network, as shown below.

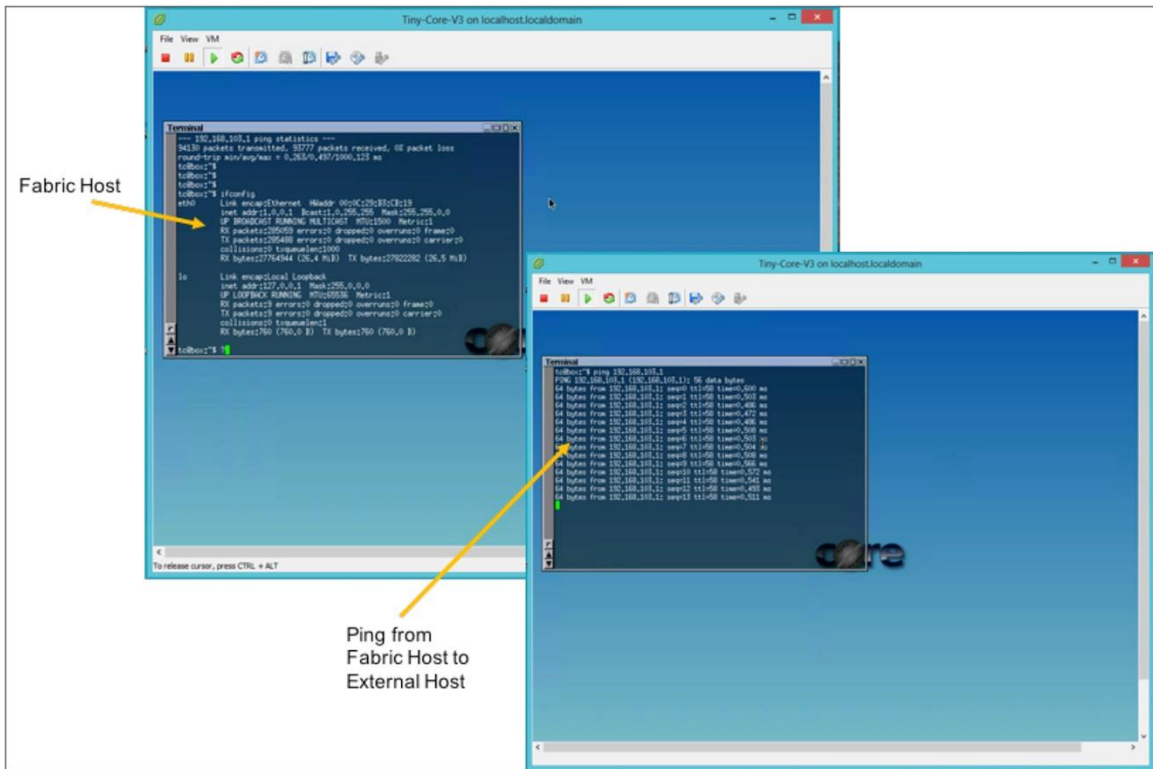


Step 13. Depending on the authentication methods available in the SD-Access fabric, the endpoints will be onboarded into the fabric. The example below shows the static authentication method of mapping a port to the IP subnet.



Step 14. Ensure connectivity between a host in fabric and a host external to the fabric domain, as shown below.



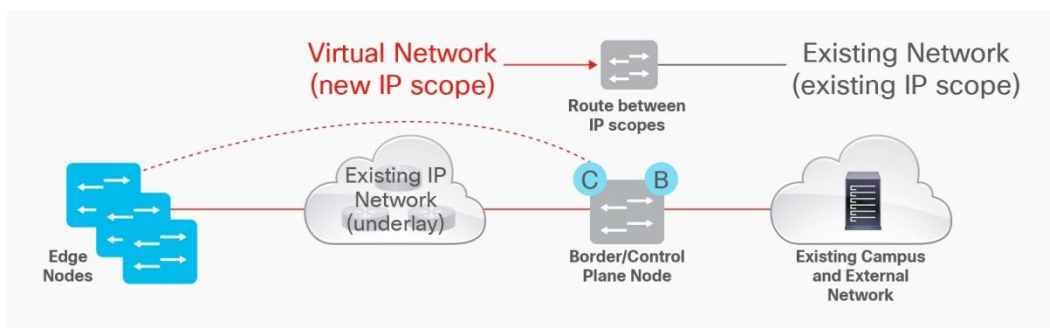


Incrementally expand the fabric edge functionality at the access layer to build out the Cisco SD-Access solution.

Option 2: Map IP Subnets in SD-Access to Outside VLAN

In contrast to the above method, another migration approach for SD-Access is to use the same subnets that exist in the traditional network when creating the SD-Access fabric (Figure 6). This approach saves users the hassle of updating their DHCP scope, firewall rules, and other services-related policies, as today they are based on IP subnets. Since we are using an incremental approach to migration (starting with one fabric edge switch and border and then growing the size of the fabric) during this migration process, certain IP subnets will have to remain extended across both traditional access and fabric access ports.

Figure 6. Using Existing IP Subnets in Cisco SD-Access Fabric



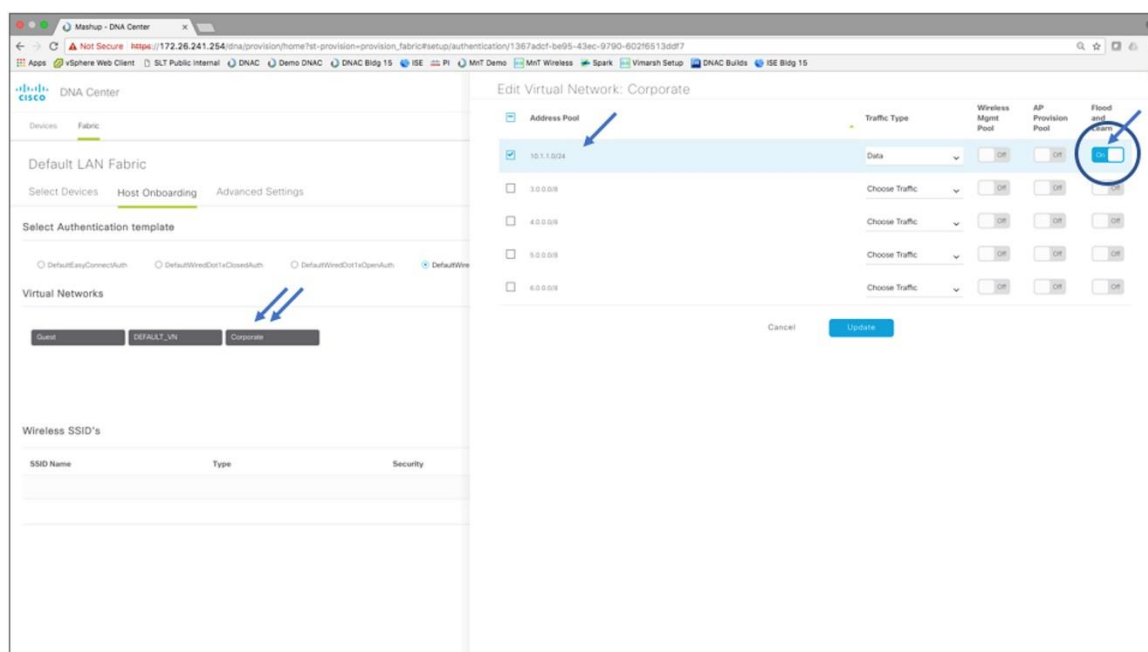
Connecting the SD-Access Fabric to a Traditional Network


The SD-Access fabric consists of a fabric edge node, intermediate node, border node, and control plane node, and these can be created in the fabric based on the procedure mentioned in the previous section.

Once this fabric is created, provision a VN via Cisco DNA Center. Allocate the SGTs to the VN as shown.

Follow the same steps as in the above option to create an SD-Access fabric domain. The only difference would be to provision the same IP pool that is present in the traditional network to onboard the different endpoints in the SD-Access fabric network – this will be the same /24 or /16 that is being used today (the example in Figure 7 uses a 10.1.1.0/24 subnet). In Cisco DNA Center, when this IP pool is created in the fabric, enable the “Flood and Learn” checkbox (as shown in Figure 7) to enable Layer 2 connectivity between the SD-Access fabric and the external traditional network. Cisco DNA Center will automatically choose a VLAN starting from 3000 in the fabric. Associate this IP pool to the appropriate VN in Cisco DNA Center. This will set up the dynamic prefix registration for specific addresses from this prefix, as well as set it up for advertisement to the external network.

Figure 7. SD-Access Fabric Connecting to Traditional Networks





Note: The above feature is available only on Cisco Catalyst 3000 and 9000 platforms starting from 16.6.1.

To configure reachability at the border node for connecting from the SD-Access fabric to the traditional network using the same subnet, perform the following steps.

- Step 1. Connect the existing network to the SD-Access border nodes via trunk links with VLANs allowed.
- Step 2. The switch virtual interface (SVI) for that particular IP subnet in the traditional network that is mapped to a particular VLAN needs to be moved to the border. In other words, we need to configure an SVI for that particular VLAN (IP subnet) on the border.
- Step 3. The IP subnet that has been created in the SD-Access fabric with a VLAN number of 3000 needs to be mapped to that same IP subnet in the traditional network with an existing VLAN (VLAN 10 in the above example). This needs to be manually configured at the border node, as shown below.

Note: The following configuration is from a fabric edge node that has been automated by Cisco DNA Center. It is shown for reference only, to compare with the needed fabric border configuration.

Fabric Edge Configuration (Automated)

```
interface Vlan3000
  description Configured from apic-em
  mac-address 0000.0c9f.fc17
  vrf forwarding Corporate
  ip address 10.1.1.254 255.255.255.0
  ip helper-address global 192.168.4.1
  no ip redirects
  ip local-proxy-arp
  ip route-cache same-interface
  lisp mobility 10_1_1_0-Corporate
!
router lisp
!
instance-id 10
  service ethernet
    eid-table vlan 3000
    database-mapping mac locator-set rloc_01899313-4fbf-45f9-9f58-ca5a157d73ca
    exit-service-ethernet
  !
exit-instance-id
!
instance-id 4098
  dynamic-eid 10_1_1_0-Corporate
    database-mapping 10.1.1.0/24 locator-set rloc_01899313-4fbf-45f9-9f58-ca5a157d73ca
    exit-dynamic-eid
  !
service ipv4
  eid-table vrf Corporate
  exit-service-ipv4
!
exit-instance-id
```

Fabric Border Configuration (Manual)

```
interface Vlan10
  description Configured MANUALLY
  mac-address 0000.0c9f.fc17
  vrf forwarding Corporate
  ip address 10.1.1.254 255.255.255.0
  ip helper-address global 192.168.4.1
  no ip redirects
  ip local-proxy-arp
  ip route-cache same-interface
  lisp mobility 10_1_1_0-Corporate
!
router lisp
!
instance-id 10
  service ethernet
    eid-table vlan 10
    database-mapping mac locator-set rloc_01899313-4fbf-45f9-9f58-ca5a157d73ca
    exit-service-ethernet
  !
exit-instance-id
!
instance-id 4098
  dynamic-eid 10_1_1_0-Corporate
    database-mapping 10.1.1.0/24 locator-set rloc_01899313-4fbf-45f9-9f58-ca5a157d73ca
    exit-dynamic-eid
  !
service ipv4
  eid-table vrf Corporate
  exit-service-ipv4
!
exit-instance-id
```

Once these steps are done, configure reachability at the border node for advertising the SD-Access fabric prefixes and the traditional network prefixes out to the external network. From the border node, configure subinterfaces – one in the GRT and one in the VN – to advertise both the SD-Access fabric and the traditional network prefixes.

Service Interworking with SD-Access

When migrating the traditional network into the SD-Access fabric, we are introducing the virtual network (VRF) concept into the fabric. Due to the introduction of this additional policy construct, we need to be able to still communicate with our services infrastructure. The services infrastructure generally resides outside the fabric domain and contains the following elements:

- Identity services (for example, AAA/RADIUS)
- Domain Name Services (DNS)
- Dynamic Host Configuration (DHCP)
- IP Address Management (IPAM)
- Monitoring tools (for example, SNMP)
- Data collectors (for example, NetFlow, syslog)
- Firewalls
- Other infrastructure elements

Since these reside outside of the fabric, the border is responsible for interconnecting the fabric with the services infrastructure. The services infrastructure will generally be deployed using one of two models:

- Shared services in GRT
- Shared services in dedicated VRF table

Shared Services in GRT

In this design option, existing services are currently deployed in the global routing table. This does not impose any specific requirements for the SD-Access fabric as such. To achieve continuity of connectivity in the fabric for the users, the GRT in the traditional portion of the network will peer using BGP/IGP with the virtual network (VRF) in the SD-Access fabric. This ensures that both of the domains can reach each other (Figures 8 and 9).

Figure 8. Configuration on the Fabric Border for External Connectivity

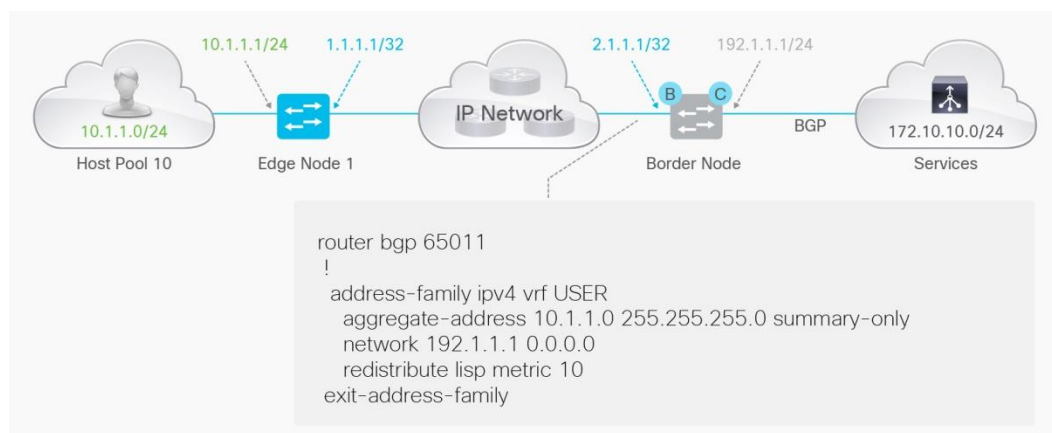
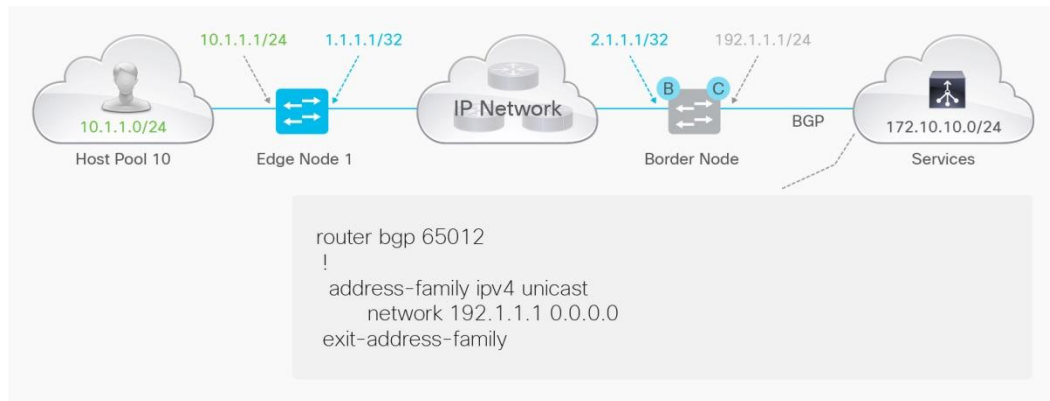


Figure 9. Configuration on the External Router for Fabric Connectivity

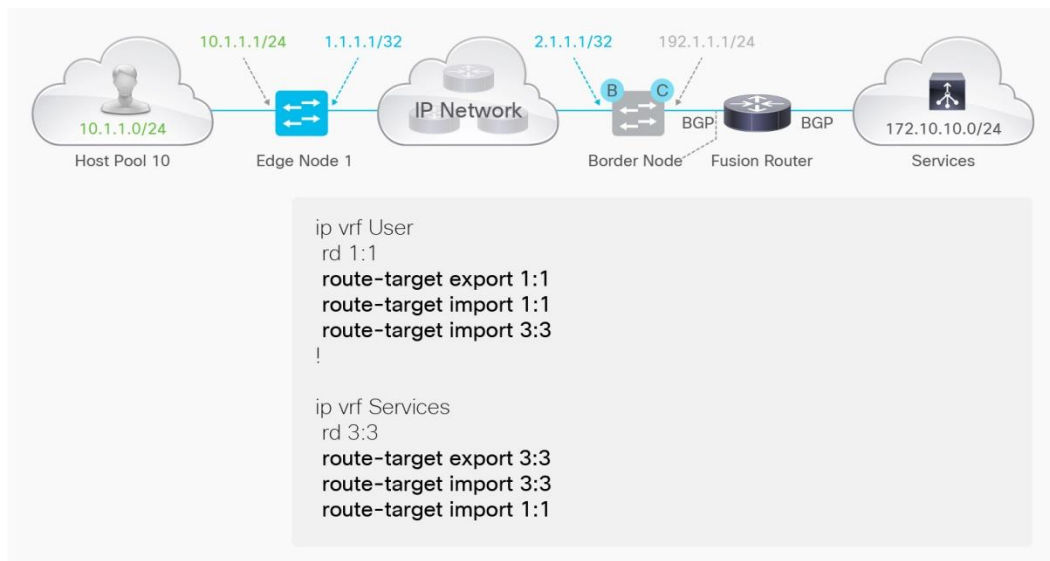


Shared Services in a Dedicated VRF Table

In this design option, the services infrastructure is placed into a dedicated VRF context of its own and VRF route leaking needs to be provided in order for the virtual network (VRF) in the SD-Access fabric to have continuity of connectivity to the services infrastructure. To achieve continuity of connectivity in the fabric for the users, a fusion router is connected to the SD-Access border through VRF-Lite using BGP/IGP, and the services infrastructure is connected to the fusion router in a services VRF (Figure 10).

A fusion router is a device that can provide VRFs with connectivity to the services, the Internet, or even inter-VRF connectivity.

Figure 10. Configuration on the Fusion Router for Fabric to External Connectivity



Summary

This guide helps you understand the different migration options available to convert a traditional network to an SD-Access fabric network.




Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at <https://www.cisco.com/go/offices>.

 Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <https://www.cisco.com/go/trademarks>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)