



Software-Defined Access for Distributed Campus Solution Adoption Prescriptive Reference – Deployment Guide

July, 2019

Table of Contents

Introduction	6
About Cisco DNA Center	6
About Cisco Digital Network Architecture	6
About Software-Defined Access	7
About This Guide	7
What is Covered in This Guide?	8
What is Not Covered in This Guide?	8
Nomenclature Conventions	8
Define	9
Fabric Architectural Components	9
About Cisco Software-Defined Access for Distributed Campus	10
Design	11
SD-Access Transit	11
Transit Control Plane Nodes	11
Transit Control Plane Deployment Location	11
Key Considerations	11
Topology Overview	12
Identity Services Engine and Shared Services Overview	13
Deploy	14
Process 1: Discovering the Network Infrastructure	14
Procedure 1: Discover Network Devices - Discovery Tool	14
Procedure 2: Define the Network Device Role - Inventory Tool	20
Process 2: Integrating Cisco DNA Center with the Identity Services Engine	21
Procedure 1: Configure Authentication and Policy Servers in Cisco DNA Center	21
Process 3: Using the Design Application	26
Procedure 1: Create the Network Hierarchy	26
Procedure 2: Define Network Settings and Services	29
Procedure 3: Create a Global IP Address Pool	33
Procedure 4: Reserve IP Address Pools	34
Procedure 5: Design Enterprise Wireless SSIDs for SD-Access Wireless	36
Procedure 6: Design a Guest Wireless SSID for SD-Access Wireless	39
Process 4: Creating Segmentation with the Cisco DNA Center Policy Application	42
Procedure 1: Add an Enterprise Overlay Virtual Network - Macro Segmentation	42

Procedure 2: Add a Guest Overlay Virtual Network – Macro Segmentation	43
Procedure 3: Create Group-Based Access Control Policies using SGTs – Micro Segmentation	44
Process 5: Deploying SD-Access with the Provision Application	47
About Provisioning in SD-Access	47
Procedure 1: Create Credentials in ISE	48
Procedure 2: Assign Network Devices to Site and Provision Network Settings	49
Procedure 3: Assign Wireless LAN Controller to Site and Provision Network Settings.....	52
Process 6: Provisioning the Fabric Overlay	54
Procedure 1: Create IP-Based Transits – Fabric Provisioning Part I.....	55
Procedure 2: Create SD-Access Transit – Fabric Provisioning Part I (Continued)	56
Procedure 3: Create a Fabric Domain – Fabric Provisioning Part 2	58
Procedure 4: Add Fabric-Enabled Sites to the Fabric Domain – Fabric Provisioning Part 2 (Continued)	59
Process 7: Assigning Fabric Roles – Fabric Provisioning Part 3	61
About Fabric Border Nodes	61
Procedure 1: Provisioning a Site Fabric Overlay.....	63
Procedure 2: Provisioning Internal Border Nodes.....	68
Procedure 3: Provisioning Redundant Layer-3 Handoffs.....	72
Process 8: Configuring Host Onboarding – Fabric Provisioning Part 4.....	75
About Authentication Templates.....	76
Procedure 1: Assign Authentication Template and Create Wired Host Pool – Host Onboarding Part 1	77
Procedure 2: Create INFRA_VN Host Pool – Host Onboarding Part 2	78
Procedure 3: Assign SSID Subnet and Assign Access Point Ports– Host Onboarding Part 3	80
Operate	83
Process 1: Providing Access to Shared Services	83
About Fusion Routers	83
Procedure 1: Use the Provision Application for BGP and VRF-Lite information – Part 1	84
Procedure 2: Configure VRF Definitions on the Fusion Routers – Command Runner Tool	85
Procedure 3: Create Layer-3 Connectivity Between Border Nodes and Fusion Routers – Fusion Routers Part I.....	88
Procedure 4: Establish BGP Adjacencies Between Fusion Routers and Border Nodes – Fusion Routers Part 2.....	91
Procedure 5: Use Route-Targets to Leak Shared Services Routes – Fusion Routers Part 3.....	94
Process 2: Provisioning and Verifying Access Points	95
Procedure 1: Access Point Verification	98
Process 3: Creating IBGP Adjacencies Between Redundant External Border Nodes	99
Procedure 1: Configure IBGP – Routing Platforms.....	100
Process 4: Providing Internet Access for SD-Access for Distributed Campus.....	102

Procedure 1: Provision the Connected-to-Internet Border Node.....	103
Internet Edge Topology Overview	109
Configuring Internet Connectivity.....	110
Procedure 2: Use the Provision Application for BGP and VRF-Lite information – Part 2.....	110
Procedure 3: Create Layer-3 Connectivity Between Border Nodes and Internet Edge Routers	111
Procedure 4: Create BGP Adjacencies Between Border Nodes and Internet Edge Routers	114
About Default Route Advertisement Using BGP	117
Procedure 5: Advertise the Default Route From the Internet Edge Routers to Fabric VNs.....	118
Procedure 6: Verify Default Route Advertisement	118
Procedure 7: Verify Default Route Via LISP – Advanced and Optional	120
Appendix A: Hardware and Software Code Versions	127
Appendix B: Additional References.....	128
Appendix C: Advanced Topology Diagrams.....	129
High-Level Overview	129
Enterprise Architecture Model Topology.....	129
Underlay Connectivity.....	130
Overlay Connectivity.....	131
BGP Autonomous System Overview	132
Fabric Role Overview.....	133
Site-1 Fabric Roles.....	133
Site-2 Fabric Roles.....	134
Site-3 Fabric Roles.....	135
Site-4 Fabric Roles.....	136
Site-5 Fabric Roles.....	137
Branch Fabric Roles.....	138
Layer-2 Overview	139
Site-1 Layer-2.....	139
Site-2 Layer-2.....	140
Site-3 Layer-2.....	141
Site-4 Layer-2.....	142
Site-5 Layer-2.....	143
Branch Layer-2.....	144
Layer-3 Overview	145
Individual Sites Loopback IP Schema.....	145
IP Address Pools	151

Shared Services and Internet Access	153
Internet Access	153
Shared Services	154
Appendix D: Route Leaking.....	155
Procedure 1: Create Layer-3 Connectivity Between Border Nodes and Fusion Routers.....	155
Procedure 2: Establish BGP Adjacencies Between Fusion Routers and Border Nodes	158
Procedure 3: Create IP Prefix lists to Match Fabric Subnets.....	160
Procedure 4: Create Route-maps that Match IP Prefix Lists.....	160
Procedure 5: Leak routes.....	161
Appendix E: IBGP Between Redundant Devices	164
IBGP for Switching Platforms	164
About this guide	168
Feedback & Discussion.....	168

Introduction

About Cisco DNA Center

Cisco DNA Center is the foundational controller and analytics platform at the heart of Cisco's Intent-Based Network (IBN) for large and midsize organizations. Intent-based networking embodies the difference between a network that needs continuous attention and one that simply understands what the organization needs and makes it happen. It is the difference between doing thousands of tasks manually and having an automated system that helps focus on business goals. Cisco DNA Center provides a centralized management dashboard for complete control of this new network. With this platform, IT can simplify network operations, proactively manage the network, provide consistent wired and wireless policy, and correlate insights with contextual cognitive analytics.

Cisco DNA Center is a dedicated hardware appliance powered through a software collection of applications, processes, services, packages, and tools, and it is the centerpiece for Cisco® Digital Network Architecture (Cisco DNA™). This software provides full automation capabilities for provisioning and change management, reducing operations by minimizing the touch time required to maintain the network. It also provides visibility and network assurance through intelligent analytics that pull telemetry data from everywhere in the network. This reduces troubleshooting time and addresses network issues through a single pane of management.

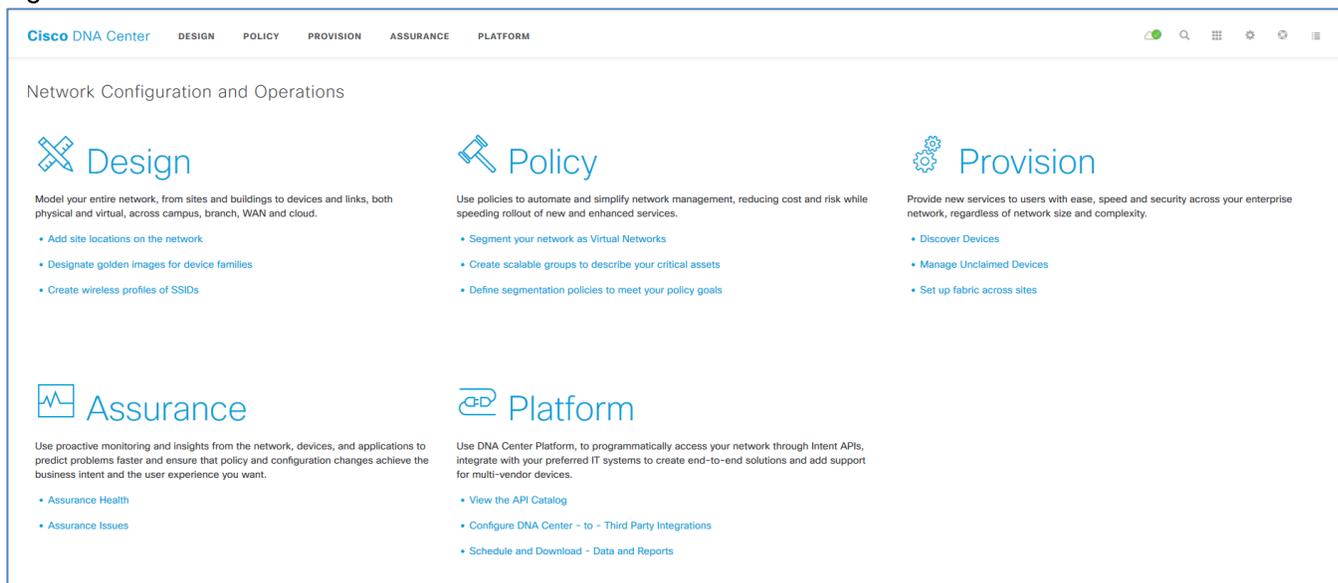
This interconnection of automation and assurance forms a continuous validation-and-verification loop, driving business intent and enabling faster innovation, lower cost and complexity, and enhanced security and compliance.

About Cisco Digital Network Architecture

Cisco® Digital Network Architecture (Cisco DNA™) provides a roadmap for digitization and a path to realization of the immediate benefits of network automation, assurance, and security. Cisco DNA is an open, extensible, software-driven architecture that accelerates and simplifies enterprise network operations while enabling business requirements to be captured and translated into network policy. The result is constant alignment of the network to the business intent.

Cisco DNA begins with the foundation of a digital-ready infrastructure that includes routers, switches, access-points, and wireless LAN controllers (WLC). The Identity Services Engine (ISE) is the key policy manager for the Cisco DNA solution. The centerpiece of Cisco DNA is the Cisco DNA Center controller which empowers simplified workflows using the Design, Provision, Policy, and Assurance applications.

Figure 1 Cisco DNA Center Dashboard



About Software-Defined Access

Cisco Software-Defined Access (SD-Access) is the Cisco DNA evolution from traditional campus LAN designs to networks that directly implement the intent of an organization. It is the intent-based networking solution for the Enterprise built on the principles of Cisco DNA. The SD-Access solution is enabled by an application package that runs as part of the Cisco DNA Center software and provides automated end-to-end segmentation to separate user, device, and application traffic. These user access policies are automated so that organizations can ensure that the right policies are established for any user or device with any application anywhere in the network.

SD-Access uses logic blocks called fabrics which leverage virtual network overlays that are driven through programmability and automation to create mobility, segmentation, and visibility. Network virtualization becomes easy to deploy through software-defined segmentation and policy for wired and wireless campus networks. Single physical networks are abstracted and can host one or more logical networks which are orchestrated through software. Error-prone manual operations in these dynamic environments are circumvented altogether, providing consistent policy for users as they move around the wired and wireless network.

About This Guide

This guide provides technical guidance for designing, deploying, and operating Software-Defined Access for Distributed Campus. It focuses on Cisco DNA Center to deploy the solution after the initial bootstrap of the network and supporting infrastructure is complete.

Figure 2 This Guide's Four Sections



This guide contains four major sections:

- Step 1:** The **DEFINE** section defines the problem being solved with the SD-Access for Distributed Campus solution and provides information on planning for the deployment and other considerations.
- Step 2:** The **DESIGN** section shows a typical deployment topology and discusses Identity Services Engine and shared services considerations.
- Step 3:** The **DEPLOY** section provides information and steps for the various workflows to deploy the solution along with recommended practices.
- Step 4:** The **OPERATE** section demonstrates the manual configurations necessary for shared services, Internet access, and BGP between redundant device types.

What is Covered in This Guide?

This guide provides guidance to SD-Access customers deploying a unified and automated policy across multiple physical locations in a metro-area network. The process, procedures, and steps listed in this guide are working configurations verified with the Cisco DNA Center, ISE, IOS, IOS-XE, and AireOS code versions listed in [Appendix A](#).

What is Not Covered in This Guide?

Although this deployment guide is about Cisco DNA Center and SD-Access, it does not cover the initial bootstrap and installation of the Appliance or of the ISE distributed deployment. Shared services installation and deployment, such as DHCP and DNS, along with the network connectivity configuration between various infrastructure components, routers, and switches, are not specifically covered. SWIM (Software Image Management), LAN automation, multicast, Layer-2 handoff, fabric-in-a-box, and Cisco Catalyst 9800 embedded wireless on Catalyst series switches are also not covered.

For more information on these items, please see additional references in [Appendix B](#).

Nomenclature Conventions

Known routes, destinations, and prefixes are locations both inside the fabric and in the shared services Domains (DHCP, DNS, WLC, and ISE) that are registered with and known to the fabric control plane nodes.

Unknown routes, destinations, and prefixes are locations on the Global Internet that are not known or registered with the fabric control plane nodes.

LISP encapsulation or LISP data plane encapsulation in the context of SD-Access refers to the [VXLAN-GPO](#) encapsulation. For brevity, this may be referred to as VXLAN encapsulation. However, this is not meant to indicate or infer that the encapsulation method is the same VXLAN in [RFC 7348](#) or associated with [VXLAN MP-BGP EVPN](#).

Define

This section provides a high-level overview of the Software-Defined Access solution and components with a focus on elements related to distributed campus.

Fabric Architectural Components

The SD-Access 1.2 solution supports provisioning of the following fabric constructs:

Fabric edge node: Equivalent of an access layer switch in a traditional campus LAN design. Endpoints, IP phones, and access points are directly connected to edge nodes.

Fabric control plane node: Based on the LISP Map-Server (MS) and Map-Resolver (MR) functionality. The control plane node's host tracking database tracks all endpoints in a fabric site and associates the endpoints to fabric nodes in what is known as an EID-to-RLOC binding in LISP.

Fabric border node: Serves as the gateway between the SD-Access fabric site and networks external to the fabric. The border node is the device physically connected to a transit or to a next-hop device connected to the outside world.

Fabric site: An independent fabric that includes a control plane node and edge node and usually includes an ISE Policy Service Node (PSN) and fabric-mode WLC. A fabric border node is required to allow traffic to egress and ingress the fabric site.

Virtual Network (VN): Ostensibly a VRF definition. VNs are created in the Policy application and provisioned to the fabric nodes as a VRF instance. VN and VRF are used interchangeably in this document.

Scalable Group Tag (SGT): A Cisco TrustSec component that operates as a form of metadata to provide logical segmentation based on group membership.

Transit: Connects a fabric site to an external network (IP-Based transit) or to one or more fabric sites (SD-Access transit). IP-Based transit networks connect the fabric to external networks using VRF-lite. SD-Access transits carry SGT and VN information inherently carrying policy and segmentation between fabric sites and do not require or use VRF-lite.

Fabric domain: Encompasses one or more fabric sites and any corresponding transit(s) associated with those sites.

Transit control plane node: The transit control plane node's database tracks all aggregate routes for the fabric domain and associates these routes to fabric sites. Its functionality is based on LISP Map-Server (MS) and Map-Resolver (MR).

Fabric in a Box: Combines the fabric border node, fabric control plane node, and fabric edge node functionality on the same switch or switch stack.

Host Pool: The binding of a reserved IP address pool to a Virtual Network which associates a segment to a VRF.

About Cisco Software-Defined Access for Distributed Campus

Cisco Software-Defined Access (SD-Access) for Distributed Campus is a metro-area solution that connects multiple, independent fabric sites together while maintaining the security policy constructs (VRFs and SGTs) across these sites. While multi-site environments and deployments have been supported with SD-Access for some time, there has not been an automated and simplistic way to maintain policy between sites. At each site's fabric border node, fabric packets were de-encapsulated into native IP. Combined with SXP, policy could be carried between sites using native encapsulation. However, this policy configuration was manual, mandated use of SXP to extend policy between sites, and involved complex mappings of IP to SGT bindings within the Identity Services Engine.

With SD-Access for Distributed Campus, SXP is not required, the configurations are automated, and the complex mappings are simplified. This solution enables inter-site communication using consistent, end-to-end automation and policy across the metro network.

Software-Defined Access for Distributed Campus uses control plane signaling from the LISP protocol and keeps packets in the fabric VXLAN encapsulation between fabric sites. This maintains the macro- and micro-segmentation policy constructs of VRFs and SGT, respectively, between fabric sites. The original Ethernet header of the packet is preserved to enable the Layer-2 overlay service of SD-Access Wireless. The result is a network that is address-agnostic because policy is maintained through group membership.

Design

This section introduces the SD-Access transit and transit control plane nodes along with key considerations, shows the deployment topology, and discusses the Identity Services Engine and shared services.

SD-Access Transit

The core components enabling the Distributed Campus solution are the SD-Access transit and the transit control plane nodes. Both are new architectural constructs introduced with this solution. The SD-Access transit is simply the physical metro-area connection between fabric sites in the same city, metropolitan area, or between buildings in a large enterprise campus.

The key consideration for the Distributed Campus design using SD-Access transit is that the network between fabric sites and to Cisco DNA Center should be created with campus-like connectivity. The connections should be high-bandwidth (Ethernet full port speed with no sub-rate services), low latency (less than 10ms as a general guideline), and should accommodate the MTU setting used for SD-Access in the campus network (typically 9100 bytes). The physical connectivity can be direct fiber connections, leased dark fiber, Ethernet over wavelengths on a WDM system, or metro Ethernet systems (VPLS, etc.) supporting similar bandwidth, port rate, delay, and MTU connectivity capabilities.

Transit Control Plane Nodes

The transit control plane nodes track all aggregate routes for the fabric domain and associate these routes to fabric sites. When traffic from an endpoint in one site needs to send traffic to an endpoint in another site, the transit control plane node is queried to determine to which site's border node this traffic should be sent. The role of transit control plane nodes is to learn which prefixes are associated with each fabric site and to direct traffic to these sites across the SD-Access transit using control-plane signaling.

Transit Control Plane Deployment Location

The transit control plane nodes do not have to be physically deployed in the Transit Area (the metro connection between sites) nor do they need to be dedicated to their own fabric site, although common topology documentation often represents them in this way. For the prescriptive configuration in this guide, a pair of ISR 4451-X routers were used as the transit control plane nodes. These routers were deployed in their own dedicated site, only accessible through the SD-Access transit Metro-E network, although this is not a requirement.

While accessible only via the Transit Area, these routers do not act as a physical-transit hop in the data packet forwarding path. Rather, they function similarly to a DNS server in that they are queried for information, even though data packets do not transit through them. This is a key consideration.

Key Considerations

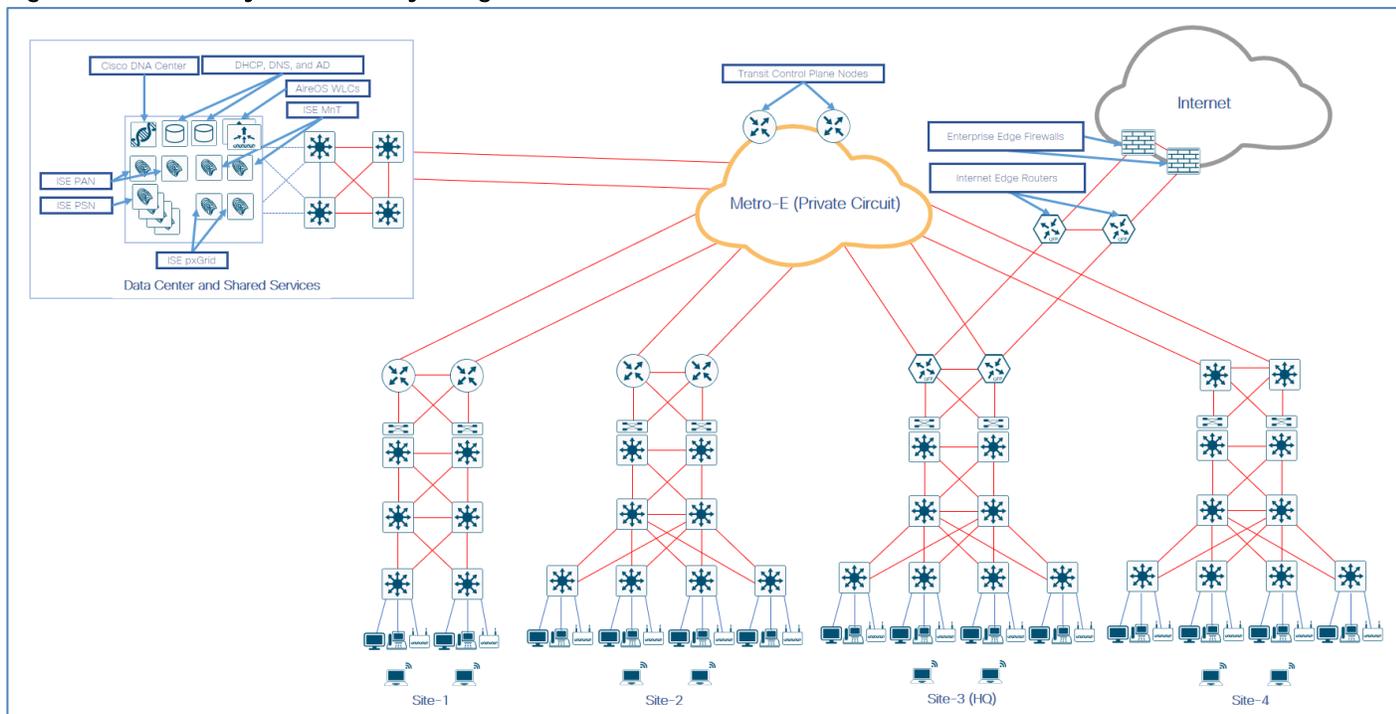
The transit between sites is best represented and most commonly deployed as direct or leased fiber over a Metro Ethernet system. While Metro-E has several different varieties (VPLS, VPWS, etc.), the edge routers and switches of each site ultimately exchange underlay routes through an Interior Gateway Routing (IGP) protocol. In SD-Access, this is commonly done using the IS-IS routing protocol, although other IGPs are supported.

IP reachability must exist between fabric sites. Specifically, there must be a known underlay route between the Loopback 0 interfaces on all fabric nodes. Existing BGP configurations and peerings on the transit control plane nodes could have complex interactions with the Cisco DNA Center provisioned configuration and should be avoided. The transit control plane nodes should have IP reachability to the fabric sites through an IGP before being discovered or provisioned.

Traversing the transit control plane nodes in the data forwarding path between sites has not been validated. Transit control plane nodes should always be deployed as a pair of devices to provide resiliency and high availability.

Topology Overview

Figure 3 Underlay Connectivity – High Level Overview



For topology diagram simplicity, redundant links from devices to the private circuit, Internet edge routers, and Internet edge firewalls are not shown although they are present in the deployment. Only the four sites are shown here with additional topology diagrams available in [Appendix C](#).

This deployment guide topology contains twelve fabric sites in a metropolitan area. It is a large enterprise campus with dispersed buildings in a similar geographic area with each building operating as an independent fabric site. Each site is connected to a Metro-E circuit at the Enterprise Edge block. A thirteenth site is the Data Center where shared services are located along with the Cisco DNA Center cluster. The Data Center site has a direct connection to both the Metro-E circuit and to Site-1 in the network. Site-3 is the main head-end location (represented as Headquarters – HQ) and is connected to both the Metro-E circuit and the upstream Internet edge routers which themselves connect to the Internet edge firewalls. NAT and ultimately access to the Internet is provided through these firewalls. The result is that all traffic directed to the Internet from any site must be carried across the Metro-E circuit and then egressed out of Site-3 (HQ).

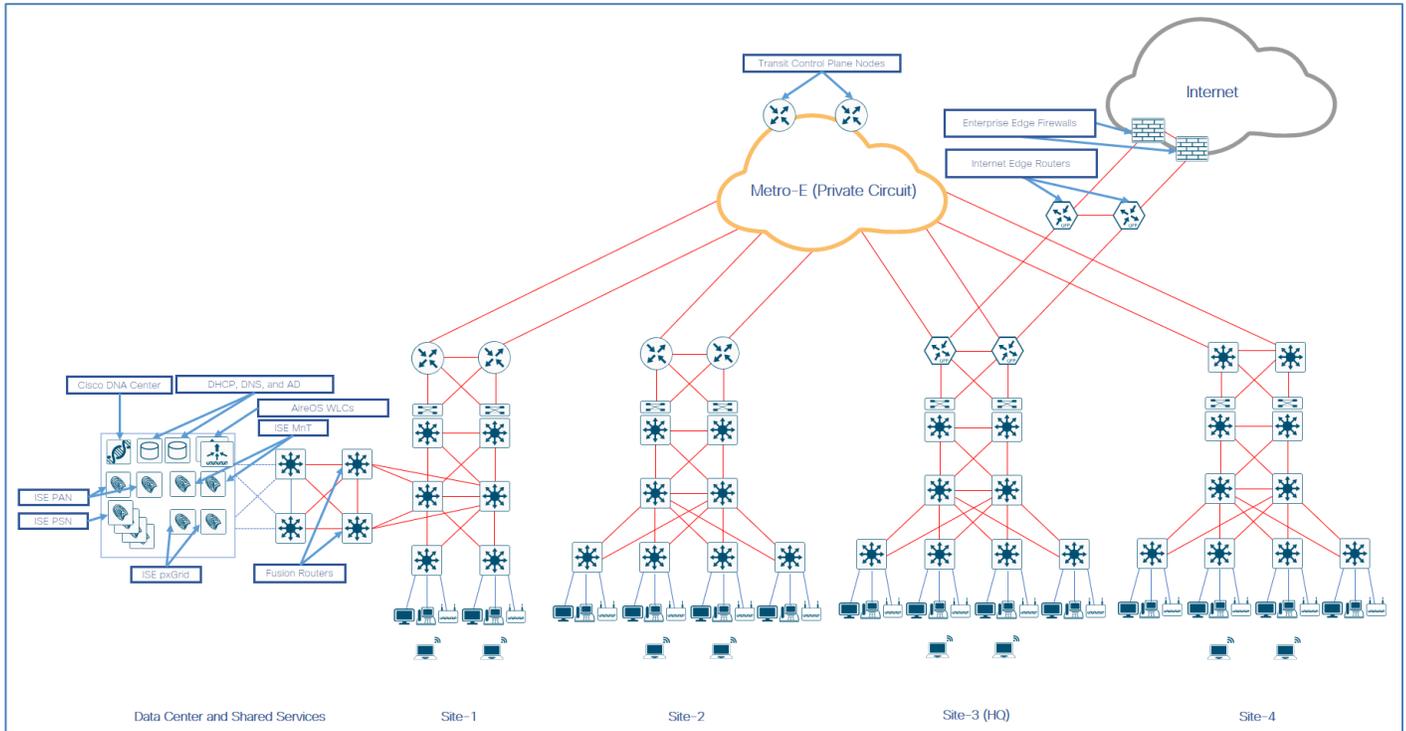
Direct Internet access (DIA) is not required at each site, although DIA is deployed depending on the design and deployment needs. DIA per site is more common in a WAN deployment than a MAN deployment and therefore the configuration has not been validated in this document.

The Enterprise edge block routers and switches for each site, which ultimately become fabric border nodes, are shown as directly connected to the Metro-E circuit as they are connected to the Service Provider equipment providing access to this circuit. The transit control plane nodes are shown located in the private circuit cloud. Like every other site, except for HQ, they are accessible only through the Metro-E circuit. Like the Enterprise edge block routers and switches at each site, the transit control plane nodes are connected to the Service Provider equipment accessing the circuit.

Within each site, the IS-IS routing protocol is used to advertise underlay routes and to provide IP reachability between Loopback 0 addresses. Through the Metro-E circuit, the Enterprise Edge routers and switches (fabric border nodes) at each site are IS-IS adjacent with each other so that there is full, end-to-end IP reachability between sites across the private circuit.

Identity Services Engine and Shared Services Overview

Figure 4 Overlay Connectivity – High Level Overview



The Identity Services Engine environment is configured as a [distributed deployment](#). There are multiple ISE virtual machines, each running an independent persona. The deployment consists of two Primary Administration Nodes (PAN), two Monitoring and Troubleshooting Nodes (MnT), and two Platform Exchange Grid (pxGrid) nodes with each pair set up as primary and secondary node. There is a Policy Services Node (PSN) dedicated to each fabric site in the network. The ISE nodes are all physically located in the Data Center site. This is not mandatory, and an optional deployment variation would include deploying the ISE PSNs physically at each site location.

Windows Active Directory (AD) is configured as an external identity store for the ISE deployment. AD is running on a Windows Server 2016 that is also providing DHCP and DNS services to the entire deployment.

To satisfy the [asymmetric packet path](#) necessary in SD-Access Wireless for Access Point to WLC registration, WLC to control plane node registration, and AP to DHCP server, the Data Center location has underlay reachability through a direct link to the Metro-E circuit along with overlay reachability to the deployment through a pair of fusion routers ultimately connected to a pair of Internal-Only fabric border nodes located in Site-1.

The Fusion Routers are a pair of Catalyst switches connected upstream to the distribution block at Site-1 and downstream to the data center switches which provide reachability to the UCS servers hosting the management virtual machines.

 Tech Tip

Fusion router technical requirements are discussed later in the document along with information on border node types.

Deploy

This section focuses on the complete deployment workflow and guidelines, from device discovery through to fabric automation.

Using the Cisco DNA Center GUI, the network is discovered and added to the inventory. Cisco DNA Center is then integrated with the Identity Services Engine through mutual certificate authentication creating a trusted relationship. The Design application is used to create site hierarchies, define common network settings, define IP address pools, configure WLAN settings, and set up the ISE guest portal.

The Policy application is used to create Scalable Group Tags (SGTs) and configure TrustSec related policies. This workflow is managed in Cisco DNA Center and pushed to ISE using the trust established during integration.

Finally, the Provision application is used to provision the configuration defined in the Design and Policy workflows to the managed devices discovered and present in the inventory. The Provision application is also used to create and define the SD-Access overlay network.

Process 1: Discovering the Network Infrastructure

Procedure 1: Discover Network Devices – Discovery Tool

Cisco DNA Center is used to discover and manage the SD-Access underlay network devices. To discover equipment in the network, the Appliance must have IP reachability to these devices, and CLI and SNMP management credentials must be configured on them. Once discovered, the devices are added to Cisco DNA Center’s inventory, allowing the controller to make configuration changes through provisioning.

The following steps show how to initiate a discovery job by supplying an IP address range or multiple ranges for scanning for network devices. IP address range discovery provides a small constraint to the discovery job which may save time. Alternatively, by providing the IP address of an initial device for discovery, Cisco DNA Center can use Cisco Discovery Protocol (CDP) to find neighbors connected to the initial device.



If using CDP for discovery, reduce the default number of hops to speed up the discovery job.

At a minimum, CLI and SNMP credentials must be provided to initiate a discovery job. Either SNMPv2c Read and Write credentials or SNMPv3 credentials are required. SNMPv3 is given priority if both are used.



Avoid using **admin** as the username for the network device’s CLI credentials. This can lead to username conflicts with the ISE administrator login, resulting in the inability to log into devices.

1. In the Cisco DNA Center dashboard, click **Discovery** in the **Tools** section.
2. In the New Discovery dialog, supply a **Discovery Name** (example: Site-1).
3. Select the **Range** radio button.
4. Enter an IP address for the start (**From**) and end (**To**) of the IP range (for example, 192.168.10.1, 192.168.10.8). For a single address, enter the same IP address for both the start and end of the range.
5. Select **Use Loopback** for the **Preferred Management IP**.

New Discovery

Discovery Name*
Site-1

▼ IP Address/Range *

Discovery Type ⓘ
 CDP Range LLDP

From* ⓘ To* ⓘ
 192.168.10.1 - 192.168.10.8 +

From* To*
 192.168.108.1 - 192.168.108.2 X

Preferred Management IP ⓘ
 Use Loopback ▼



If additional ranges are in the deployment, click **+**, enter the additional range, and repeat for any remaining ranges (example, 192.168.108.1, 192.168.108.2 used here to discover the WLCs).

6. Scroll down to the Credentials section.
7. Click **+ Add Credentials** on the far right.
The Add Credentials slide-in pane appears.

▼ Credentials *

ⓘ At least one CLI credential and one SNMP credential are required.

ⓘ Netconf is mandatory for enabling Wireless Services on eWLC and Wireless capable devices such as Cat9k.

■ global ■ task-specific

[+ Add Credentials](#)

8. Select the **CLI** tab.
9. Provide a **Username** and **Password** and then confirm the **Password**.
10. Provide a password in the **Enable Password** field and confirm it.
11. (Optional) Select the **Save as global settings** check box
12. Click **Save**.

Add Credentials

CLI SNMPv2c SNMPv3 SNMP PROPERTIES HTTP(S) NETCONF

Username*
Operator

Password*
●●●●●●●●

Confirm Password*
●●●●●●●●

Enable Password
●●●●●●

Confirm Enable Password
●●●●●●

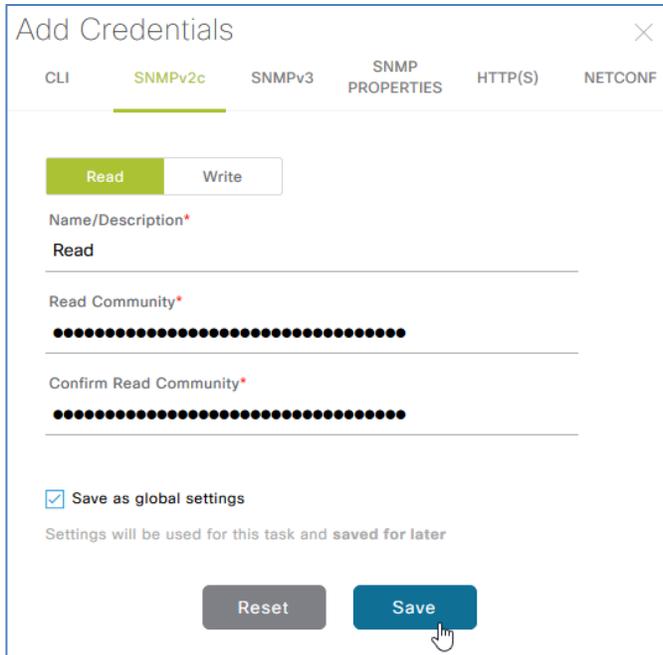
Save as global settings
Settings will be used for this task and saved for later

Reset Save

 Tech tip

The *Save as global settings* check box is selected if the devices in the network share common administrator credentials. When a new discovery job is created, the global credentials are automatically populated and are indicated with a purple border around the credential name. Up to five global credentials can be saved.

13. Select the relevant **SNMP** tab.
14. For **SNMPv2c**, click **Read**.
15. Enter a name/description.
16. Enter a **Read Community** string and confirm it.
17. (Optional) Select the **Save as global settings** check box.
18. Click **Save**.

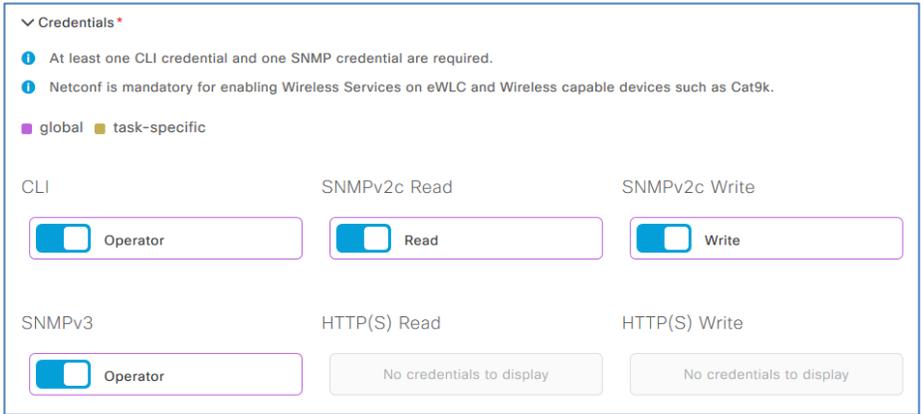


The screenshot shows the 'Add Credentials' dialog box with the 'SNMPv2c' tab selected. The 'Read' button is highlighted in green. The 'Name/Description' field contains the text 'Read'. The 'Read Community' and 'Confirm Read Community' fields are filled with a series of black dots. The 'Save as global settings' checkbox is checked. Below the checkbox, the text reads 'Settings will be used for this task and saved for later'. At the bottom, there are 'Reset' and 'Save' buttons, with a mouse cursor pointing at the 'Save' button.

19. For **SNMPv2c**, click **Write**.
20. Enter a name/description.
21. Enter a **Write Community** string and confirm it.
22. (Optional) Select the **Save as global settings** check box.
23. Click **Save**.



31. Verify that the defined Credentials are displayed in the Discovery job.



32. Scroll down and click **Advanced** to expand the section.

33. Under **Protocol Order**, ensure that **SSH** is selected and is above **Telnet** in the protocol order.

34. Click **Start**.

The discovery details are displayed while the Discovery runs.

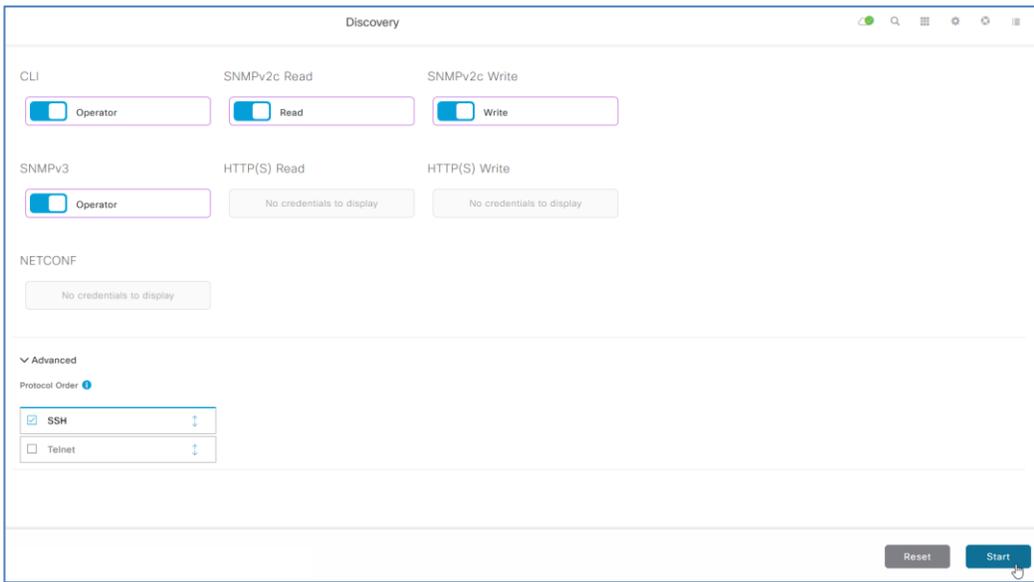


Figure 5 Site-1 Discovery Job Details

History ▾

Filter

IP Address	Device Name	Status	ICMP ▲	SNMP	CLI	HTTP(S)	NETCONF
192.168.108.2	WLC-8540-02	✓	✓	✓	✓	●	●
192.168.10.5	3850-24P-01	✓	✓	✓	✓	●	●
192.168.108.1	WLC-8540-01	✓	✓	✓	✓	●	●
192.168.10.4	6832-02	✓	✓	✓	✓	●	●
192.168.10.6	3850-24P-02	✓	✓	✓	✓	●	●
192.168.10.3	6832-01	✓	✓	✓	✓	●	●
192.168.10.2	C-6840-04	✓	✓	✓	✓	●	●
192.168.10.7	ISR-4451-07	✓	✓	✓	✓	●	●
192.168.10.8	ISR-4451-08	✓	✓	✓	✓	●	●
192.168.10.1	C-6840-03	✓	✓	✓	✓	●	●

Show ▾ Showing 1 to 10 of 10 Page ▾ of 1

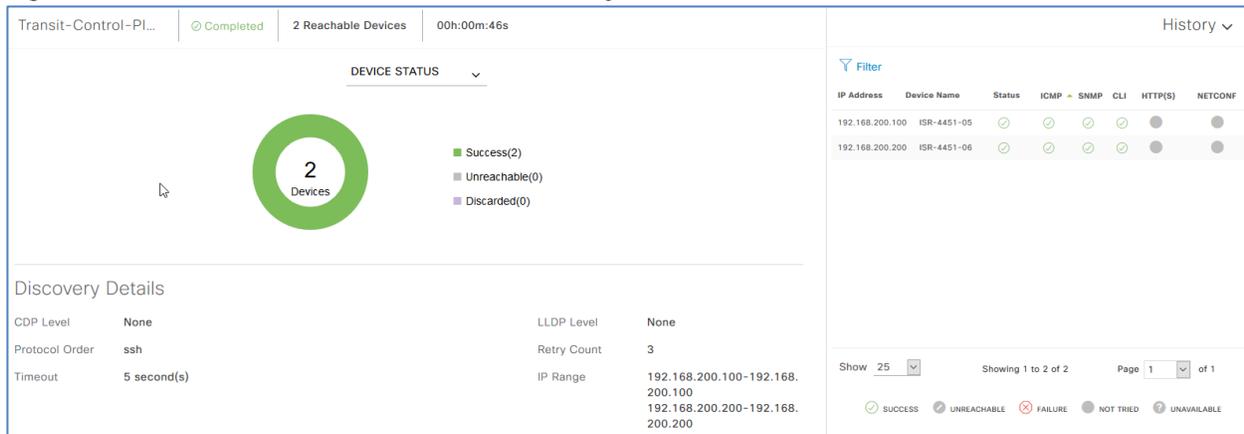
✓ SUCCESS ● UNREACHABLE ✗ FAILURE ● NOT TRIED ? UNAVAILABLE

Tech tip

SSH is the preferred discovery type over Telnet. Telnet is most commonly used in lab settings and not in production environments. If the environment uses Telnet, select the check box, and drag and drop Telnet higher in the **Protocol Order**. If both SSH and Telnet are selected, Cisco DNA Center attempts to connect to the discovered devices using both options in the order they are listed. Once Cisco DNA Center has a successful connection with a device, it will not attempt the next protocol type.

35. If there are any discovery failures, inspect the devices list, resolve the problem, and restart the discovery for those devices.
36. Repeat these steps to discover additional devices in the network and add them to the inventory.

Figure 6 Transit Control Plane Node Discovery Job



Tech tip

When using the Cisco Catalyst 6800 Series switch, particularly with a large configuration or many files in flash, Discovery timeouts may be avoided by adding the following command in configuration mode on the switch:

```
snmp mib flash cache
```

Procedure 2: Define the Network Device Role – Inventory Tool

The Cisco DNA Center **Inventory** tool is used to add, update, and/or delete devices that are managed by the Appliance. The **Inventory** tool can also be used to perform additional functions and actions, such as updating device credentials, resyncing devices, and exporting the device configurations. It also provides the ability to launch the **Command Runner** tool.

The Inventory tool has many uses and significant device information is collected and exported by exposing the various columns. However, the tool has a single use in this deployment guide – modifying the device role.

About Device Roles

The device role is used to position devices in the Cisco DNA Center topology maps under the **Fabric** tab in the Provision application and in the Topology tool. The device positions in these applications and tools are shown using the classic three-tiered Core, Distribution, and Access layout.

Changing a device role does not change the function of the device, nor does it impact the SD-Access configuration that is provisioned. In SD-Access, it is simply a GUI construct to help with network layout visualization shown in later steps.

Table 1 Device Roles

Device Role	Topology Position
Internet (Not Selectable)	Top Row
Border Router	Below Internet (displayed with lines connected to the Internet Cloud)
Core	Third Row
Distribution	Fourth Row
Access	Bottom Row
Unknown	To the side of the Bottom Row

To define the device role:

1. After the discovery process finishes successfully, navigate to the main Cisco DNA Center dashboard, and click **Inventory** in the **Tools** section. The discovered devices are displayed.
2. Verify that the **Last Sync Status** of the devices is **Managed** before making device role changes.

<input type="checkbox"/>	Device Name	IP Address ▲	Reachability Status	Uptime	Last Updated	Resync Interval	Last Sync Status
<input type="checkbox"/>	3850-24P-01	192.168.10.5	✔ Reachable	13 days 22 hrs 16 mins	7 minutes ago	01:00:00	Managed

3. To update the device role, click the blue pencil icon in the **Device Role** column, and select the required role.

<input type="checkbox"/>	C-6840-03	192.168.10.1	✔ Reachable	24 days 20 hrs 05 mins	an hour ago	02:00:00	Managed	CORE
<input type="checkbox"/>	C-6840-04	192.168.10.2	✔ Reachable	24 days 20 hrs 02 mins	an hour ago	02:00:00		CORE
<input type="checkbox"/>	ISR-4451-07	192.168.10.7	✔ Reachable	24 days 23 hrs 41 mins	an hour ago	02:00:00		BORDER ROUTER
<input type="checkbox"/>	ISR-4451-08	192.168.10.8	✔ Reachable	24 days 23 hrs 37 mins	2 hours ago	02:00:00		BORDER ROUTER

Update Device Role ✕

The following device roles are assigned to the devices discovered in Site-1 and are used to create the Fabric topology layout shown later in this document.

Figure 7 Device Roles - Inventory Tool – Site-1

Device Name	IP Address	Reachability Status	Uptime	Last Updated	Resync Interval	Last Sync Status	Device Role	Site
3850-24P-01	192.168.10.5	✔ Reachable	7 days 21 hrs 09 mins	19 minutes ago	00:25:00	Managed	ACCESS	Unassigned
3850-24P-02	192.168.10.6	✔ Reachable	32 days 20 hrs 07 mins	19 minutes ago	00:25:00	Managed	ACCESS	Unassigned
6832-01	192.168.10.3	✔ Reachable	32 days 23 hrs 54 mins	19 minutes ago	00:25:00	Managed	DISTRIBUTION	Unassigned
6832-02	192.168.10.4	✔ Reachable	32 days 23 hrs 54 mins	19 minutes ago	00:25:00	Managed	DISTRIBUTION	Unassigned
C-6840-03	192.168.10.1	✔ Reachable	32 days 18 hrs 43 mins	19 minutes ago	00:25:00	Managed	CORE	Unassigned
C-6840-04	192.168.10.2	✔ Reachable	32 days 18 hrs 42 mins	19 minutes ago	00:25:00	Managed	CORE	Unassigned
ISR-4451-07	192.168.10.7	✔ Reachable	32 days 23 hrs 04 mins	19 minutes ago	00:25:00	Managed	BORDER ROUTER	Unassigned
ISR-4451-08	192.168.10.8	✔ Reachable	32 days 23 hrs 02 mins	19 minutes ago	00:25:00	Managed	BORDER ROUTER	Unassigned

Process 2: Integrating Cisco DNA Center with the Identity Services Engine

The Identity Services Engine (ISE) is the authentication and policy server required for Software-Defined Access. Once integrated with Cisco DNA Center using pxGrid, information sharing between the two platforms is enabled, including device information and group information. This allows Cisco DNA Center to define policies that are pushed to ISE and then rendered into the network infrastructure by the ISE Policy Service Nodes (PSNs).

When integrating the two platforms, a trust is established through mutual certificate authentication. This authentication is completed seamlessly in the background during integration and requires both platforms to have accurate NTP sync.

Procedure 1: Configure Authentication and Policy Servers in Cisco DNA Center

The following steps outline how to integrate Cisco DNA Center and the Identity Services Engine. This integration must be done before completing the workflows in the Design and Policy applications so that the results of the workflows can be provisioned to the network equipment to use ISE as the AAA server for users and endpoints via RADIUS and for device administrators via TACACS+.



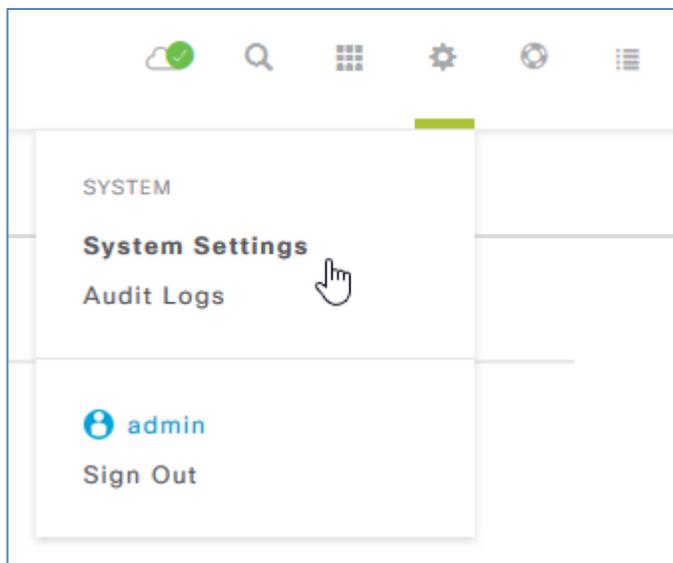
Many organizations use TACACS+ for management authentication to network devices. TACACS+ was used for this prescriptive guide, although it is not required. TACACS+ is not a mandatory requirement for Software-Defined Access. If TACACS+ is not used in the deployment, do not select the option in **Advanced Settings** when integrating Cisco DNA Center and ISE.

During the integration, all Scalable Group Tags (SGTs) present in ISE are pulled into Cisco DNA Center. Whatever policy is currently configured in the TrustSec matrices of ISE are also pulled into Cisco DNA Center. This is sometimes referred to as *Day 0 Brownfield Support*: if policies are present in ISE at the point of integration, those policies are populated in the Cisco DNA Center Policy application.

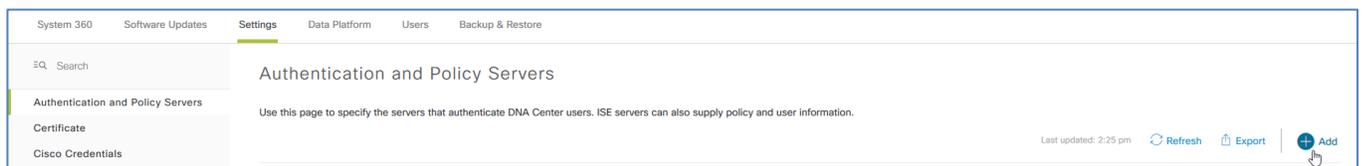


Except for SGTs, anything TrustSec and TrustSec Policy-related that is created directly on ISE OOB (out-of-band) from Cisco DNA Center after the initial integration is not available or displayed in DNA Center. Once the platforms are integrated, TrustSec must be managed from Cisco DNA Center, not from ISE.

1. From the main Cisco DNA Center dashboard select the  gear icon in the top-right corner.
2. Click System Settings.



3. Navigate to **Settings > Authentication and Policy Servers**.
4. Click the **+ Add** button.



5. In the **Add AAA/ISE** dialog, enter the IP Address of the ISE Primary Administration Node (PAN) in the **Server IP Address** field (example: 198.51.100.121).
6. Enter a **Shared Secret**.
7. Toggle the **Cisco ISE** selector to **On** to reveal additional settings required for integration with ISE.

Add AAA/ISE server

Server IP Address*
198.51.100.121

Shared Secret*
●●●●●●●●●●

Cisco ISE server
 On

8. Enter the ISE Super Admin **Username** and **Password**.
9. Enter the ISE fully qualified domain name for **FQDN** field.
10. Enter a meaningful **Subscriber Name** (example: **dnac**) as it will be displayed later in the ISE GUI.
11. Leave the **SSH Key** blank.
12. Click **View Advanced Settings** to expand it.

Add AAA/ISE server

Server IP Address*
198.51.100.121

Shared Secret*
●●●●●●●●●●

Cisco ISE server
 On

Username*
admin

Password*
●●●●●●●●

FQDN*
ISE-PAN-121.dna.local

Subscriber Name*
dnac

SSH Key

Virtual IP Address(es) ⓘ

View Advanced Settings

13. Select the **RADIUS** and **TACACS** protocols.
14. Click **Apply**.

Hide Advanced Settings

Protocol

RADIUS TACACS

Authentication Port*

1812

Accounting Port*

1813

Port*

49

Retries*

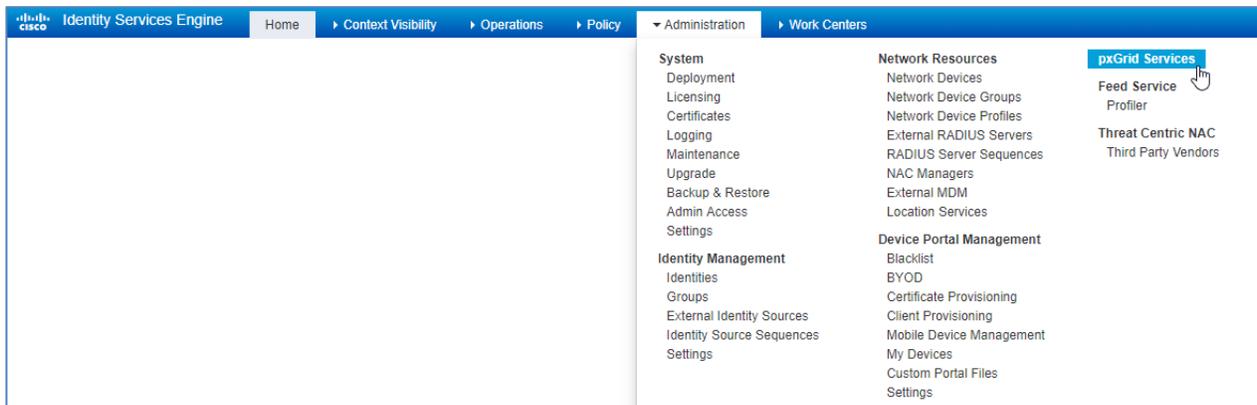
1

Timeout (seconds)*

2

Cancel Apply

15. During the establishment of communication, Cisco DNA Center displays **Creating AAA server...**
Verify that the **Status** displays as **INPROGRESS**.
16. Log in to the **ISE GUI**.
17. Navigate to **Administration > pxGrid Services**.



18. Please wait at minimum of three minutes.
19. The client named **dnac** (the value entered as the Subscriber name) is now showing **Pending** in the **Status** column.

Cisco Identity Services Engine							
Home		Context Visibility	Operations	Policy	Administration	Work Centers	
System		Identity Management	Network Resources	Device Portal Management	pxGrid Services	Feed Service	Threat Centric NAC
All Clients Web Clients Capabilities Live Log Settings Certificates							
<input checked="" type="checkbox"/> Enable <input type="checkbox"/> Disable <input checked="" type="checkbox"/> Approve <input type="checkbox"/> Group <input type="checkbox"/> Decline <input type="checkbox"/> Delete <input checked="" type="checkbox"/> Refresh Total Pending Approval(1)							
<input type="checkbox"/>	Client Name	Client Description	Capabilities	Status	Client Group(s)	Auth Method	Log
<input type="checkbox"/>	ise-admin-ise-mnt-124		Capabilities(1 Pub, 1 Sub)	Online (XMPP)	Administrator	Certificate	View
<input type="checkbox"/>	ise-admin-ise-pxgrid-125		Capabilities(1 Pub, 1 Sub)	Online (XMPP)	Administrator	Certificate	View
<input type="checkbox"/>	ise-mnt-ise-mnt-123		Capabilities(2 Pub, 1 Sub)	Online (XMPP)	Administrator	Certificate	View
<input type="checkbox"/>	ise-mnt-ise-mnt-124		Capabilities(2 Pub, 1 Sub)	Online (XMPP)	Administrator	Certificate	View
<input type="checkbox"/>	ise-admin-ise-psn-127		Capabilities(0 Pub, 1 Sub)	Online (XMPP)	Administrator	Certificate	View
<input type="checkbox"/>	ise-pubsub-ise-pxgrid-125		Capabilities(0 Pub, 0 Sub)	Online (XMPP)		Certificate	View
<input type="checkbox"/>	ise-fanout-ise-pxgrid-125		Capabilities(0 Pub, 0 Sub)	Online (XMPP)	Administrator	Certificate	View
<input type="checkbox"/>	ise-pubsub-ise-pxgrid-126		Capabilities(0 Pub, 0 Sub)	Online (XMPP)		Certificate	View
<input type="checkbox"/>	ise-admin-ise-psn-128		Capabilities(0 Pub, 1 Sub)	Online (XMPP)	Administrator	Certificate	View
<input type="checkbox"/>	ise-admin-ise-pan-121		Capabilities(2 Pub, 2 Sub)	Online (XMPP)	Administrator	Certificate	View
<input type="checkbox"/>	ise-bridge-ise-pxgrid-126		Capabilities(0 Pub, 5 Sub)	Online (XMPP)	Administrator	Certificate	View
<input type="checkbox"/>	ise-admin-ise-pxgrid-126		Capabilities(1 Pub, 1 Sub)	Online (XMPP)	Administrator	Certificate	View
<input type="checkbox"/>	ise-admin-ise-pan-122		Capabilities(0 Pub, 1 Sub)	Online (XMPP)	Administrator	Certificate	View
<input type="checkbox"/>	ise-fanout-ise-pxgrid-126		Capabilities(0 Pub, 0 Sub)	Online (XMPP)	Administrator	Certificate	View
<input type="checkbox"/>	ise-admin-ise-mnt-123		Capabilities(2 Pub, 1 Sub)	Online (XMPP)	Administrator	Certificate	View
<input type="checkbox"/>	dnac		Capabilities(0 Pub, 0 Sub)	Pending	Session	Certificate	View

- If the Subscriber has not appeared, please use the green refresh button. Integration may take as long as ten minutes although it commonly takes less than five minutes.
- Select the checkbox beside the subscriber (**dnac**).
- In the **Total Pending Approval** drop-down, select **Approve All**.

Cisco Identity Services Engine							
Home		Context Visibility	Operations	Policy	Administration	Work Centers	
System		Identity Management	Network Resources	Device Portal Management	pxGrid Services	Feed Service	Threat Centric NAC
All Clients Web Clients Capabilities Live Log Settings Certificates							
<input checked="" type="checkbox"/> Enable <input type="checkbox"/> Disable <input checked="" type="checkbox"/> Approve <input type="checkbox"/> Group <input type="checkbox"/> Decline <input type="checkbox"/> Delete <input checked="" type="checkbox"/> Refresh Total Pending Approval(1)							
<input type="checkbox"/>	Client Name	Client Description	Capabilities	Status	Client Group(s)	Auth Method	Log
<input type="checkbox"/>	ise-mnt-ise		Capabilities(2 Pub, 1 Sub)	Online (XMPP)		Certificate	View

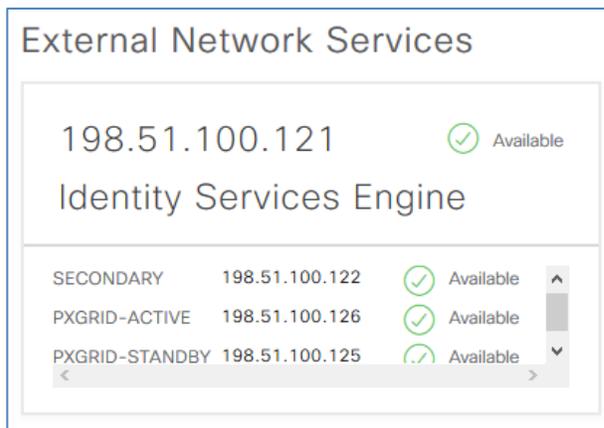
- Click the second **Approve All** to double confirm.

i Are you sure you want to Approve clients(listed below) that are currently pending approval?

- A success message displays. Verify that the subscriber status changes to **Online (XMPP)**.

Cisco Identity Services Engine							
Home		Context Visibility	Operations	Policy	Administration	Work Centers	
System		Identity Management	Network Resources	Device Portal Management	pxGrid Services	Feed Service	Threat Centric NAC
All Clients Web Clients Capabilities Live Log Settings Certificates Permissions							
<input checked="" type="checkbox"/> Enable <input type="checkbox"/> Disable <input checked="" type="checkbox"/> Approve <input type="checkbox"/> Group <input type="checkbox"/> Decline <input type="checkbox"/> Delete <input checked="" type="checkbox"/> Refresh Total Pending Approval(0)							
<input type="checkbox"/>	Client Name	Client Description	Capabilities	Status	Client Group(s)	Auth Method	Log
<input type="checkbox"/>	ise-admin-ise-psn-127		Capabilities(0 Pub, 1 Sub)	Online (XMPP)	Internal	Certificate	View
<input type="checkbox"/>	ise-admin-ise-mnt-124		Capabilities(1 Pub, 1 Sub)	Online (XMPP)	Internal	Certificate	View
<input type="checkbox"/>	ise-admin-ise-psn-129		Capabilities(0 Pub, 1 Sub)	Online (XMPP)	Internal	Certificate	View
<input type="checkbox"/>	ise-mnt-ise-mnt-124		Capabilities(2 Pub, 1 Sub)	Online (XMPP)	Internal	Certificate	View
<input type="checkbox"/>	ise-admin-ise-psn-130		Capabilities(0 Pub, 1 Sub)	Online (XMPP)	Internal	Certificate	View
<input type="checkbox"/>	ise-bridge-ise-pxgrid-126		Capabilities(0 Pub, 4 Sub)	Online (XMPP)	Internal	Certificate	View
<input type="checkbox"/>	ise-mnt-ise-mnt-123		Capabilities(2 Pub, 1 Sub)	Online (XMPP)	Internal	Certificate	View
<input type="checkbox"/>	ise-fanout-ise-pxgrid-125		Capabilities(0 Pub, 0 Sub)	Online (XMPP)	Internal	Certificate	View
<input type="checkbox"/>	ise-fanout-ise-pxgrid-126		Capabilities(0 Pub, 0 Sub)	Online (XMPP)	Internal	Certificate	View
<input type="checkbox"/>	ise-admin-ise-psn-128		Capabilities(0 Pub, 1 Sub)	Online (XMPP)	Internal	Certificate	View
<input type="checkbox"/>	ise-pubsub-ise-pxgrid-125		Capabilities(0 Pub, 0 Sub)	Online (XMPP)	Internal	Certificate	View
<input type="checkbox"/>	ise-admin-ise-pxgrid-125		Capabilities(0 Pub, 1 Sub)	Online (XMPP)	Internal	Certificate	View
<input type="checkbox"/>	ise-pubsub-ise-pxgrid-126		Capabilities(0 Pub, 0 Sub)	Online (XMPP)	Internal	Certificate	View
<input type="checkbox"/>	ise-admin-ise-pxgrid-126		Capabilities(0 Pub, 1 Sub)	Online (XMPP)	Internal	Certificate	View
<input type="checkbox"/>	ise-admin-ise-pan-122		Capabilities(0 Pub, 1 Sub)	Online (XMPP)	Internal	Certificate	View
<input type="checkbox"/>	ise-admin-ise-pan-121		Capabilities(2 Pub, 2 Sub)	Online (XMPP)	Internal	Certificate	View
<input type="checkbox"/>	ise-admin-ise-mnt-123		Capabilities(2 Pub, 1 Sub)	Online (XMPP)	Internal	Certificate	View
<input checked="" type="checkbox"/>	dnac		Capabilities(0 Pub, 2 Sub)	Online (XMPP)	Session	Certificate	View

Once established, the current communication status between ISE and Cisco DNA Center can be viewed by navigating from the gear icon to **System Settings > System 360**. Under **External Network Services**, the Cisco ISE server shows an **Available** status.



Process 3: Using the Design Application

Cisco DNA Center provides a robust Design application to allow customers of varying sizes and scales to easily define their physical sites and common resources. Using an intuitive hierarchical format, the Design application removes the need to redefine the same resource – such as DHCP, DNS, and AAA servers – in multiple places when provisioning devices. The network hierarchy created in the Design application should reflect the actual physical network hierarchy of the deployment.

The Design application is the building block for every other workflow in both Cisco DNA Center and Software-Defined Access. The configuration and items defined in this section are used and provisioned in later steps. The Design application begins with the creation of the network hierarchy of Areas, Buildings, and Floors. Once the hierarchy has been created, network settings are defined. These include DHCP and DNS servers, AAA servers and NTP servers, and when applicable, SNMP, Syslog, and Netflow servers. These servers are defined once in the Design application and provisioned to multiple devices in later steps. This allows for faster innovation without the repetitive typing of the same server configuration on the network infrastructure. After network settings are defined, IP address pools are designed, defined, and reserved. These pools are used for automation features such as border handoff and LAN automation and are used in host onboarding in later steps. The final step in the Design application in this guide is the creation and configuration of the wireless network settings including a Guest Portal in ISE.

Procedure 1: Create the Network Hierarchy

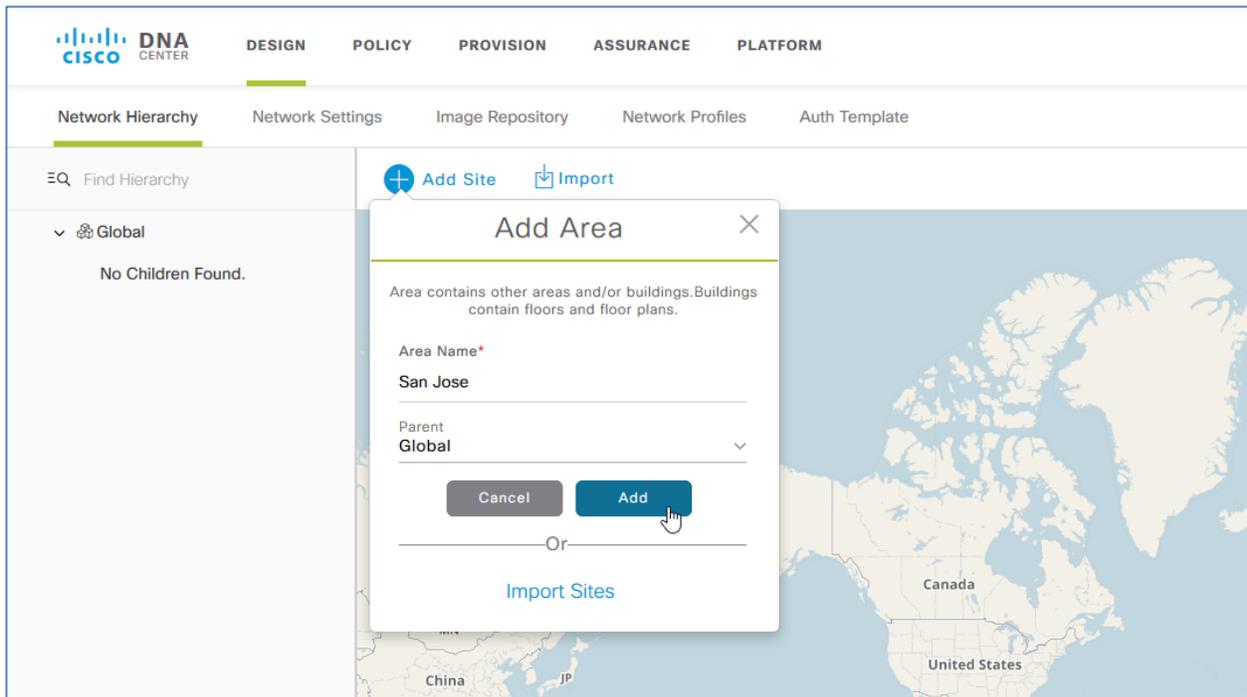
Using Cisco DNA Center, create a network hierarchy of areas, buildings, and floors that reflect the physical deployment. In later steps, discovered devices are assigned to buildings so that they are displayed hierarchically in the topology maps using the device role defined earlier in the Inventory tool.

Areas are created first. Within an Area, sub-areas and buildings are created. To create a building, the street address must be known to determine the coordinates and thus place the building on the map. Alternatively, use latitude and longitude coordinates without the street address. Floors are associated with buildings and support the importation of floor maps. For SD-Access Wireless, floors are mandatory as this is where Access Points are assigned, as clarified in later steps. Buildings created in these steps each represent a fabric site in the later Provisioning application procedures.

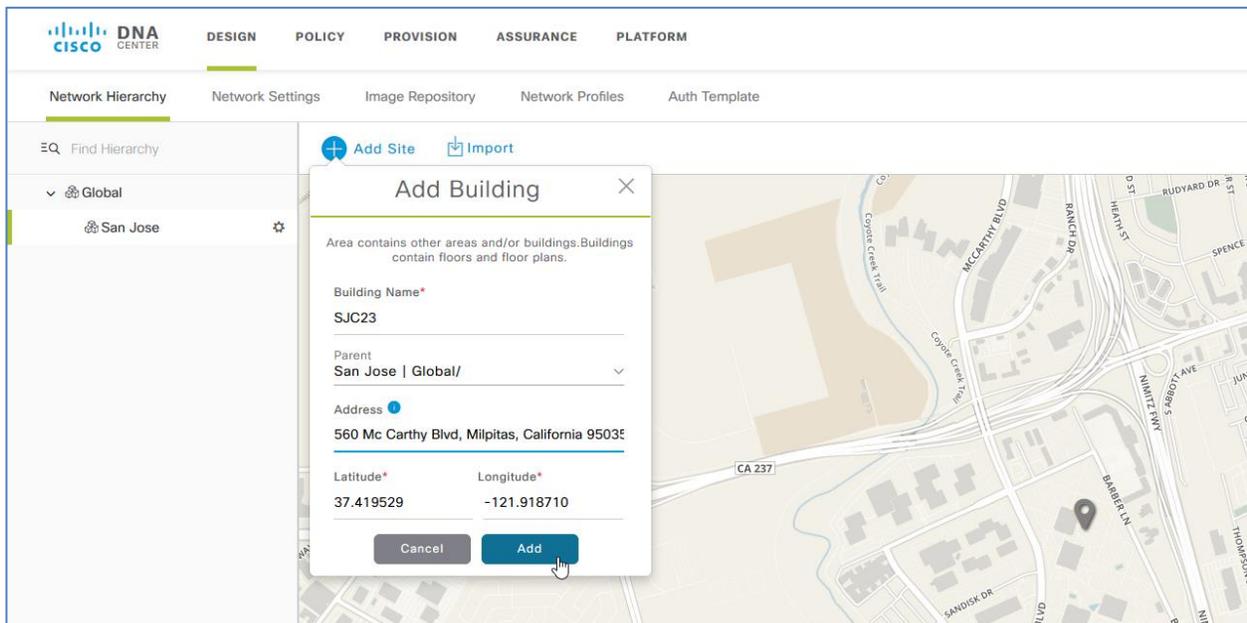
To create the network hierarchy:

1. From the main Cisco DNA Center dashboard, click the **DESIGN** tab.
2. Select the **Network Hierarchy** tab.
3. Click **+Add Site** to begin the network hierarchy design.
4. Select **Add Area** from the drop-down list.
5. Supply an appropriate **Area Name** (example: San Jose).

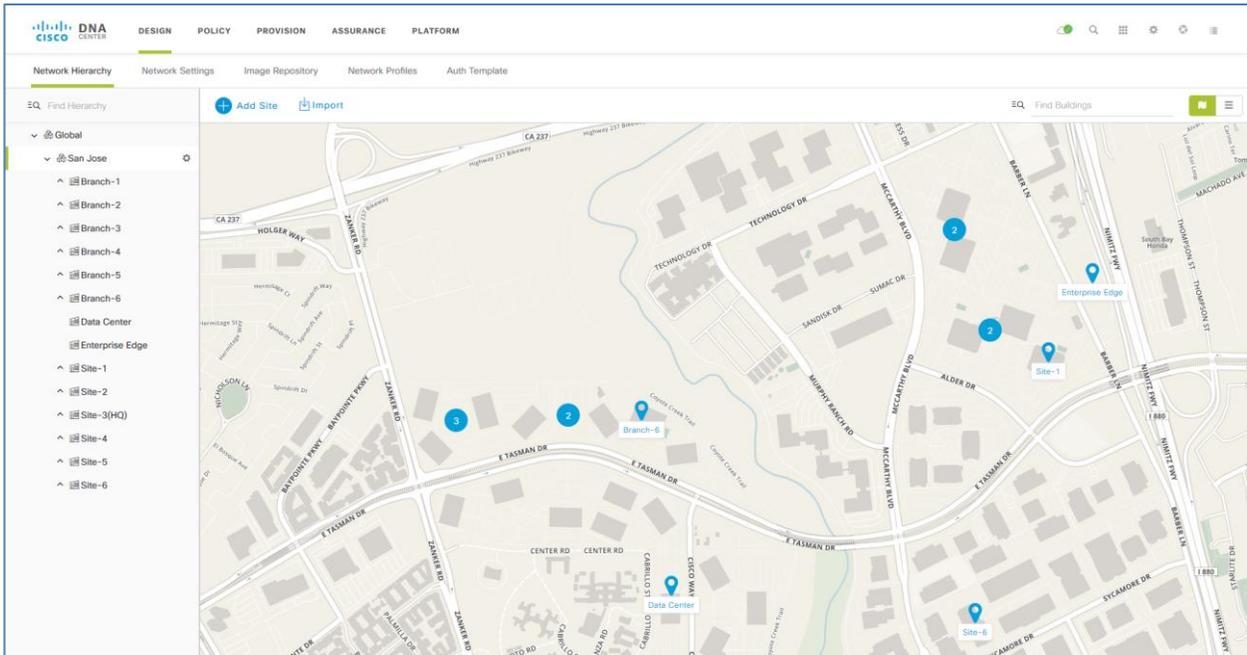
6. Click **Add**.



7. Click **+Add Site**.
8. Select **Add Building** from the drop-down list.
9. Provide an appropriate Building Name (example: SJC23).
10. Select the area created in the previous step as the **Parent**.
11. Complete the **Address** or **Latitude** and **Longitude** information to assign a location.
12. Click **Add**.



13. Repeat the previous steps to add sites and buildings, creating a hierarchy that reflects the physical network hierarchy and topology for the organization.



- To support SD-Access Wireless, select the gear icon next to a building in the hierarchy, choose **Add Floor**, and complete the wizard with the applicable details.

Add Floor ✕

Floor Name*
Site1-Floor1

Parent
Site-1

Type (RF Model)
Cubes And Walled Offices

Floor Image
Drag floor plan here
or
Upload file
SJ-23_1.dxf

(Supported formats DXF, DWG, JPG, GIF, PNG)

Width (ft)
 Length (ft)
 Height (ft)

100 60 10

Cancel
Add



Floor map diagrams are supported in the following formats: .DXF, .DWG, .JPG, .GIF, or .PNG.

The RF Model Type is related to Access Point heat maps. If required, the RF Model type can be modified after floor creation.

Procedure 2: Define Network Settings and Services

In Cisco DNA Center, common network resources and settings are saved in the Design application's Network Settings tab. Saving allows information pertaining to the enterprise to be stored so it can be reused throughout various Cisco DNA Center workflows. Items are defined once so that they can be used many times.

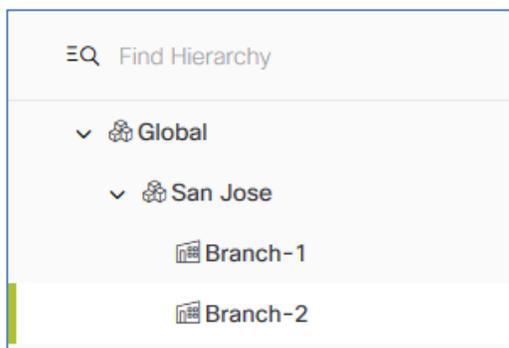
Configurable network settings in the Design application include AAA server, DHCP server, DNS server, Syslog server, SNMP server, Netflow collector, NTP server, time zone, and Message of the Day. Several of these items are applicable to a Cisco DNA Center Assurance deployment. For SD-Access, AAA and DHCP servers are mandatory and DNS and NTP servers should always be used.

By default, when clicking the Network Settings tab, newly configured settings are assigned as Global network settings. They are applied to the entire hierarchy and inherited by each site, building, and floor. In Network Settings, the default selection point in the hierarchy is **Global**.

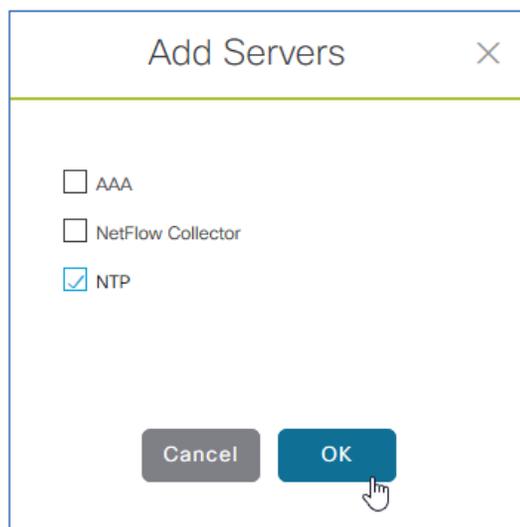
It is possible to define specific network settings and resources for specific sites. In fact, each fabric site in this deployment has its own dedicated ISE Policy Service Node, as shown in the next steps. For this prescriptive deployment, NTP, DHCP, and DNS have the same configuration for each site and are therefore defined at the Global level.



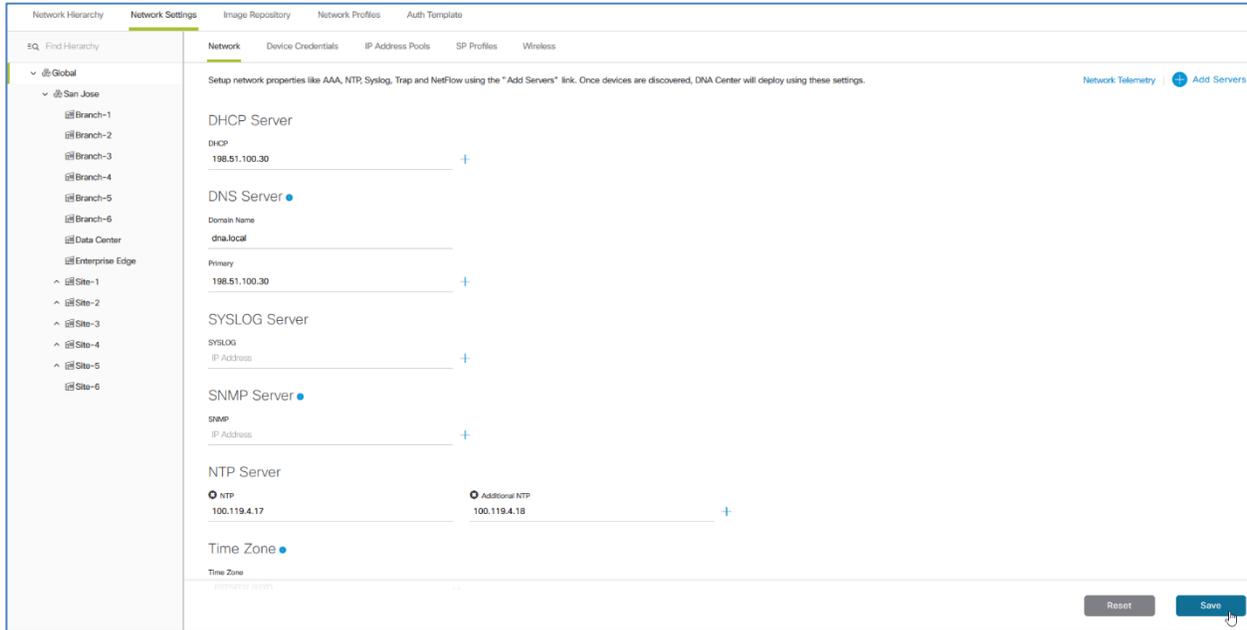
The vertical green bar in the site hierarchy indicates the currently selected area, building, or floor.



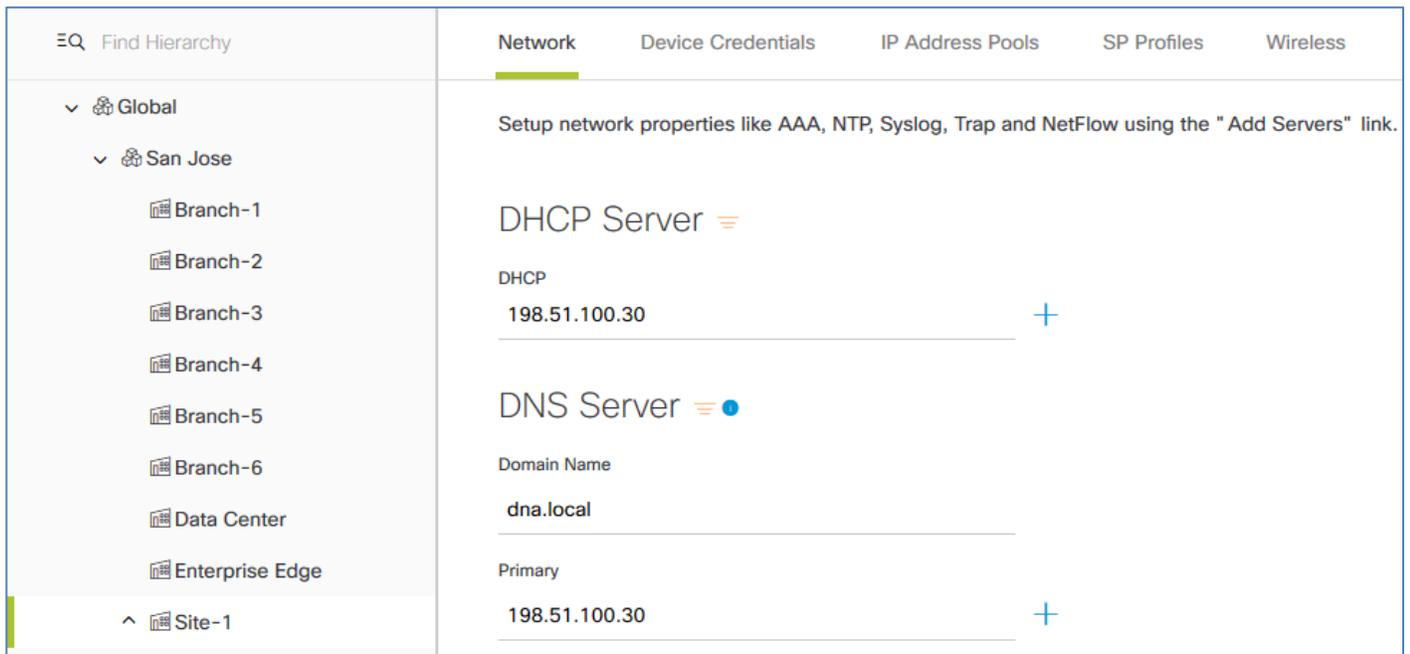
1. Navigate to **DESIGN > Network Settings > Network**.
2. Near the top, next to **Network Telemetry**, click **+ Add Servers**.
3. Select the **NTP** check box.
4. Click **OK**.



5. Within the left pane in the site hierarchy, select **Global**.
6. Enter the DHCP server IP address (example: 198.51.100.30).
7. Under **DNS Server**, enter the domain name (example: **dna.local**) and the IP address of the **Primary** DNS server (example: 198.51.100.30).
8. Under **NTP Server**, enter the NTP Server IP address (example: 100.119.4.17).
9. Click the **+** button under **Additional NTP** to add an additional NTP field.
10. Enter a second NTP Server IP address (example: 100.119.4.18).
11. Click **Save**.

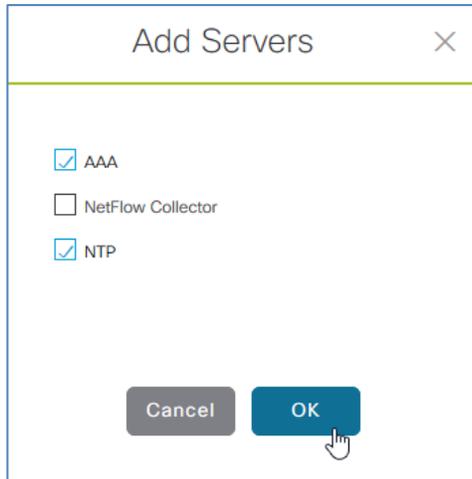


12. Navigate to a building (example: Site-1) in the hierarchy using the pane on the left. The DHCP, DNS, and NTP servers previously defined under **Global** are inherited  at the site level.



13. Near the top, next to **Network Telemetry**, click **+ Add Servers**.
14. Select the **AAA** check box and verify that the NTP check box is already selected.

15. Click **OK**.

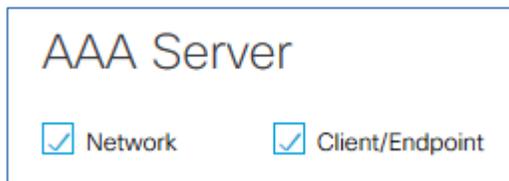


16. Verify that the configuration pane is updated with **AAA Server**.



For this prescriptive validation, ISE AAA services are used for both the network infrastructure device management and the client endpoints connecting to the infrastructure. The distributed ISE PSNs are defined to use RADIUS for the client endpoint authentication and TACACS+ for the network infrastructure.

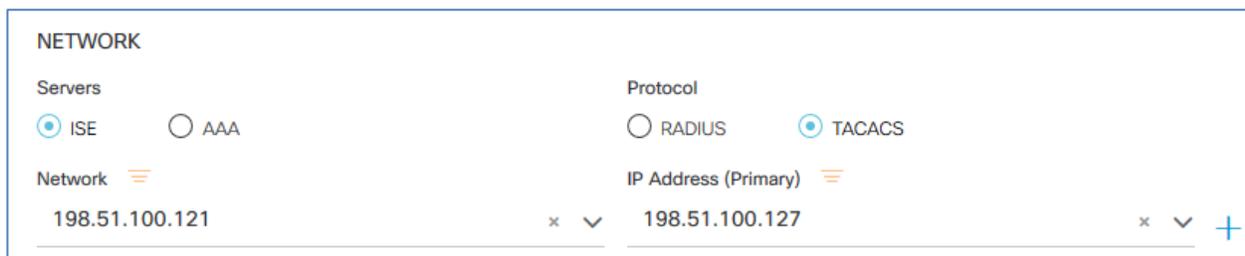
17. Under **AAA Server** select the **Network** and **Client/Endpoint** check boxes.



18. Under **NETWORK > Servers**, select the **ISE** radio button.
19. In the **Network** drop-down, select the prepopulated ISE server.
This is the IP address of the primary PAN.
20. Under **Protocol**, select the **TACACS** radio button.
21. In the **IP Address (Primary)** drop-down, select the IP address of the ISE PSN for the site.



If the site has a secondary PSN, click the **+** button, and then in the **IP Address (Additional)** dropdown, select the redundant ISE PSN.



22. Under **CLIENT/ENDPOINT > Servers**, select the **ISE** radio button.
23. In the **Client/Endpoint** drop-down, select the prepopulated ISE server.
This is the IP address of the Primary PAN.
24. Under **Protocol**, ensure that the **RADIUS** radio button is selected.

25. In the **IP Address (Primary)** drop-down, select the IP address of the ISE PSN for the site.



If the site has a secondary PSN, click the **+** button, and then in the **IP Address (Additional)** drop-down, select the redundant ISE PSN.

CLIENT/ENDPOINT

Servers ISE AAA

Protocol RADIUS TACACS

Client/Endpoint

IP Address (Primary)

26. Click **Save**.

The ISE servers for AAA and the servers for DHCP, DNS, and NTP for the selected level in the hierarchy are saved for later provisioning steps.

AAA Server

Network Client/Endpoint

NETWORK

Servers ISE AAA

Protocol RADIUS TACACS

Client/Endpoint

IP Address (Primary)

Save

27. Repeat the steps above for each site in the Distributed Campus deployment.

Figure 8 Site-2 Network Settings

DHCP Server

Network Client/Endpoint

NETWORK

Servers ISE AAA

Protocol RADIUS TACACS

Client/Endpoint

IP Address (Primary)

Save

Procedure 3: Create a Global IP Address Pool

This section provides information about global IP address pools and shows how to define the global IP address pools that are referenced during the pool reservation process.

About Global IP Address Pools

The Cisco DNA Center IPAM (IP Address Management) tool is used to create and reserve IP address pools for LAN and Border Automation, ENCS (NFV) workflows, and for the binding of the subnet segment to a VN in host onboarding. IP address pools are defined at the global level and then reserved at the area, building, or floor level. Reserving an IP address pool tells Cisco DNA Center to set aside that block of addresses for one of the special uses listed above.



IP address pool reservation is not available at the global level. It must be done at the area, building, or floor level.

IP address pools that will be used for DHCP must be manually defined and configured on the DHCP server. Cisco DNA Center does not provision the actual DHCP server, even if it is a Cisco device. It reserves pools as a visual reference for use in later workflows. DHCP scopes on the DHCP server should be configured with any additional DHCP options required to make a device work. For example, Option 150 is used to direct an IP phone to a TFTP server to receive their configuration and Option 43 is commonly used for Access Points to direct them to their corresponding wireless LAN controller.

The DHCP and DNS servers are the same for every site in this deployment. Because of the number of pools and sites, not every pool is listed in the chart below. The full IP address pool list can be found in [Appendix C](#).

Table 2 IP Address Pools – Site-1 through Site-4

Location	Usage at Site	Pool Name	Network/Mask	IP Gateway	DHCP Server	DNS Server
Global	Global Pool	ONE-SEVEN-TWO	172.16.0.0/12	172.16.0.1	198.51.100.30	198.51.100.30
Site-1	Access Points	AP-SITE-01	172.16.110.0/24	172.16.110.1	198.51.100.30	198.51.100.30
Site-1	BGP Handoff	BGP-SITE-01	172.16.111.0/24	172.16.111.1	-	-
Site-1	CAMPUS VN	CORP-SITE-01	172.16.112.0/24	172.16.112.1	198.51.100.30	198.51.100.30
Site-1	GUEST VN	GUEST-SITE-01	172.16.113.0/24	172.16.113.1	198.51.100.30	198.51.100.30
Site-1	RESEARCH VN	RESEARCH-SITE-01	172.16.114.0/24	172.16.114.1	198.51.100.30	198.51.100.30
Site-2	Access Points	AP-SITE-02	172.16.120.0/24	172.16.120.1	198.51.100.30	198.51.100.30
Site-2	-	BGP-SITE-02	172.16.121.0/24	172.16.121.1	-	-
Site-2	CAMPUS VN	CORP-SITE-02	172.16.122.0/24	172.16.122.1	198.51.100.30	198.51.100.30
Site-2	GUEST VN	GUEST-SITE-02	172.16.123.0/24	172.16.123.1	198.51.100.30	198.51.100.30
Site-2	RESEARCH VN	RESEARCH-SITE-02	172.16.124.0/24	172.16.124.1	198.51.100.30	198.51.100.30
Site-3	Access Points	AP-SITE-03	172.16.130.0/24	172.16.130.1	198.51.100.30	198.51.100.30
Site-3	BGP Handoff	BGP-SITE-03	172.16.131.0/24	172.16.131.1	-	-
Site-3	CAMPUS VN	CORP-SITE-03	172.16.132.0/24	172.16.132.1	198.51.100.30	198.51.100.30
Site-3	GUEST VN	GUEST-SITE-03	172.16.133.0/24	172.16.133.1	198.51.100.30	198.51.100.30
Site-3	RESEARCH VN	RESEARCH-SITE-03	172.16.134.0/24	172.16.134.1	198.51.100.30	198.51.100.30
Site-4	Access Points	AP-SITE-04	172.16.140.0/24	172.16.140.1	198.51.100.30	198.51.100.30
Site-4	-	BGP-SITE-04	172.16.141.0/24	172.16.141.1	-	-
Site-4	CAMPUS VN	CORP-SITE-04	172.16.142.0/24	172.16.142.1	198.51.100.30	198.51.100.30
Site-4	GUEST VN	GUEST-SITE-04	172.16.143.0/24	172.16.143.1	198.51.100.30	198.51.100.30
Site-4	RESEARCH VN	RESEARCH-SITE-04	172.16.144.0/24	172.16.144.1	198.51.100.30	198.51.100.30

IP address pools must be created at the global level first to be reserved later. Because the deployment is not integrating with a third-party IPAM tool where pools may already be configured, a single, large aggregate block is created at the Global level.

Location	Usage at Site	Pool Name	Network/Mask	IP Gateway	DHCP Server	DNS Server
Global	Global Pool	ONE-SEVEN-TWO	172.16.0.0/12	172.16.0.1	198.51.100.30	198.51.100.30

To create a global IP address pool:

1. Navigate to **DESIGN > Network Settings > IP Address Pools**.
2. In the site hierarchy on the left, select **Global**, and click **+ Add IP Pool**.
3. Enter the **IP Pool Name** (example: ONE-SEVEN-TWO).
4. Enter the **IP Subnet** (example: 172.16.0.0).
5. Use the drop-down to select the applicable **CIDR Prefix** (example: /12).
6. Enter the **Gateway IP Address** (example 172.16.0.1).
7. **Do not** select **Overlapping**.
8. Click **Save**.



Because an aggregate IP address pool was defined at the global level, there is increased flexibility and ease in reserving pools. It requires precise subnet planning and may not be available for all deployments – particularly if integrating with a third-party IPAM tool that has API integration with Cisco DNA Center. If an aggregate subnet is not available, then the same IP address pools need to be defined at the global level and then reserved at the applicable site level.

Procedure 4: Reserve IP Address Pools

The defined global IP address pool is used to reserve pools at the area, building, or floor in the network hierarchy. For single-site deployments, the entire set of global IP address pools is reserved for that site. In an SD-Access for Distributed Campus deployment, each site has its own assigned subnets that do not overlap between the sites.

To reserve IP address pools at the site level:

1. In the Cisco DNA Center dashboard, navigate to **DESIGN > Network Settings > IP Address Pools**.
2. In the site hierarchy on the left, select an area, building, or floor to begin the IP address pool reservation (example: Site-1).
3. In the top right, click **+ Reserve IP Pool** to display the **Reserve IP Pool** wizard.
4. Enter the **IP Pool Name** (example: AP-SITE-01).

5. In the **Type** drop-down, select **LAN**.
6. In the **Global IP Pool** drop-down, select a global pool (example: ONE-SEVEN-TWO).
7. Under **CIDR Notation / No. of IP Addresses** select the portion of the address space to use (example: 172.16.110.0 and /24 (255.255.255.0)).
8. Define the **Gateway IP Address** (example: 172.16.110.1).
9. Use the drop-downs to select the previously defined **DHCP Server(s)** and **DNS Servers(s)**.
10. Click **Reserve**.

Reserve IP Pool

IP Pool Name *
AP-SITE-01

Type
LAN

Global IP Pool *
ONE-SEVEN-TWO (172.16.0.0/12)

CIDR Notation / No. of IP Addresses *
172.16.110.0 /24 (255.255.255.0) or No. of IP Addresses

Gateway IP Address
172.16.110.1

DHCP Server(s)
x 198.51.100.30

DNS Server(s)
x 198.51.100.30

Overlapping

Cancel Reserve

11. Repeat the previous steps for all address blocks required to be reserved in the hierarchy for *each* site in the network.

The full reservations for Site-1 and Site-3 are shown below.

Figure 9 IP Address Pool Reservations – Site-1

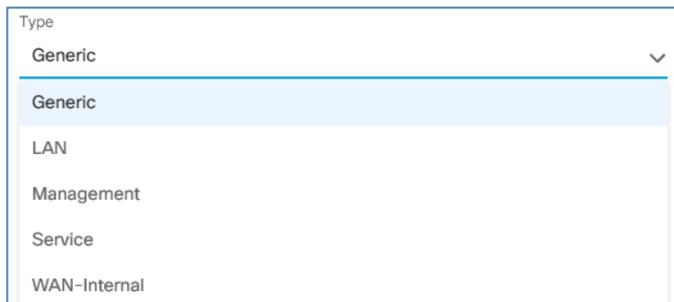
Name	IP Subnet Mask	Type	Global IP Pool	Gateway	DHCP Server	DNS Server	Free Count	Inherited fr...	Actions
AP-SITE-01	172.16.110.0/24	LAN	ONE-SEVEN-TWO (17...	172.16.110.1	198.51.100.30	198.51.100.30	256 of 256		Edit Release
BGP-SITE-01	172.16.111.0/24	LAN	ONE-SEVEN-TWO (17...	172.16.111.1			256 of 256		Edit Release
CORP-SITE-01	172.16.112.0/24	LAN	ONE-SEVEN-TWO (17...	172.16.112.1	198.51.100.30	198.51.100.30	0 of 256		Edit Release
GUEST-SITE-01	172.16.113.0/24	LAN	ONE-SEVEN-TWO (17...	172.16.113.1	198.51.100.30	198.51.100.30	256 of 256		Edit Release
RESEARCH-SITE-01	172.16.114.0/24	LAN	ONE-SEVEN-TWO (17...	172.16.114.1	198.51.100.30	198.51.100.30	256 of 256		Edit Release

Figure 10 IP Address Pool Reservations – Site-3

Name	IP Subnet Mask	Type	Global IP Pool	Gateway	DHCP Server	DNS Server	Free Count	Inherited fr...	Actions
AP-SITE-03	172.16.130.0/24	LAN	ONE_SEVEN_TWO (17...	172.16.130.1	198.51.100.30	198.51.100.30	256 of 256		Edit Release
BGP-SITE-03	172.16.131.0/24	LAN	ONE_SEVEN_TWO (17...	172.16.131.1			256 of 256		Edit Release
CORP-SITE-03	172.16.132.0/24	LAN	ONE_SEVEN_TWO (17...	172.16.132.1	198.51.100.30	198.51.100.30	256 of 256		Edit Release
GUEST-SITE-03	172.16.133.0/24	LAN	ONE_SEVEN_TWO (17...	172.16.133.1	198.51.100.30	198.51.100.30	256 of 256		Edit Release
RESEARCH-SITE-03	172.16.134.0/24	LAN	ONE_SEVEN_TWO (17...	172.16.134.1	198.51.100.30	198.51.100.30	256 of 256		Edit Release

Tech Tip

When selecting **Type** from the drop-down, the available options are **Generic**, **LAN**, **Management**, **Service**, and **WAN-Internal**. For the SD-Access workflows, use either **Generic** or **LAN** during the reservation process. **LAN** is used in the deployment. The remaining options are specific to NFV/ENCs provisioning and should only be used for that solution.



Procedure 5: Design Enterprise Wireless SSIDs for SD-Access Wireless

Cisco DNA Center is used to automate the creation of SSIDs, SSID security, AP profiles and parameters, and wireless Quality of Service (QoS) settings. Defining these elements in Cisco DNA Center, rather than defining each of these (per SSID) in the

GUI of the wireless LAN controller, saves time and reduces potential manual configuration errors.

The wireless workflow in the Design application is a two-step process:

1. Create the SSID and its parameters.
2. Create the wireless profile.

Creating the SSID involves defining the name along with the type of wireless network (voice, data, or both), wireless band, security type, and advanced options.

The wireless profile defines whether the wireless network is fabric or non-fabric and defines to which sites and locations within the hierarchy the profile is assigned. Wireless configurations, like IP address pools, are configured at the global level of the hierarchy.



The wireless network creation shown below uses recommend practices for security and roaming. Please adjust the options as necessary for a particular deployment.

To create the Enterprise wireless network:

1. In the Cisco DNA Center dashboard, navigate to **DESIGN > Network Settings > Wireless**.
2. In the **Enterprise Wireless** section click **+ Add**.
3. In the **Create an Enterprise Wireless Network** wizard, complete the following steps:
 - a. Enter the **Wireless Network Name (SSID)** (example: CORP).
 - b. Under **TYPE OF ENTERPRISE NETWORK**, select **Voice and Data** and **Fast Lane**.
 - c. Ensure **BROADCAST SSID:** is set to **On**.
 - d. Select the **WIRELESS OPTION** (example: Dual band operation with band select).
 - e. For **LEVEL OF SECURITY**, select **WPA2 Enterprise**.
 - f. Under **ADVANCED SECURITY OPTIONS**, select **Adaptive**.

The screenshot shows the 'Create an Enterprise Wireless Network' wizard in Cisco DNA Center. It is divided into two steps: '1 Enterprise Wireless Network' and '2 Wireless Profiles'. The current step is '1'. The form contains the following fields and options:

- Wireless Network Name(SSID) ***: A text input field containing 'CORP'.
- TYPE OF ENTERPRISE NETWORK ***: A group of radio buttons with 'Voice and Data' selected.
- BROADCAST SSID:**: A toggle switch set to 'On'.
- WIRELESS OPTION**: A group of radio buttons with 'Dual band operation with band select' selected.
- LEVEL OF SECURITY ***: A group of radio buttons with 'WPA2 Enterprise' selected.

At the bottom, there are three buttons: 'Previous', 'Cancel', and 'Next'. A mouse cursor is pointing at the 'Next' button. Below the buttons, the text 'ADVANCED SECURITY OPTIONS' is visible.

4. Click **Next** to continue in the wizard to the **Wireless Profiles** section.
5. Click the **+ Add** next to **Profiles**.

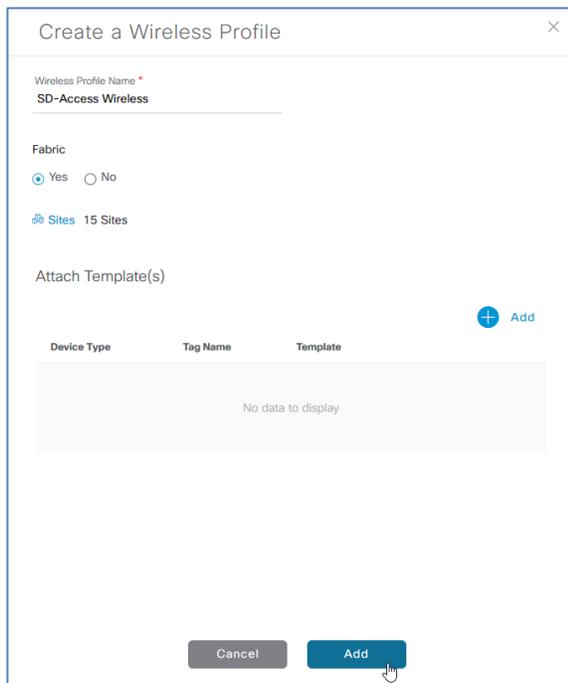
Create an Enterprise Wireless Network

- 1 Enterprise Wireless Network
- 2 Wireless Profiles

Profiles

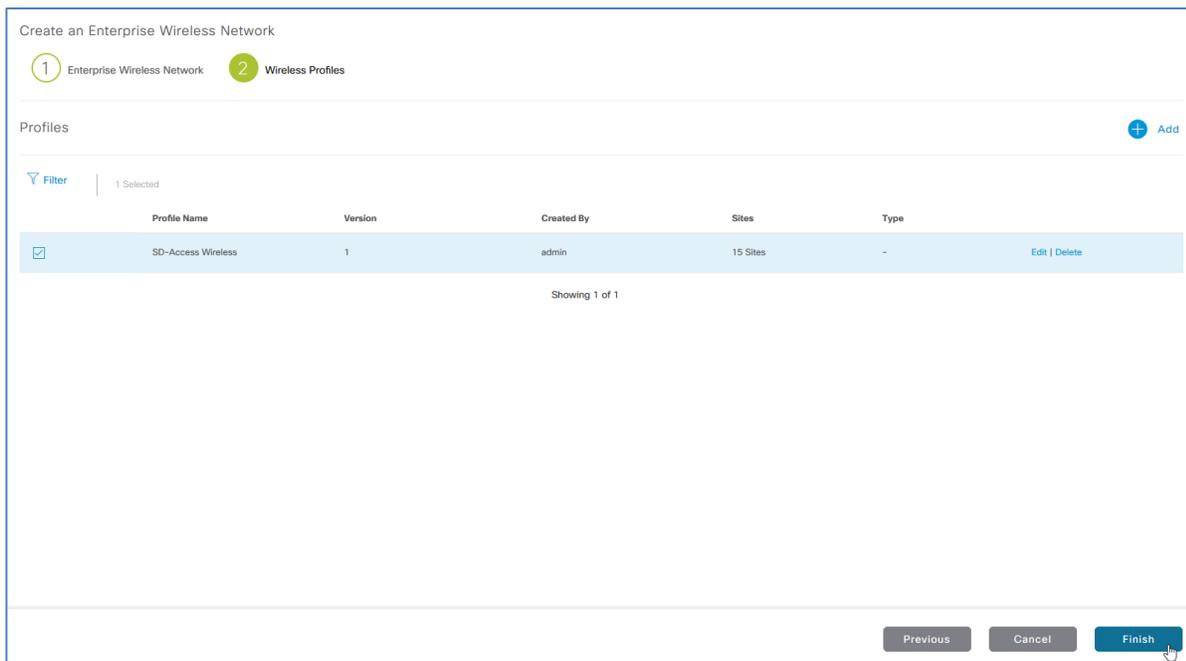


6. In the **Create Wireless Profile** dialog, complete the following steps:
 - a. Enter a **Wireless Profile Name** (example: SD-Access Wireless).
 - b. Under **Fabric**, select **Yes**.
 - c. Click **Sites** to open an additional **Sites** dialog.
 - d. Select the location(s) where the SSID broadcasts (example: San Jose)
 - e. Click **OK** to close the **Sites** dialog.
 - f. Click **Add** to close the **Create Wireless Profile** Wizard.



The screenshot shows the "Create a Wireless Profile" dialog box. At the top, it says "Create a Wireless Profile" with a close button (X). Below that, the "Wireless Profile Name" field is filled with "SD-Access Wireless". Under the "Fabric" section, the "Yes" radio button is selected. The "Sites" section shows "15 Sites" with a plus icon. Below that, the "Attach Template(s)" section has a table with columns "Device Type", "Tag Name", and "Template". The table is currently empty, showing "No data to display". To the right of the table is an "Add" button with a plus icon. At the bottom of the dialog, there are "Cancel" and "Add" buttons. A mouse cursor is pointing at the "Add" button.

7. Click **Finish**.



8. Repeat this procedure for any additional Enterprise wireless SSIDs.



The same wireless profile can be used for multiple SSIDs. Each SSID is broadcast at specific sites and locations as needed by defining only those sites in the workflow.

Procedure 6: Design a Guest Wireless SSID for SD-Access Wireless

Designing the guest wireless SSID in Cisco DNA Center is similar to designing the Enterprise wireless SSID. The primary difference is the Guest Web Authentication section in the workflow. Cisco DNA Center supports External Web Authentication and Central Web Authentication.

External Web Authentication uses the specified URL to redirect the guest users. Central Web Authentication uses the Guest Portal sequence of the Identity Services Engine to redirect guest users to the captive portal hosted on ISE.

The Guest wireless SSID creation workflow is a three-step process:

1. Create the SSID and its parameters.
2. Create the wireless profile.
3. Create the portal.

When creating the SSID, the SSID name is defined along with the Authentication Type which must be either **Web Authentication/Captive Portal** or **Open**. The profile defines whether the wireless network is fabric or non-fabric and defines to which sites and locations within the hierarchy the profile is assigned. The Guest Portal of ISE is customized and configured through Cisco DNA Center.

To create the Guest wireless network:

1. Navigate to **DESIGN > Network Settings > Wireless**.
2. In the Guest Wireless section, click **+ Add**.
3. In the **Create a Guest Wireless Network** wizard, complete the following steps:
 - a. Enter the **Wireless Network Name (SSID)** (example: GUEST).
 - b. Under **LEVEL OF SECURITY**, select **Web Auth**.

- c. Under **AUTHENTICATION SERVER**, select **ISE Authentication**.
- d. Leave the other default selections.

Network Device Credentials IP Address Pools SP Profiles **Wireless**

Create a Guest Wireless Network

1 Guest Wireless Network 2 Wireless Profiles 3 Portal Customization

Wireless Network Name (SSID) *
GUEST

BROADCAST SSID:

LEVEL OF SECURITY *
 Web Auth Open
Most secure
Guest users are redirected to a Web Portal for authentication

AUTHENTICATION SERVER *
 ISE Authentication External Authentication

What kind of portal are you creating today ?
Self Registered

Where will your guests redirect after successful authentication ?
Original URL

Previous Cancel **Next**

4. Click **Next** to continue in the wizard to **Wireless Profiles**.
5. In the **Wireless Profiles** section, select the **Profile Name** corresponding to the deployment location (example: SD-Access Wireless).

Network Device Credentials IP Address Pools SP Profiles **Wireless**

Create a Guest Wireless Network

1 Guest Wireless Network 2 **Wireless Profiles** 3 Portal Customization

Profiles

Filter 1 Selected

Profile Name	Version	Created By
<input checked="" type="checkbox"/> SD-Access Wireless	85	admin

Showing 1 of 1

Wireless Profile Name *
SD-Access Wireless

Fabric
 Yes No

Sites 14 Sites
Sites without configured ISE will be unselected automatically

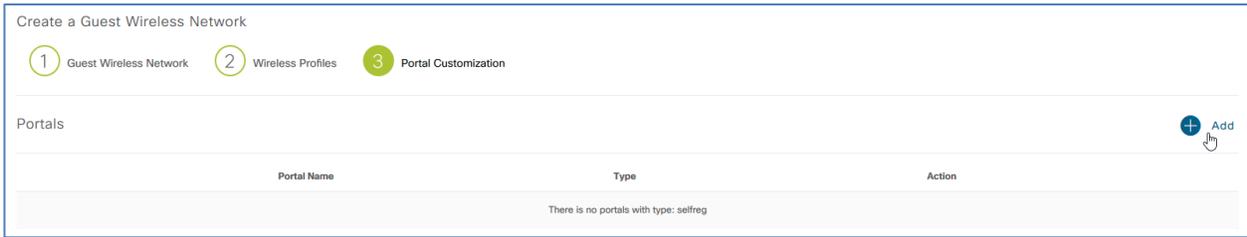
Attach Template(s)

+ Add

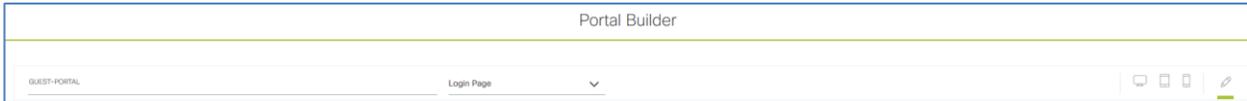
Device Type	Tag Name	Template
No data to display		

Cancel **Save**

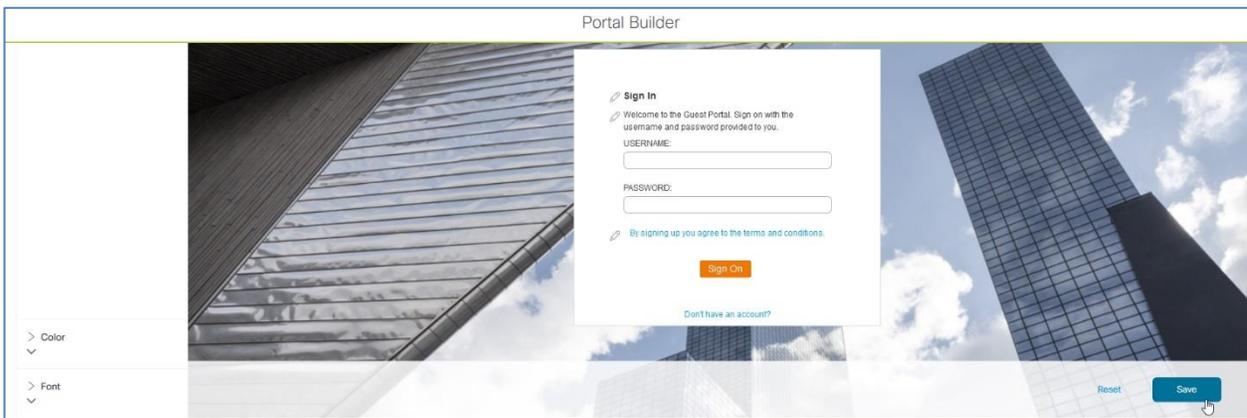
6. In the slide-in pane, keep the default **Fabric** selection of **Yes**.
7. Keep the other default selections.
If a new profile is created, please select the **Sites** accordingly.
8. At the bottom of the panel click **Save**.
9. Click **Next** to continue in the wizard to **Portal Customization**.
10. In the **Portals** screen, click **+ Add**.



- The **Portal Builder** wizard appears.
Supply a **Guest Portal name** (example: GUEST-PORTAL).

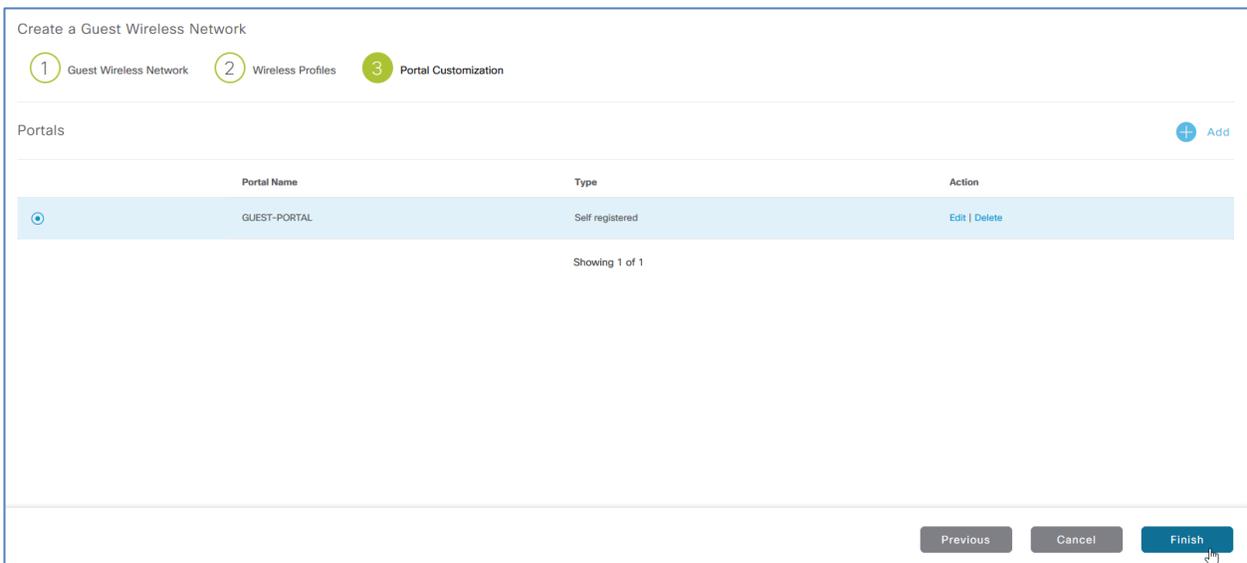


- Make any desired customizations.
- Use the scroll bar to navigate to the bottom of the portal.
- Once there, click **Save**.



- A guest web authentication portal is generated for the site(s), and the GUI returns to the previous screen.

Click **Finish** to complete the guest wireless LAN design.



Process 4: Creating Segmentation with the Cisco DNA Center Policy Application

SD-Access supports two levels of segmentation – macro and micro. Macro segmentation uses overlay networks - VRFs. Micro segmentation uses scalable group tags (SGTs) to apply policy to groups of users or devices.

In a University example, students and faculty machines may both be permitted to access printing resources, but student machines should not communicate directly with faculty machines, and printing devices should not communicate with other printing devices. This micro-segmentation policy can be accomplished using the Policy application in Cisco DNA Center which leverages APIs to program the ISE TrustSec Matrix.

For a deeper exploration of designing segmentation for SD-Access with additional use cases, see the [Software-Defined Access Segmentation Design Guide](#).

The Policy application supports creating and managing virtual networks, policy administration and contracts, and scalable group tag creation. Unified policy is at the heart of the SD-Access solution, differentiating it from others. Therefore, deployments should set up their SD-Access policy (virtual networks and contracts) before doing any SD-Access provisioning.

The general order of operation for SD-Access is Design, Policy, and Provision, corresponding with the order of the applications seen on the Cisco DNA Center dashboard.

In this section, the segmentation for the overlay network is defined. (Note that the overlay network will only be fully created until the host onboarding stage). This process virtualizes the overlay network into multiple self-contained virtual networks (VNs). After VN creation, the TrustSec policies are created to define which users and groups within a VN are able to communicate.

Use these procedures as prescriptive examples for deploying macro and micro segmentation policies using Cisco DNA Center.

Procedure 1: Add an Enterprise Overlay Virtual Network – Macro Segmentation

Virtual networks are created first, then group-based access control policies are used to enforce policy within the VN.

1. From the main Cisco DNA Center dashboard, navigate to **POLICY > Virtual Network**.
2. Click the **+** to create a new virtual network.
3. Enter a **Virtual Network Name** (example: CAMPUS).
4. Drag scalable groups from the **Available Scalable Groups** pane into the **Groups in the Virtual Network** pane (example: Employees, Contractors, and Network Services).
5. Click **Save**.

The screenshot shows the 'Create or Modify Virtual Network' configuration page. At the top, there are 'Reset' and 'Save' buttons. The 'Virtual Network Name' field contains 'CAMPUS'. There is a checkbox for 'Guest Virtual Network' which is currently unchecked. The interface is divided into two main panes: 'Available Scalable Groups' on the left and 'Groups in the Virtual Network' on the right. The 'Available Scalable Groups' pane shows a search bar and a grid of 14 group icons: AU (Auditors), BY (BYOD), DE (Developers), DS (Development Servers), GU (Guests), PC (PCI Servers), PO (Point of Sale), PS (Production Services), PU (Production Users), QS (Quarantined Systems), TS (Test Servers), TS (TrustSec Devices), and UN (Unknown). The 'Groups in the Virtual Network' pane shows a search bar and three group icons: EM (Employees), CO (Contractors), and NS (Network Services). A mouse cursor is hovering over the 'Save' button.

6. Verify that the VN with associated groups is defined and appears in the list on the left. These virtual network definitions are now available for provisioning the fabric in later steps.
7. Repeat this procedure for each overlay network.



Four common virtual networks appear in most networks: Campus/Staff, Guest/BYOD, BMS/IoT, and Records/Research with Campus/Staff and Guest/BYOD appearing in almost all production networks. This deployment guide will use these four common VNs.

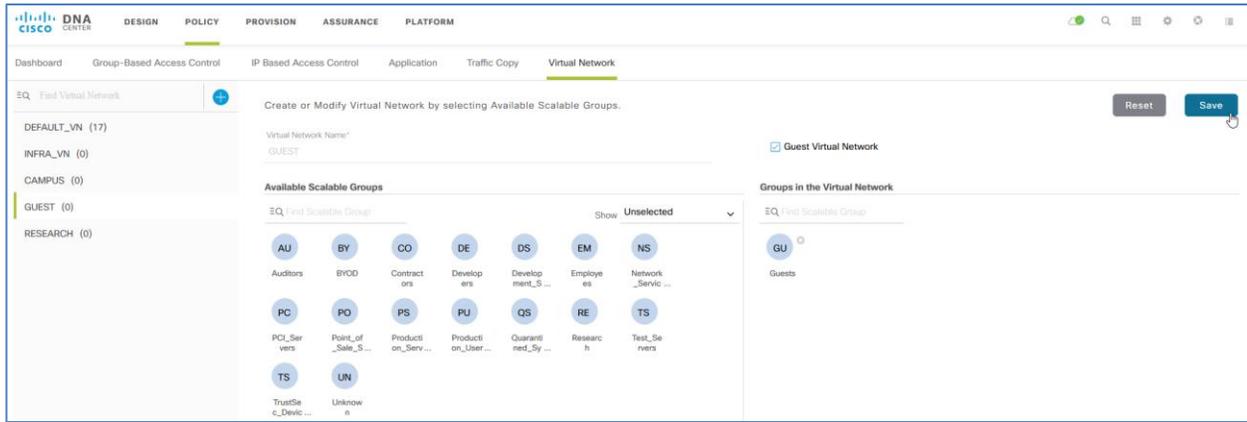


A common configuration convention is to use ALL CAPS for any user-defined elements. The VNs defined in the Policy application are provisioned to the devices as a VRF definition. Using ALL CAPS to identify these user-defined variables can assist in troubleshooting and monitoring. This convention is a best practice recommendation.

Procedure 2: Add a Guest Overlay Virtual Network – Macro Segmentation

Creating the guest virtual overlay network is very similar to creating of enterprise overlay virtual network. However, there must be a method to indicate that this virtual network is for guest access in order to trigger the use of the captive portal and guest SSID(s) created in earlier steps.

1. From the main Cisco DNA Center dashboard, click the **POLICY > Virtual Network**.
2. Click the **+** to create a new virtual network.
3. Enter a **Virtual Network Name** (example: GUEST).
4. Drag scalable groups from the **Available Scalable Groups** pane into the **Groups in the Virtual Network** pane (example: GUEST).
5. Select the **Guest Virtual Network** check box.
With the check box selected, this virtual network can be used for guest SSID, guest portals, and other guest-related workflows.
6. Click **Save**.



About Group-Based Access Control Policies

SD-Access segments the physical network into virtual networks, and security policies are defined to segment traffic inside of the VNs. Cisco DNA Center allows the administrator to explicitly deny or allow traffic between groups (SGTs) within each virtual network. This policy is created in Cisco DNA Center, pushed to ISE, and then finally pushed down to the switches to enforce the policy.

Group-Based Access Control Components

The SD-Access 1.2 solution supports creation and provisioning of the following policy constructs:

Security Group Access Control List (SGACL): Policy enforcement through which the administrator can control the operations performed by the user based on the security group assignments and destination resources. Cisco DNA Center refers to these as **Group-Based Access Control** policies to express the intent of these constructs.

Access Contract: Policy enforcement through which the administrator can control the operations performed by the user based on destination port and protocol information.



Tech Tip

SGACLs are sometimes referred to conceptually as *Layer-3 Policies* as they enforce traffic based on Layer-3 information learned from the SGT-to-IP bindings.

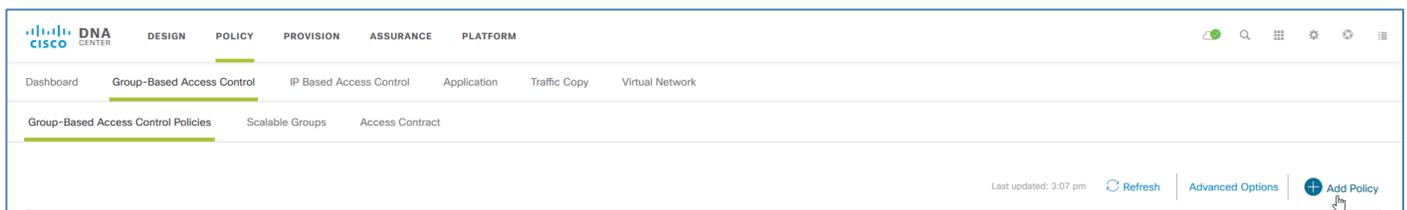
Access Contracts are sometimes referred to conceptually as *Layer-4 Policies* as they enforce traffic based on the Layer-4 TCP/UDP port information. Cisco DNA Center has two predefined access contracts, **permit** and **deny**. Respectively, they either allow all traffic or deny all traffic between the selected groups.

Procedure 3: Create Group-Based Access Control Policies using SGTs – Micro Segmentation

The following steps demonstrate how use Cisco DNA Center to create micro-segmentation security policies with just a few clicks. This simple example shows a basic policy that can be used to deny users from the Contractors group from communicating with the Employees group.

Using a single check box, Cisco DNA Center deploys this policy in ISE bidirectionally – also preventing the Employees group from communicating with the Contractors group.

1. In Cisco DNA Center, navigate to **POLICY > Group-Based Access Control > Group-Based Access Control Policies**.
2. Click **+ Add Policy**.



- From the **Available Scalable Groups** pane, drag the **Contractors** group and drop it into the **Source** pane.
- Drag the **Employees** group into the **Destination** pane.

The screenshot shows a configuration interface for Scalable Groups. On the left, the 'Available Scalable Groups' pane contains a search bar and a grid of 16 groups: AU (Auditors), BY (BYOD), CO (Contractors), DE (Developers), DS (Development_S...), EM (Employees), GU (Guests), NS (Network_Service...), PC (PCI_Servers), PO (Point_of_Sale_S...), PS (Producti_on_Serv...), PU (Producti_on_User...), QS (Quaranti_ned_Sy...), TS (Test_Se_rvers), TS (TrustSe_c_Devic...), and UN (Unknow_n). On the right, the 'Source Scalable Groups' pane contains CO (Contractors), and the 'Destination Scalable Groups' pane contains EM (Employees).

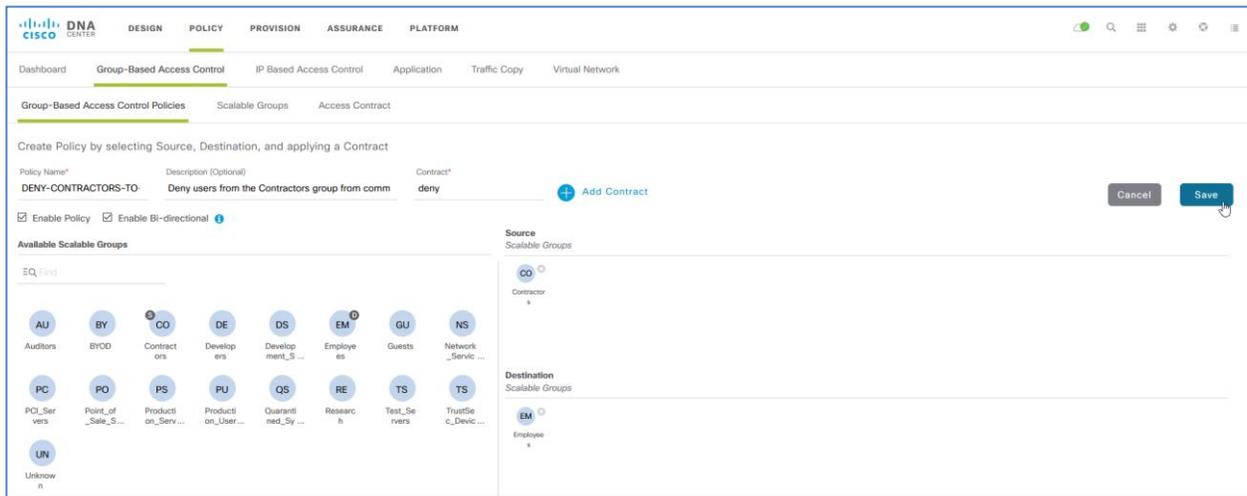
- Input a **Policy Name** (example: DENY-CONTRACTORS-TO-EMPLOYEES).
- Enter an optional **Description**.
- Select **Enable Policy**.
- Select **Enable Bi-directional**.

The screenshot shows the 'Create Policy' form. The 'Policy Name*' field contains 'DENY-CONTRACTORS-TO-', the 'Description (Optional)' field contains 'Deny users from the Contractors group from comm', and the 'Contract*' field is empty. The 'Enable Policy' and 'Enable Bi-directional' checkboxes are checked. An 'Add Contract' button is visible on the right.

- Click **+ Add Contract** to show the available **Access Contracts**.
- Select **deny**.
- Click **OK**.

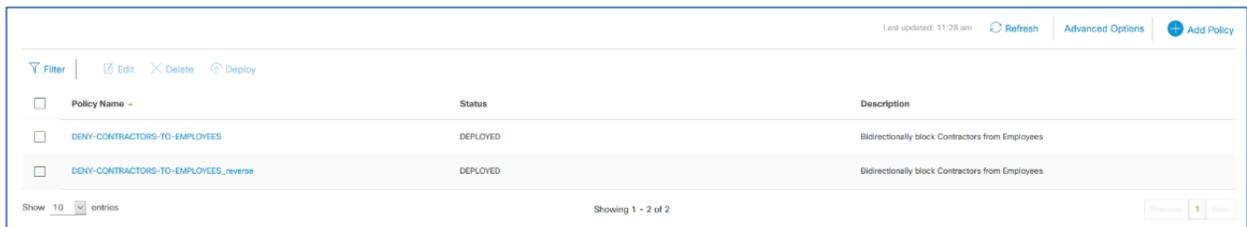
The screenshot shows the 'Access Contracts' dialog box. The 'deny' radio button is selected, and the 'OK' button is highlighted.

12. Click **Save**.



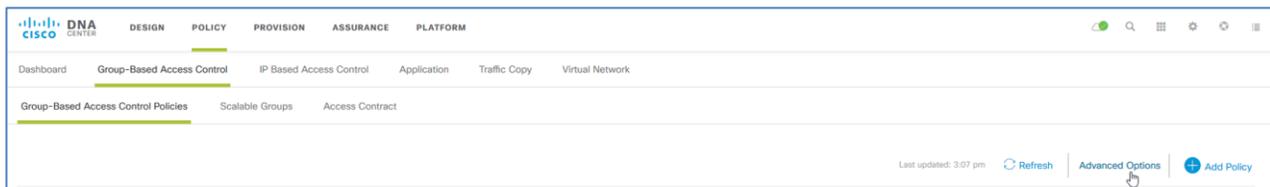
13. Verify that the policy is created and listed with a status of **DEPLOYED**.

14. Verify that the bidirectionally reverse policy is also created and **DEPLOYED**.



15. The policies are now available to be applied to fabrics and are also available in ISE.

To view them and cross-launch ISE using Cisco DNA Center, navigate to **POLICY > Group-Based Access Control > Group-Based Access Control Policies > Advanced Options**.



16. Log into ISE in the new browser tab that opens.

This cross-launch opens directly to the ISE TrustSec Matrix.

17. Verify that the policy is programmatically pushed to ISE.

The screenshot shows the Cisco Identity Services Engine (ISE) Production Matrix. The interface includes a navigation menu at the top with options like Home, Context Visibility, Operations, Policy, Administration, and Work Centers. The main area displays a grid with 'Source' on the vertical axis and 'Destination' on the horizontal axis. The sources listed are BYOD (15/000F), Contractors (5/0005), Developers (8/0008), Development_Ser... (12/000C), Employees (4/0004), and Guests (6/0006). The destinations listed are Auditors (9/0009), BYOD (15/000F), Contractors (5/0005), Developers (8/0008), Development_Ser... (12/000C), Employees (4/0004), Guests (6/0006), Network_Ser... (3/0003), and PCL_Servers (14/000E). Two cells in the grid are highlighted in red and labeled 'Deny IP': one for the source 'Employees' and destination 'Employees', and another for the source 'BYOD' and destination 'Contractors'. The matrix is titled 'Production Matrix' and shows 'Populated cells: 2'. At the bottom, there are controls for 'Default', 'Enabled', 'SGACLs: Permit IP', and a description: 'Default egress rule'.



The **Advanced Options** in Cisco DNA Center is a shortcut to logging into ISE and navigating to **Work Centers > TrustSec > TrustSec Policy > Egress Policy > Matrix**.

Process 5: Deploying SD-Access with the Provision Application

The Provision application in Cisco DNA Center is used to take the network components defined in the Design application and the segmentation created in the Policy application and deploy them to the devices. The Provision application has several different workflows that build on one another. This application can also be used for SWIM and LAN automation although these are beyond the scope of this document.

The process begins by assigning and provisioning devices to a site. Once devices are provisioned to a site, the fabric overlay workflows can begin. This starts through the creation of transits, the formation of a fabric domain, and the assignment of sites, buildings, and/or floors to this fabric domain. Once assigned, Cisco DNA Center lists these locations as *Fabric-Enabled*.

Once a site is fabric-enabled, devices provisioned to that site can be assigned a fabric role. This creates the core infrastructure for the fabric overlay. Host onboarding is completed afterwards to bind VNs and reserved IP address pools together, completing the overlay configuration.

About Provisioning in SD-Access

Provisioning in the SD-Access solution consists of three workflows:

1. Assign devices to site, building, or floor.
2. Provision devices added to site, building, or floor.
3. Provision devices to the fabric roles.

Assigning a device to a site causes Cisco DNA Center pushes certain site-level network settings configured in the Design

application to the devices, whether used as part of the SD-Access fabric overlay or not. Specifically, Netflow Exporter, SNMP server and traps, and syslog server information configured in the Design application, are provisioned on the devices.

After provisioning devices to a site, the remaining network settings from the Design application are pushed down to the devices. These include time zone, NTP server, and AAA configuration.

When devices are defined in fabric roles, the fabric overlay configuration is pushed to the devices.

The above workflows are described as separate actions as they support different solution workflows. The first is required for the Assurance solution, and all three workflows are required for the SD-Access solution.



Both assigning a device to a site and provisioning a device to a site can be completed as a single step or as two independent steps. This document showcases the single step option. The configuration and provisioning does not change if these steps are done together or independently.

Procedure 1: Create Credentials in ISE

When devices are provisioned to a site, they receive several configurations, including the centralized AAA server configuration pointing to ISE. Authentication against ISE is preferred over local login credentials within this configuration.

Devices provisioned to the site authenticate and authorize their console and VTY lines against the configured ISE server. To maintain the ability to manage the devices after provisioning, the credentials defined during the discovery job must be available in ISE, either directly or through an external identity source such as Active Directory.

The following steps show how to create internal user credentials in the Identity Services Engine that match the user credentials defined during the discovery job.

1. Log into the ISE GUI.
2. Navigate to **Administration > Identity Management > Identities**.
3. Click **+Add**.
4. Enter the **Name** (example: Operator).
This name must match what was used for Cisco DNA Center discovery job.
5. Enter the associated **Login Password** and **Re-Enter Password**.
6. Click **Submit**.

The screenshot shows the Cisco Identity Services Engine (ISE) interface for creating a new Network Access User. The breadcrumb navigation is: Home > Context Visibility > Operations > Policy > Administration > Work Centers > Identity Management > Network Resources > Device Portal Management > pxGrid Services > Feed Service > Threat Centric NAC > Identities > Users > Network Access Users List > New Network Access User.

The form includes the following sections and fields:

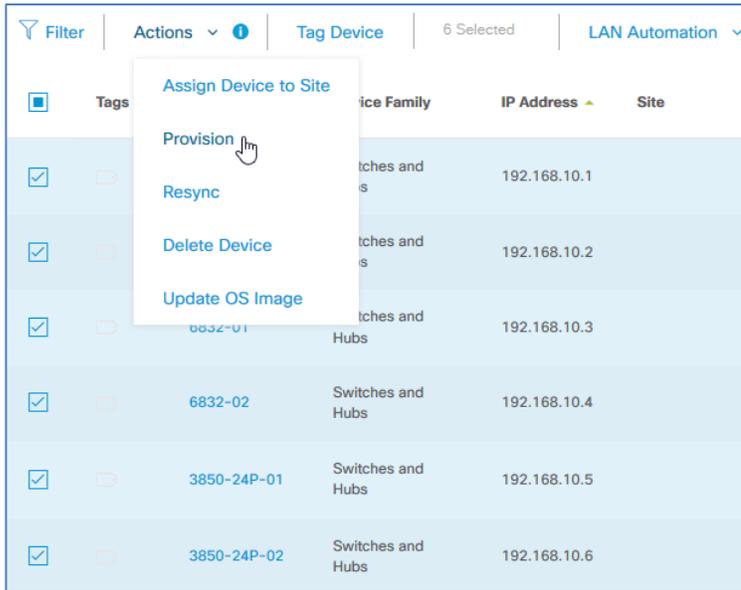
- Network Access User:**
 - * Name: Operator
 - Status: Enabled
 - Email: [Empty field]
- Passwords:**
 - Password Type: Internal Users
 - * Login Password: [Masked] [Generate Password]
 - Re-Enter Password: [Masked] [Generate Password]
 - Enable Password: [Empty] [Generate Password]
- User Information:**
 - First Name: [Empty]
 - Last Name: [Empty]
- Account Options:**
 - Description: [Empty]
 - Change password on next login:
- Account Disable Policy:**
 - Disable account if date exceeds: 2019-06-03 (yyyy-mm-dd)
- User Groups:**
 - Select an item [Dropdown]

At the bottom of the form are **Submit** and **Cancel** buttons.

Procedure 2: Assign Network Devices to Site and Provision Network Settings

The first step of the provisioning process begins by selecting devices and assigning them to a site, building, or floor previously created with the Design application. Both assigning and provisioning to site is completed in a single procedure rather than doing each step independently.

1. In Cisco DNA Center, navigate to **PROVISION > Devices > Inventory**.
2. Select the devices to be assigned and provisioned to a site.
3. Click **Actions**.
4. Click **Provision**.



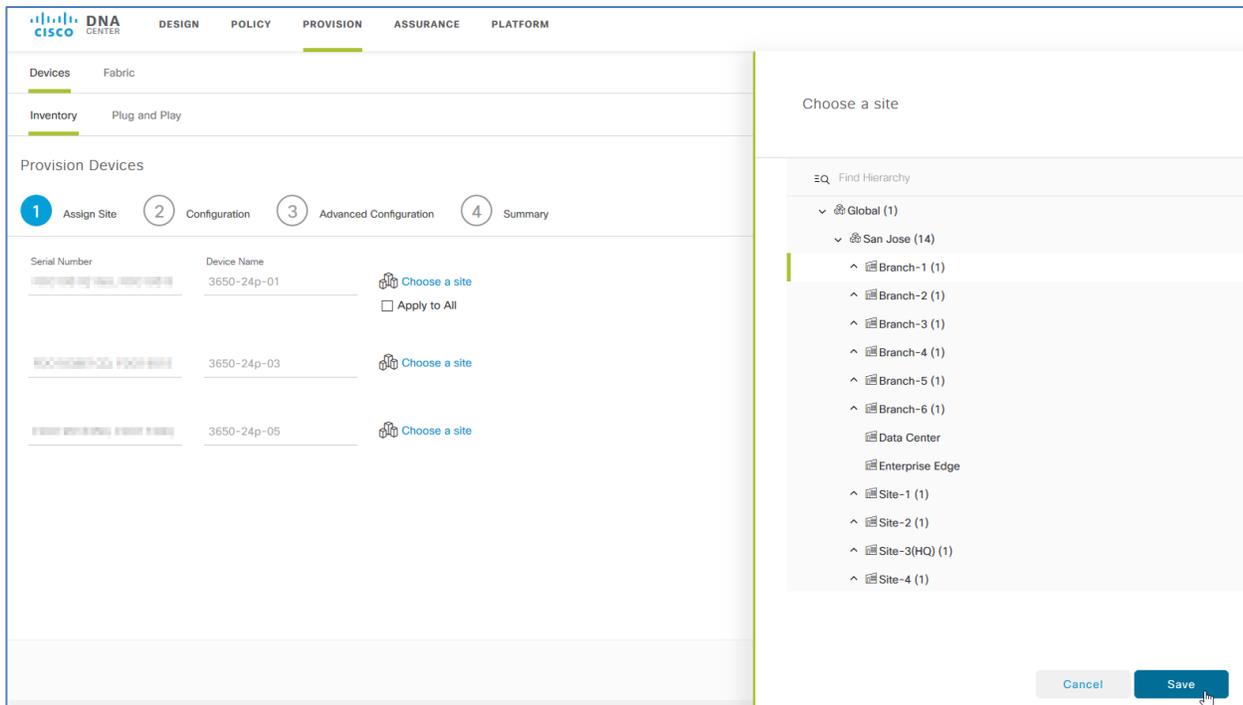
5.



The most commonly used approach is to provisioning devices that share common site assignments in multiple small batch groups. To do this, devices must be of the same type (example: *Switches and Hubs*) to provision them at the same time.

In IOS-XE \geq 16.8.1, Catalyst 9000 series switches appear as *Switches and Hubs (WLC Capable)* in the GUI. Cisco DNA Center categorizes these devices differently than the other Catalyst Switches – *Switches and Hubs* – and therefore they must be provisioned separately.

6. In the first page of the **Provision Devices** wizard, click **Choose a site**.
7. In the slide-in pane on the right, select the site assignment for the devices.
8. Click **Save**.



If the devices are all in the same site, the **Apply to All** check box can be selected.

Different site locations can be selected for each device where applicable.

Serial Number	Device Name	Site Location
██████████	C-6840-03	Global/San Jose/Site-1
██████████	C-6840-04	Global/San Jose/Site-1
██████████	6832-01	Global/San Jose/Site-1
██████████	6832-02	Global/San Jose/Site-1
██████████	3850-24P-01	Global/San Jose/Site-1

9. Use **Next** to skip the **Configuration** and **Advanced Configuration** screens.
10. In the **Summary** screen, review the details for each device.
11. Finally, click **Deploy**.

Device Name	Value
Device Name	C-6840-03
Platform Id	C6840-X-LE-40G
Device IP	192.168.10.1
Device Location	

Network Setting	Value
NTP Server	100.119.4.17
NTP Server	100.119.4.18
AAA Network Primary Server	198.51.100.127
AAA Client Primary Server	198.51.100.127
SYSLLOG Server	
SNMP Trap Server	
DNS Domain Name	@na.local
DNS Primary Server	198.51.100.30

12. In the **Provision Device** slide-in pane, leave the default selection of **Now**.
13. Click **Apply**.

When

Now Later

Cancel Apply

14. Configuration of each device begins, and status messages appear as each device is provisioned successfully. The Device Inventory screen updates with **Last Provisioned Time** and **Provision Status**.

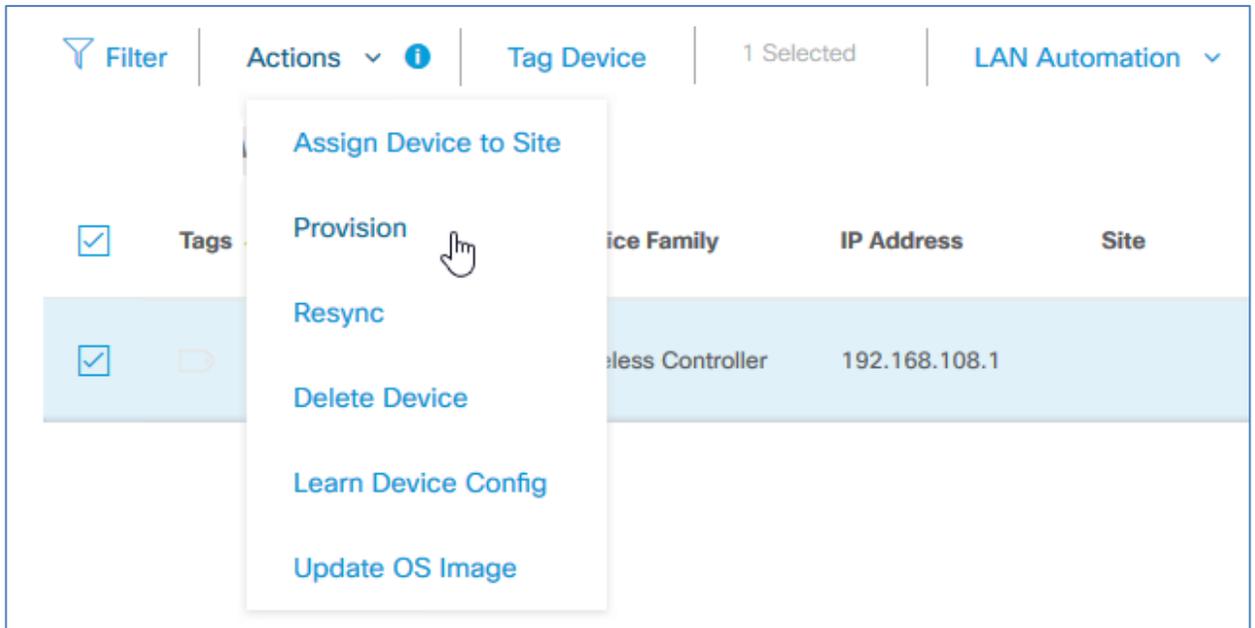
Use the **Refresh** button to view the final status of the devices once the status messages fade.

15. Repeat the previous steps to assign and provision the remainder of the devices to the applicable sites for the deployment.

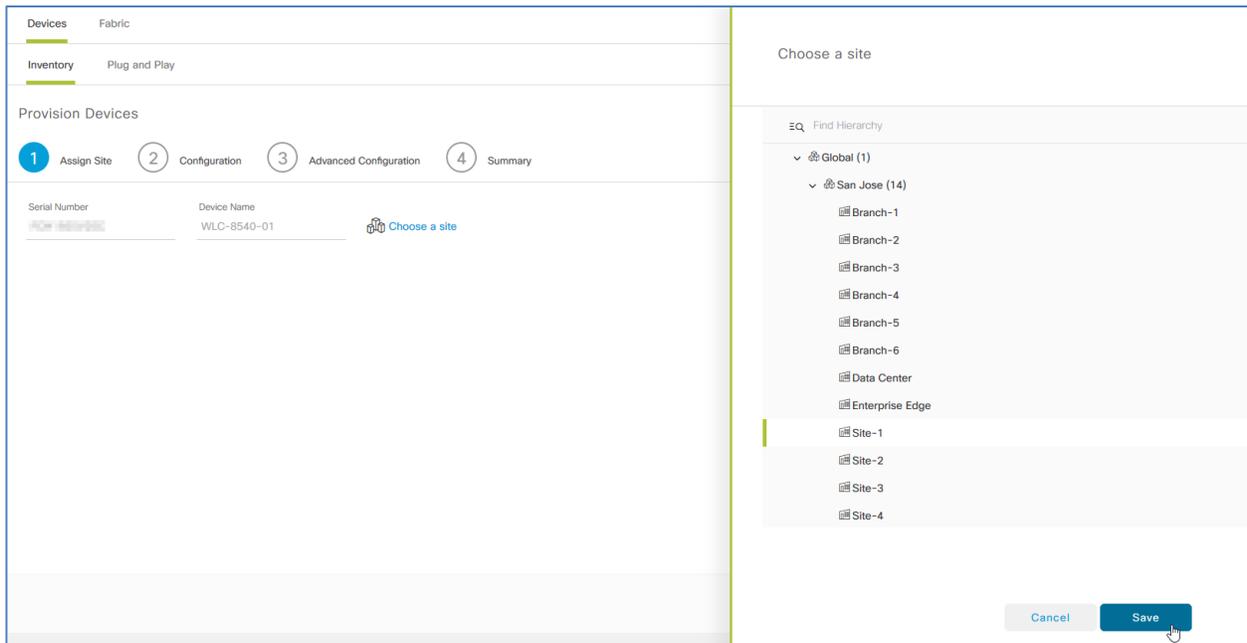
Procedure 3: Assign Wireless LAN Controller to Site and Provision Network Settings

The workflow for provisioning site assignment and network settings for the WLC is slightly different from the procedure for the routers and switches. The WLC must be assigned as the Active or Guest Anchor WLC. Potentially, APs from multiple locations may be associating with the same WLC, so the WLC must also be set to manage wireless for specific locations.

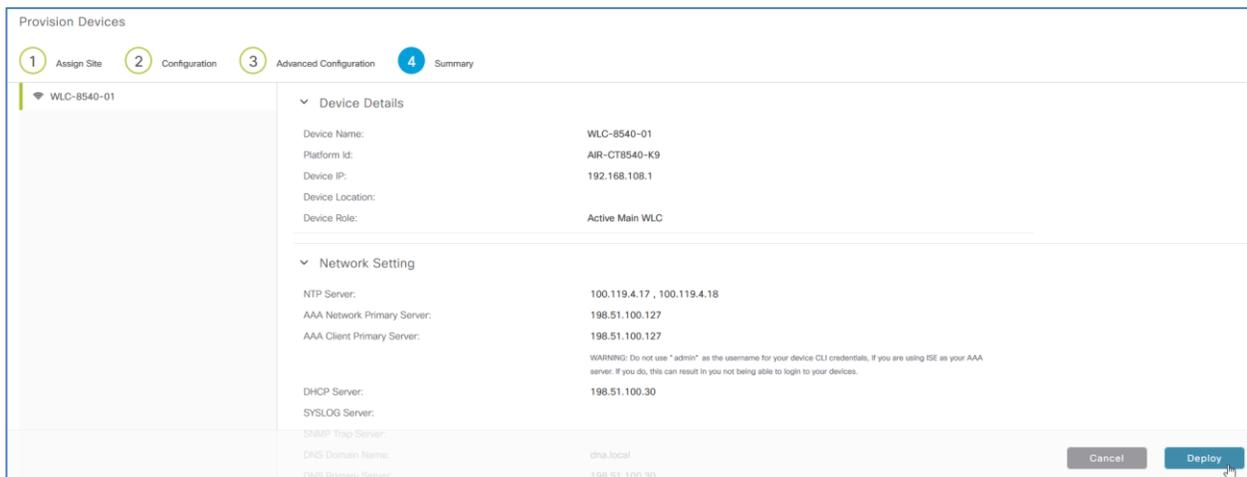
1. In Cisco DNA Center, navigate to **PROVISION > Devices > Inventory**.
2. Select a Wireless LAN Controller.
3. Click **Actions > Provision**.



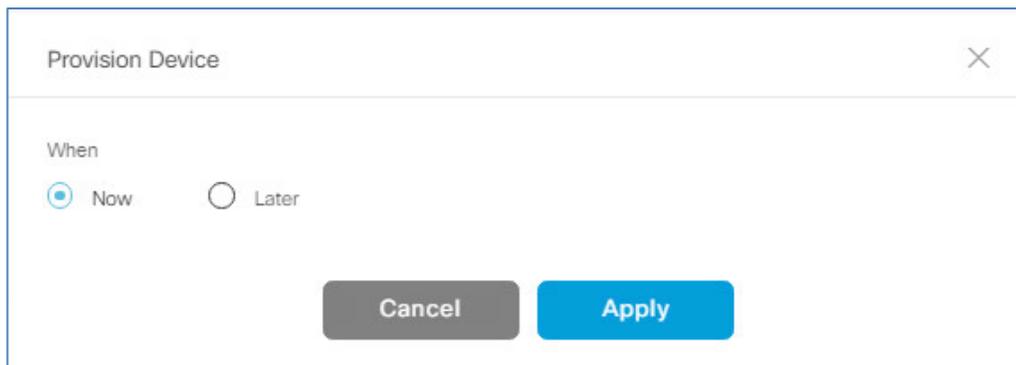
4. In the first page of the **Provision Devices** wizard, click **Choose a site**.
5. In the slide-in pane on the right, select the site assignment for the device. This defines WLC will be present in the fabric topology maps.
6. Click **Save**.



7. Click **Next**.
8. Select the appropriate **WLC Role** (example: Active Main WLC).
9. If the WLC is managing APs in locations other than the site to which it is assigned, these locations can be selected by clicking **Managing location(s)**.
10. Use **Next** to skip the **Advanced Configuration** screen.
11. In the **Summary** screen, review the details for the device.
12. Finally, click **Deploy**.



13. In the **Provision Device** slide-in pane, leave the default selection of **Now**.
14. Click **Apply**.

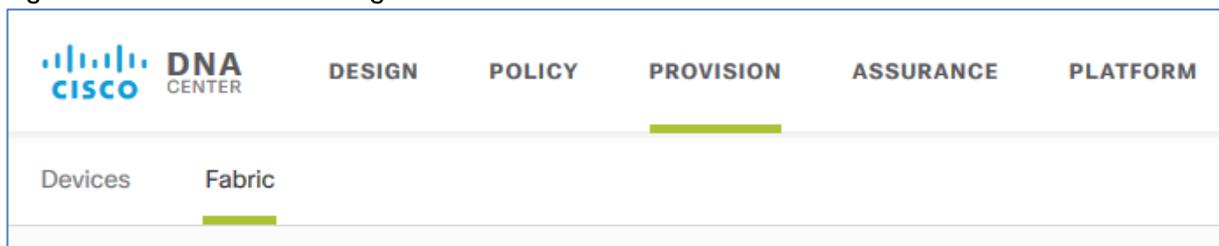


15. Configuration of the WLC begins, and status messages appear as the device is provisioned successfully. The Device Inventory screen updates with **Last Provisioned Time** and **Provision Status**. Use the **Refresh** button to view the final status of the device once the status messages fade.

Process 6: Provisioning the Fabric Overlay

Creating the fabric overlay is a multi-step workflow. Each of fabric overlay steps are managed under the **Fabric** tab of the Provision application.

Figure 11 Fabric Provisioning Tab



Provisioning the fabric overlay involves the following steps:

1. Create transits.
2. Create the fabric domain.
3. Assign and provision fabric roles.
4. Set up host onboarding.

The first step is to create the transit. The transit simply represents a connection that joins a fabric site to an external network (IP-Based transit) or to one or more fabric sites (SD-Access transit). Both IP-based transits and SD-Access transits are created in this guide. The transits are referenced later during the provisioning of fabric border nodes.



The SD-Access transit will be used between the sites in the domain, and the IP-based transit will be used to connect to the Data Center and shared service and to the Internet.

Please see [Appendix C](#) for detailed topology diagrams.

After the transit creation, a fabric domain is created. A fabric domain is a logical organizational construct in Cisco DNA Center. It contains fabric-enabled sites and their associated transits. Sites are added to the fabric domain one at a time. During fabric role provisioning, a transit is associated with a border node, joining the transit and site together under the domain.

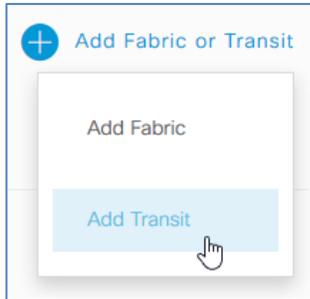
With the domain created and sites added to it, devices in each site are then added to fabric roles. Border node provisioning is addressed separately in the document due to the different types (Internal, External, and Anywhere) and the separate transit options. Once the SD-Access overlay infrastructure (fabric devices) is provisioned, host onboarding and port assignment are completed allowing endpoint and Access Point connectivity.

Procedure 1: Create IP-Based Transits – Fabric Provisioning Part I

The IP-based transit represents the remote BGP autonomous system (AS). The local BGP AS is configured as part of the fabric border provisioning in subsequent steps. The topology in this prescriptive guide includes two IP-based transits and an SD-Access transit.

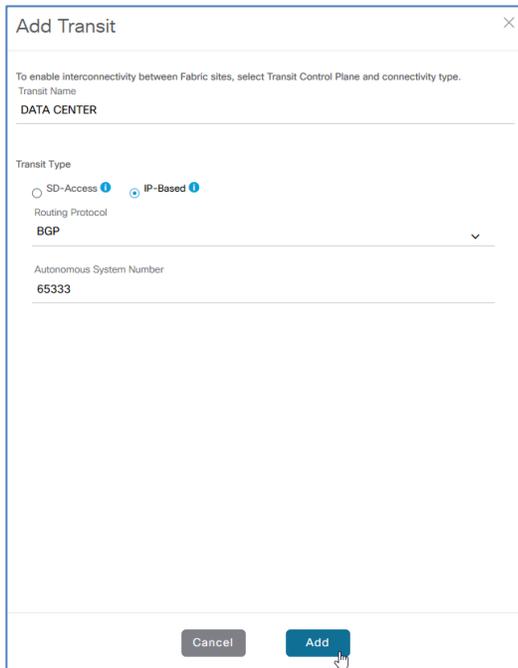
The IP-based transits will be used to automate the border connectivity between Site-1 and the Data Center and between Site-3 and the Enterprise edge (Internet edge) respectively. The SD-Access transit will be used to interconnect the various sites in the distributed deployment thus allowing for unified policy configuration, constructs, and enforcement between these sites.

1. In the Cisco DNA Center dashboard, navigate to **PROVISION > Fabric**.
2. At the top right, click **+ Add Fabric Domain or Transit**.
3. Click **Add Transit**.



4. In the slide-in pane, select the **Transit Type** of **IP-Based**.
5. For **Routing Protocol** select **BGP** – this is the only option currently.
6. Supply a **Transit Name** (example: DATA CENTER).
7. Enter an **Autonomous System Number** for the remote BGP AS (example: 65333).
8. Click **Add**.

A status message appears, and the transit is created.

A screenshot of a configuration form titled "Add Transit" with a close button (X) in the top right corner. The form contains the following fields:

- A note: "To enable interconnectivity between Fabric sites, select Transit Control Plane and connectivity type."
- "Transit Name": A text input field containing "DATA CENTER".
- "Transit Type": Two radio buttons, "SD-Access" (unselected) and "IP-Based" (selected).
- "Routing Protocol": A dropdown menu showing "BGP".
- "Autonomous System Number": A text input field containing "65333".

At the bottom of the form, there are two buttons: "Cancel" and "Add". A mouse cursor is pointing at the "Add" button.

9. Repeat the previous steps to create a second **IP-Based** transit.
A different **Transit Name** is required (example: INTERNET EDGE).

Add Transit

To enable interconnectivity between Fabric sites, select Transit Control Plane and connectivity type.

Transit Name
INTERNET EDGE

Transit Type
 SD-Access **i** IP-Based **i**

Routing Protocol
BGP

Autonomous System Number
65535

Cancel **Add**

Procedure 2: Create SD-Access Transit – Fabric Provisioning Part I (Continued)

The SD-Access transit is the circuit between sites in an SD-Access for Distributed Campus deployment. It uses the signaling from transit control plane nodes to direct traffic between fabric sites.



Tech Tip

Transit control plane nodes should always be deployed as a redundant pair of devices.

Unlike creating an IP-based transit, there is no requirement to define a BGP Autonomous System number, as BGP private AS 65540 is reserved for use on the transit control plane nodes and automatically provisioned by Cisco DNA Center. Using this AS, the SD-Access transit control plane nodes will form EBGP adjacencies with each site border in the Distributed Campus deployment.

When creating an SD-Access transit, the workflow also requires defining the devices used as the transit control plane nodes. Upon creating the transit and defining the nodes, they are provisioned automatically. Additional steps are not required to provision these devices.

1. In the Cisco DNA Center dashboard, navigate to **PROVISION > Fabric**.
2. At the top right, click + **Add Fabric Domain or Transit**.
3. Click **Add Transit**.
4. In the slide-in pane, select the **Transit Type** of **SD-Access**.
5. Supply a **Transit Name** (example: METRO-E-SDA).
6. Using the first **Site for the Transit Control Plane** drop-down, select the applicable device as a Transit Control Plane.
7. Using the second **Site for the Transit Control Plane** drop-down, select the second applicable device as a Transit Control Plane.
8. Click **Add**.

Add Transit
✕

To enable interconnectivity between Fabric sites, select Transit Control Plane and connectivity type.

Transit Name
METRO-E-SDA

Transit Type

SD-Access ? IP-Based ?

Site for the Transit Control Plane
Global/San Jose/Enterprise Edge ▼ 🗑️

Transit Control Plane
ISR-4451-05 ▼ 🗑️

Site for the Transit Control Plane
Global/San Jose/Enterprise Edge ▼ 🗑️

Transit Control Plane
ISR-4451-06 ▼ 🗑️

Cancel
Add

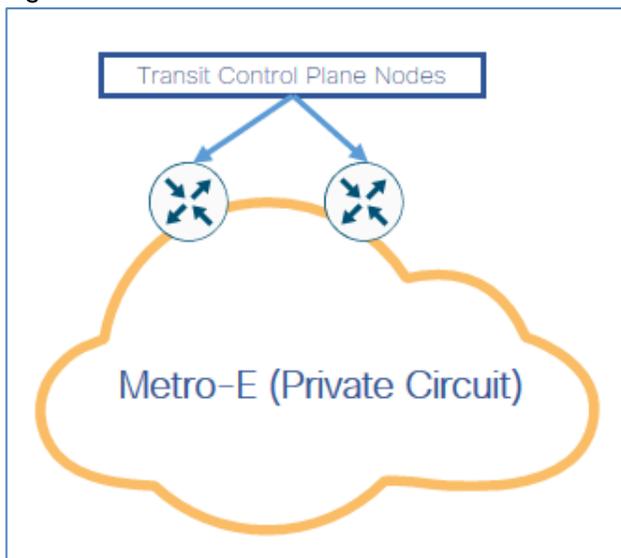


Tech tip The **Site for the Transit Control Plane** drop-down shows all the devices not currently provisioned for a fabric role in the selected site.

About Transit Control Plane Node Location

As noted in the [Topology Overview](#) section, a pair of ISR 4451-x routers were used as the transit control plane nodes. These routers were deployed (assigned and provisioned to) their own, dedicated site. These transit control plane nodes are accessible and have IP reachability via the Metro-Ethernet circuit connecting all the sites.

Figure 12 Transit Control Plane Node Site



This is just one design option. Transit control plane nodes could physically be in the same site as other equipment in the deployment. The primary requirements are to have IP reachability and that the transit control plane nodes cannot also operate as another role in the fabric (example: both transit control plane node and site border).

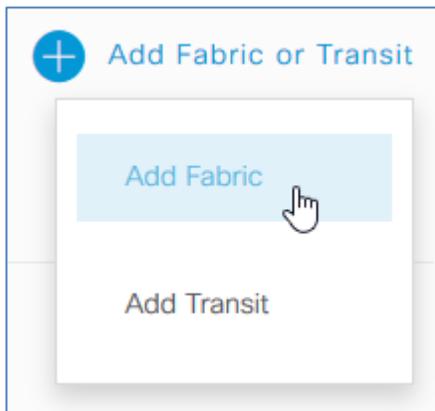
Procedure 3: Create a Fabric Domain – Fabric Provisioning Part 2

A fabric domain is an administrative construct in Cisco DNA Center. It is combination of fabric sites along with their associated transits. Transits are bound to sites later in the workflow when the fabric borders are provisioned. A fabric domain can include all sites in a deployment or only certain sites depending on the physical and geographic topology of the network and the transits. The prescriptive deployment includes a single fabric domain that will encompass the buildings (sites) created during previous steps in the Design application.



Cisco DNA Center 1.2.0 introduced the ability to have multiple fabric sites within the topology maps in the Provision application. This introduced a change in the workflow from Cisco DNA Center 1.1.x. Each site must be enabled for fabric, as shown in the following steps. A fabric domain must be created first, and then individual sites are added as fabric-enabled sites. This allows the devices assigned and provisioned to that site to appear in the topology maps.

1. In the Cisco DNA Center dashboard, navigate to **PROVISION > Fabric**.
2. At the top right, click **+ Add Fabric or Transit**.
3. Click **Add Fabric**.



4. In the slide-in pane, define a **Fabric Name** (example: SJC).
5. Select a single location from the **Hierarchy** (example: Site-1).
6. Click **Add**.

Add Fabric

Name the Fabric and choose a location for common policy enforcement. All sites in the chosen location will be added to the Fabric.

Fabric Name
SJC

Select a location to create a Fabric. All sites in the chosen location will be added to the Fabric

Find Hierarchy

- Global (1)
 - San Jose (14)
 - Branch-1
 - Branch-2
 - Branch-3
 - Branch-4
 - Branch-5
 - Branch-6
 - Data Center
 - Enterprise Edge
 - Site-1
 - Site-2

Cancel Add



The **Fabric Name** (fabric domain name) must be different from any other area, building, or floor name in the hierarchy.



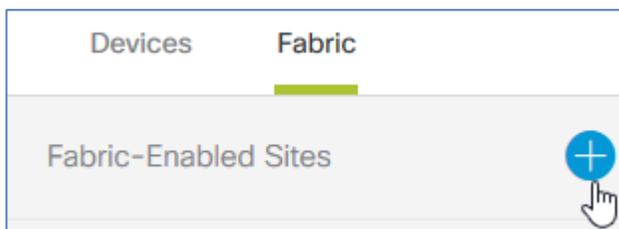
For added granularity in the topology maps, in fabric site provisioning, and in fabric role provisioning, the buildings (each site) are added individually to the domain rather than adding the parent area (San Jose). Steps for adding sites individually will become clearer in the next procedure.

If the parent is selected rather than individual sites, all equipment assigned and provisioned within locations encompassed by that parent area will be displayed on the same topology maps. When the sites are added individually, the topology maps only display the equipment assigned and provisioned within that location.

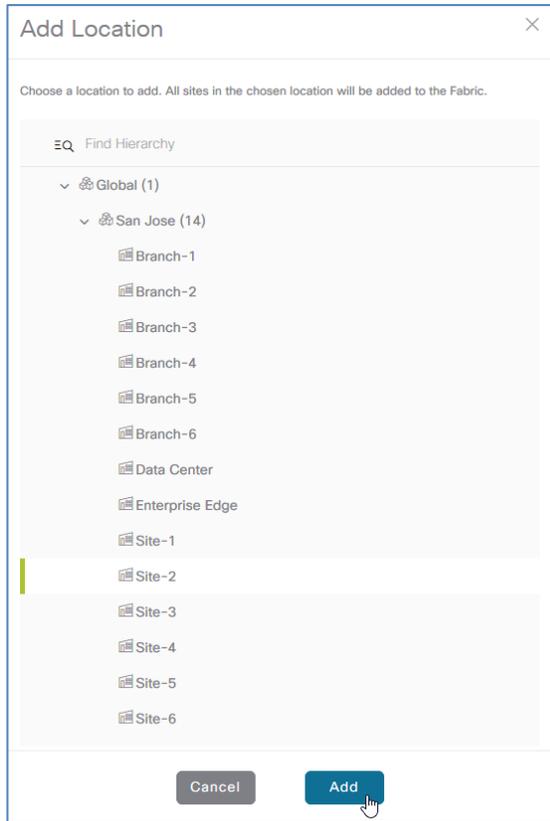
Procedure 4: Add Fabric-Enabled Sites to the Fabric Domain – Fabric Provisioning Part 2 (Continued)

Each building created in the Design application is intended to represent a separate fabric site. In the previous steps, a fabric domain was created, and a single site was added to that domain. The following steps will add the remainder of the sites to the domain.

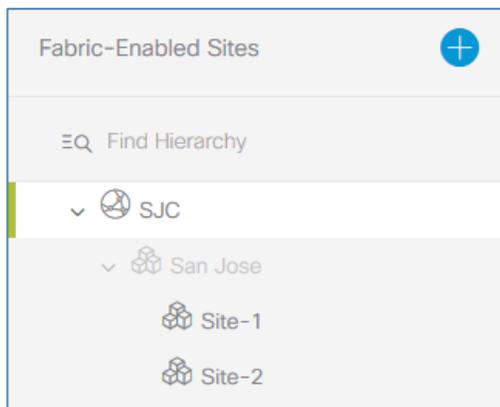
1. In the Cisco DNA Center dashboard, navigate to **PROVISION > Fabric**.
2. Select the newly created fabric domain (example: SJC).
3. Select the **+** button next to **Fabric-Enabled Sites**.



4. In the **Add Location** slide-in pane, select an additional site in the hierarchy to add to as a fabric-enabled site to this fabric domain.



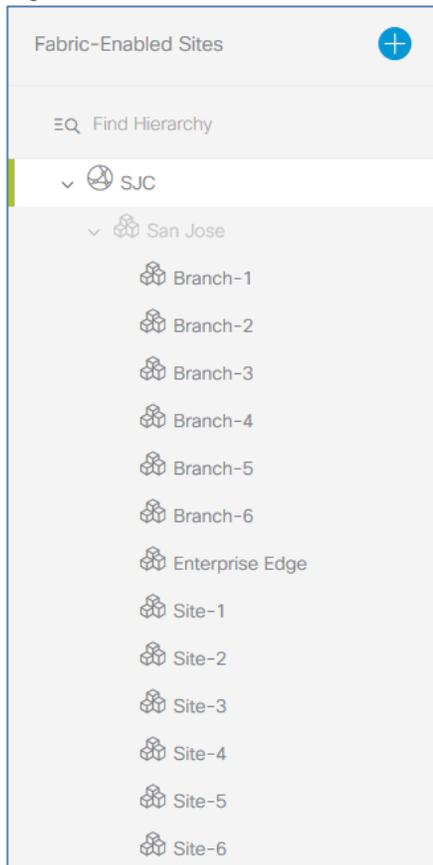
5. Click **Add**.
6. Verify that the site is added to the domain hierarchy on the left.



7. Repeat these steps to add all the sites that are associated with this fabric domain.

The final hierarchy of fabric-enabled sites for the fabric domain is shown below.

Figure 13 Fabric-Enabled Sites



Process 7: Assigning Fabric Roles – Fabric Provisioning Part 3

A fabric overlay consists of three different fabric nodes: control plane node, border node, and edge node.

To function, a fabric must have an edge node and control plane node. This allows endpoints to traverse their packets across the overlay to communicate with each other (policy dependent). The border node allows communication from endpoints inside the fabric to destinations outside of the fabric along with the reverse flow from outside to inside.

About Fabric Border Nodes

When provisioning a border node, there are number of different automation options in the Cisco DNA Center GUI.

A border node can have a Layer-3 handoff, a Layer-2 handoff, or both (platform dependent). A border node can be connected to an IP transit, to an SDA transit, or both (platform dependent). A border node can provide connectivity to the Internet, connectivity outside of the fabric site to other non-Internet locations, or both. It can operate strictly in the border node role or can also operate as both a border node and control plane node. Finally, border nodes can either be routers or switches which creates slight variations in the provisioning configuration to support fabric DHCP and the Layer-3 handoff.

With the number of different automation options, it is important to understand the use case and intent of each selection to have a successful deployment.



The Layer-2 handoff is used for migration purposes. It connects an SD-Access network (site) to a legacy (non-SD-Access network) and allows for communication between endpoints in both networks along with overlapping IP address space between them. Layer-2 handoffs are currently supported on the Catalyst 9000 Series Switches although this feature is not covered in this prescriptive deployment guide. For additional information, please see the Cisco Live BRKCRS links in [Appendix B](#).

Border Nodes – Network Types

When provisioning a device as a border node, there are three options to indicate the type of network(s) to which the border node is connected:

- Rest of Company (Internal)
- Outside World (External)
- Anywhere (Internal & External)

An **Internal** border is connected to the known routes in the deployment such as a Data Center. As an **Internal** border, it will register these known routes with the site-local control plane node which directly associates these prefixes with the fabric.

An **External** border is connected to unknown routes such as the Internet, WAN, or MAN. It is the gateway of last resort for the local site's fabric overlay. A border connected to an SD-Access transit must always use the **External** border functionality. It may also use the **Anywhere** border option as described below.

An **Anywhere** border is used when the network uses one set of devices to egress the site. It is directly connected to both known and unknown routes. A border node connected to an SD-Access transit may use this option if it is also connected to a [fusion router](#) to provide access to shared services.



Review the [Topology Diagrams](#). This deployment does not utilize any borders operating in the **Anywhere** role. The distinction is important to understand for production deployments. Selection of the wrong border node type can impact the success of the deployment or result in unnecessary configurations that consume and utilize resources.

Although the site border nodes in Site-2, Site-3, Site-4 and Site-5 are physically traversed to access both shared services (known routes) and the Internet (unknown routes), these border nodes are not connected to the fusion routers nor are they required to register the shared services routes with their site-local control plane nodes. They are therefore not **Anywhere** borders.

Border selection type is **critical** for a successful deployment.

Border Nodes – Fabric Roles

When provisioning a border node, the device can also be provisioned as a control plane node. Alternatively, the control plane node role and border node role can be independent devices. While the resulting configuration on the devices is different based on these selections, Cisco DNA Center abstracts the complexity, understands the user's intent, and provisions the appropriate resulting configuration to create the fabric overlay.

In the GUI, Cisco DNA Center refers to these provisioning options as **Add as CP**, **Add as Border**, and **Add as CP+Border**.

Which option should be deployed is ultimately a design discussion based on the existing technology, processes, and routing that is operating on the intended border node device, along with the scale requirements of the site and the physical layout of the topology itself. For further discussion, please see the Software-Defined Access Design and Deployment CVDs in [Appendix B](#) along with Cisco Live [BRKCRS-2810](#) and [BRKCRS-2811](#).



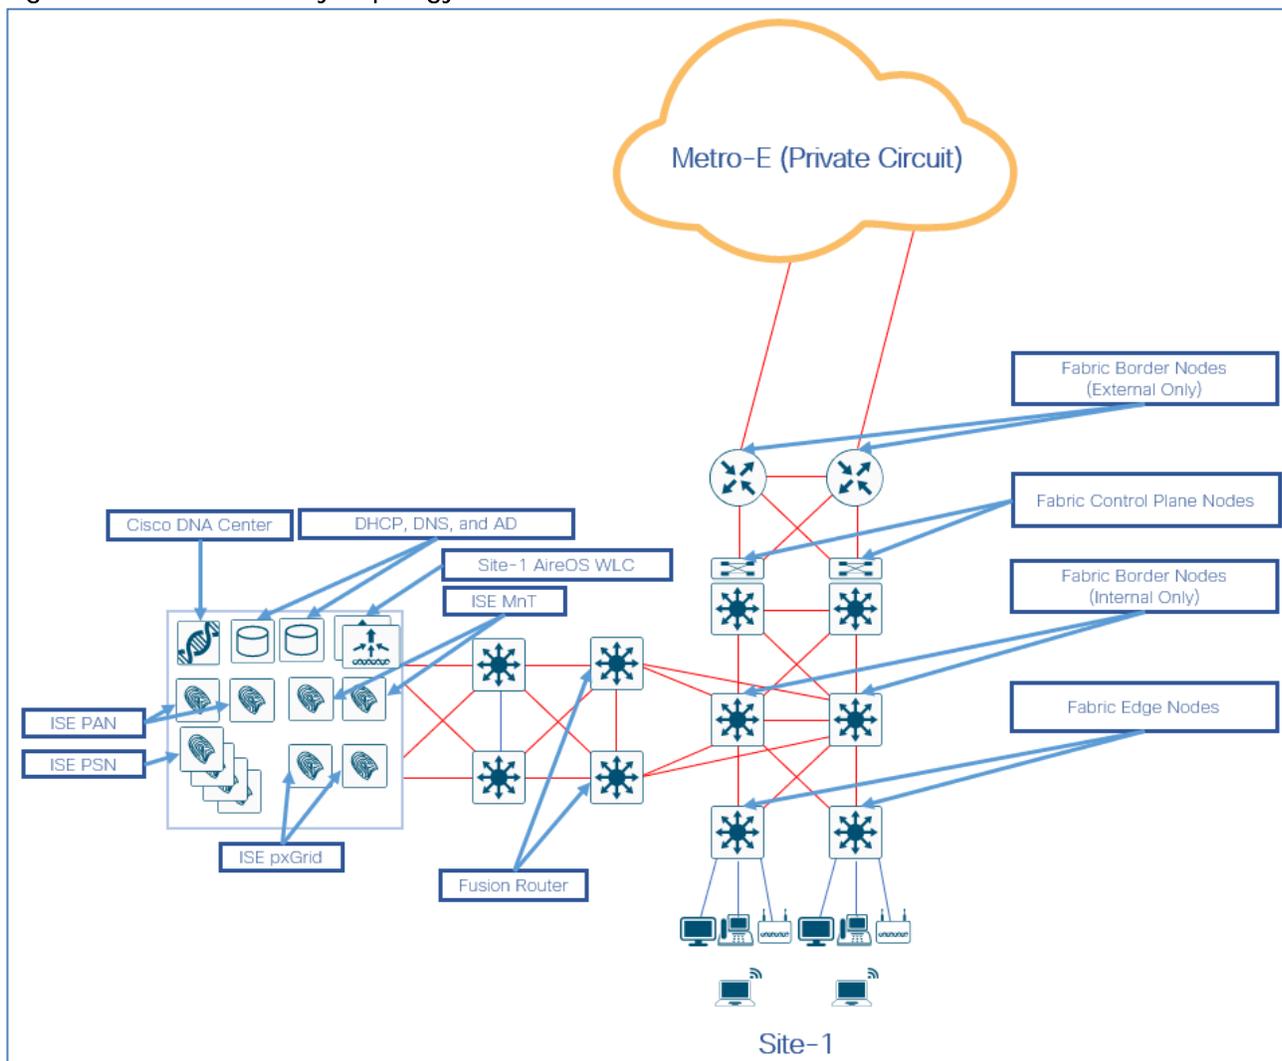
There are no restrictions regarding External, Internal, and Anywhere along with **Add as Border** or **Add as CP+Border**. For example, a device (platform dependent) can operate as an Internal-Only border node with the control plane node function on the same device or any other variation of these selections. Please note the pop-up **Warning** shown in steps below for additional details.

Procedure 1: Provisioning a Site Fabric Overlay

The following steps will provision the control plane nodes, edge nodes, and the WLC in fabric roles in Site-1. The (external) border nodes connected to the SD-Access transit will also be provisioned. In the deployment, two devices in Site-1 are (internal) border nodes connecting to the Data Center through the fusion routers. Their provisioning is addressed in the next procedure.

Site-1 through Site-4 will use independent, dedicated devices in the control plane node roles and the border node roles. Site-5 will use the control plane node and border node roles operating on the same devices – sometimes referred to as collocated. Please see [Appendix C](#) for additional and advanced topology diagrams.

Figure 14 Site-1 Overlay Topology



Tech Tip

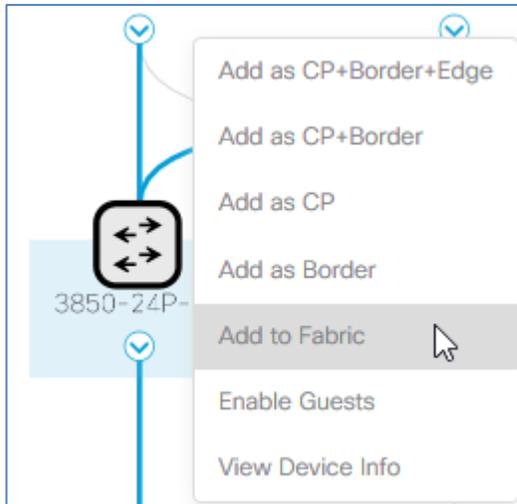
If combining both fabric roles on a device, please select **Add as CP+Border** where applicable in the workflows.

Dedicating a device to control plane node function and border node function is not required. However, there are scale implications when collocating these roles depending on what other routing and switching functions are being performed on the border node. There are also packet forwarding and processing considerations if the border node is also the core switch of the network.

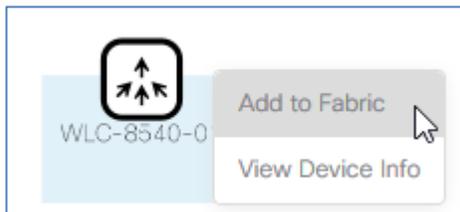
Core switches have traditionally been [designed](#) to provide the necessary scalability, load sharing, fast convergence, and high-speed capacity with minimal packet processing using the **less is more approach**. Broadly speaking, core switches should have minimum configuration wherever possible to meet these design requirements.

Use the appropriate planning, preparation, and piloting if using a core switch in a fabric border node or collocated role.

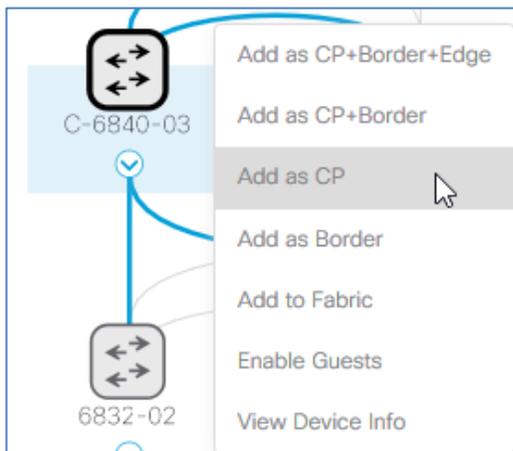
1. In the Cisco DNA Center dashboard, navigate to **PROVISION > Fabric**.
2. Select the **Fabric Domain** (SJC).
3. Select the **Fabric Site** (Site-1) from the **Fabric-Enabled Sites** on the right.
The topology map loads showing the devices provisioned to that site.
4. In the fabric topology view, click a node that will be used as a fabric edge node.
5. Select **Add to Fabric** from the popup menu.



6. Repeat this step for additional devices to be used as fabric edge nodes.
7. Click a Wireless LAN Controller.
8. Select **Add to Fabric** from the popup menu.

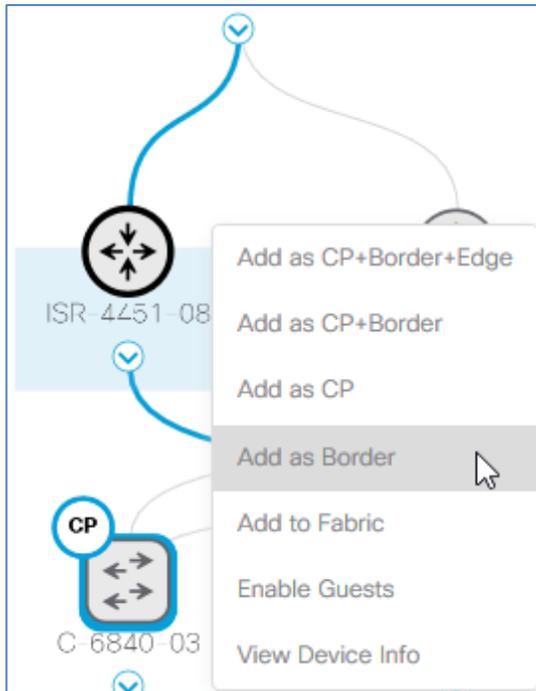


9. Click a device that will operate as a control plane node without border node functionality.
10. Select **Add as CP** from the popup menu.

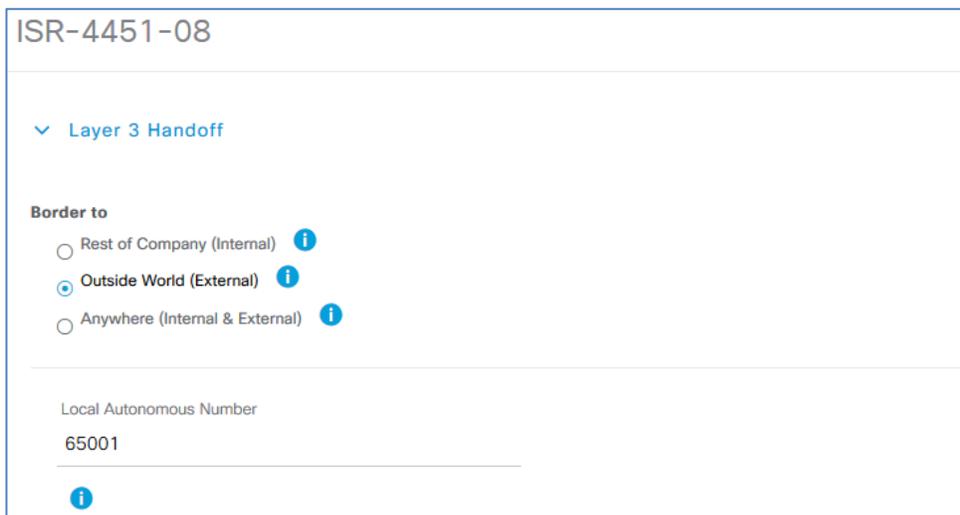


11. Repeat this step for any redundant, dedicated control plane node without border functionality.
12. Click a device that will operate as the external fabric border node.
This device should be connected to the metro-area private circuit and not to the fusion routers.

13. Select **Add as Border** from the popup menu.
A slide-in pane appears.



14. As the border is connected to unknown routes or the metro transit, select **Outside World (External)** in the slide-in pane.
15. Supply the **Local BGP Autonomous System Number** (example: 65001).



16. Under **Transit**, select the previously created **SDA: Transit** (example: METRO-E-SDA).
17. Click **Add**.
The transit is added.

▼ **Transits**

SDA: METRO-E-SDA ▼

Add

18. Click the **blue Add** Button at the bottom of the slide-in pane.

ISR-4451-08

▼ **Layer 3 Handoff**

Border to

Rest of Company (Internal) ⓘ

Outside World (External) ⓘ

Anywhere (Internal & External) ⓘ

Local Autonomous Number

65001

ⓘ

Select IP Address Pool ▼

ⓘ

Is this site connected to Internet?

▼ **Transits**

SDA: METRO-E-SDA ▼

Add

> **Layer 2 Handoff**

Cancel **Add**

19. A **Warning** pop-up provides additional information regarding fabric border nodes. Click **OK**.



Warning

Border to “Outside World (External)” or “Anywhere (Internal & External)” :

- With IP Transit : always enable the Connected to Internet option.
- With SDA Transit or SDA and IP Transits: enable the Connected to Internet if this site provides internet connectivity to other sites.
- Without any Transits, enable the Connected to Internet option if this site connects to the internet.

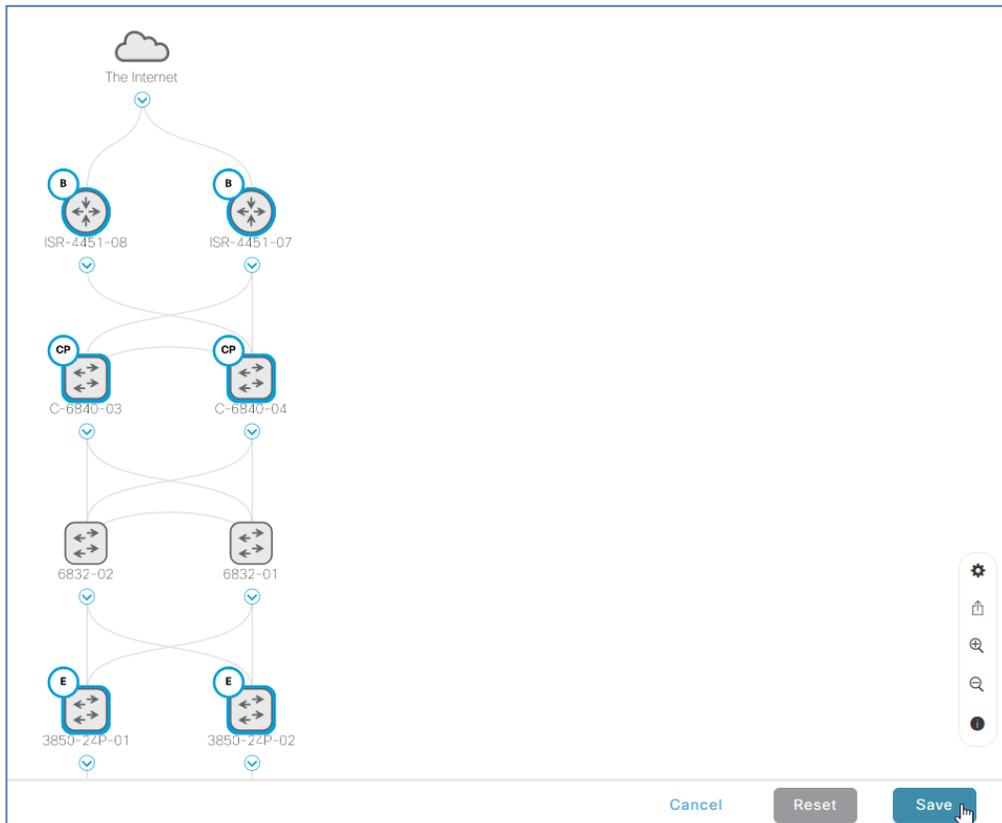
Border to “Rest of Company (Internal)” :

- Always disable the Connected to Internet option.

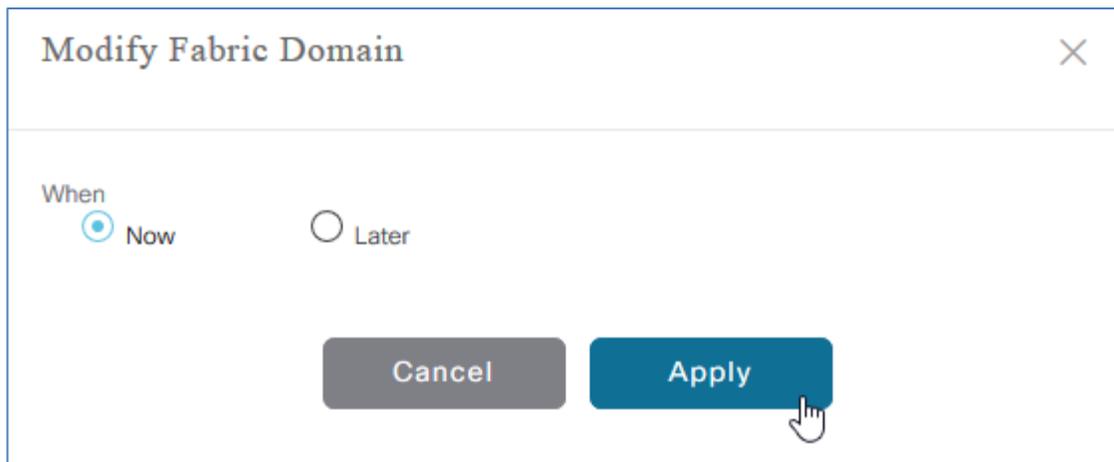
Cancel OK

20. Repeat these steps for any other **External-Only** border nodes in the site.
21. Verify that the devices are highlighted in **blue** and have a fabric-role label (example: B, CP, E). This indicates the *intent* to provision these devices to the fabric.

In the bottom right, click **Save**.



22. In the **Modify Fabric Domain** slide-in pane, leave the default selection of **Now**.
23. Click **Apply**.



Tech Tip

The fabric topology maps display the configuration and provisioning *intent*.

They do not represent or express the configuration and provisioning *state*.

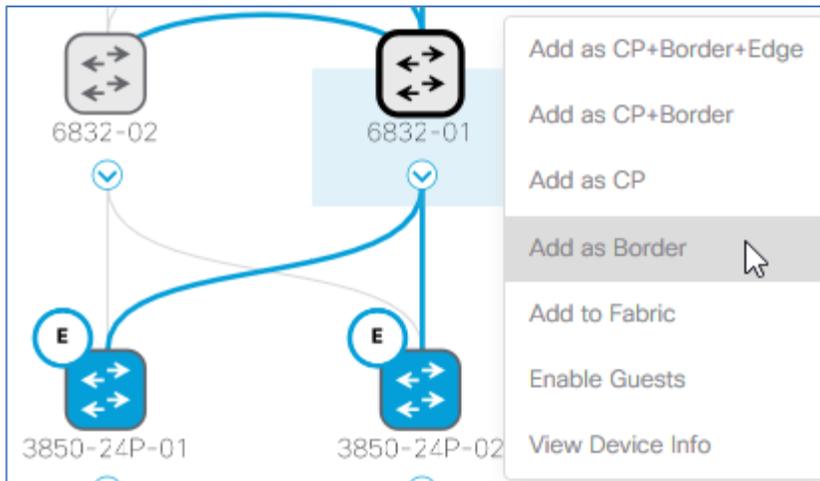
Procedure 2: Provisioning Internal Border Nodes

Internal border nodes connect to known routes in the network such as the Data Center. Internal border nodes are most often used to provide access to shared services networks such as DHCP and DNS. Most commonly, they are directly connected to another device such as a fusion router (described in the next section). Internal border nodes vary from external border nodes in that they register prefixes to the host tracking database (control plane node) in a similar way that fabric edge nodes register prefixes and endpoints. This registration notifies the rest of the fabric site and ultimately the fabric domain that these prefixes are accessible through this internal border node.

Because these internal border nodes are directly connected to the fusion routers, they will utilize the IP-based transit (DATA CENTER) previously created. This will instruct Cisco DNA Center to provision a Layer-3 handoff configuration using VRF-lite for use in later steps.

The following steps will provision internal borders with a Layer-3 handoff which extends the fabric VNs to the next hop fusion router. This will allow the endpoints in the fabric to access shared services once the fusion router configuration is completed in the [Operate](#) section.

1. In the Cisco DNA Center dashboard, navigate to **PROVISION > Fabric**.
2. Select the **Fabric Domain (SJC)**.
3. Select the **Fabric Site (Site-1)** from the **Fabric-Enabled Sites** on the right. The topology map loads showing the devices provisioned to that site.
4. Select the device to be used as an internal border.
5. Select **Add as Border** from the popup menu. A slide-in pane appears.



6. As the border is connected only to known routes, select **Rest of Company (Internal)**.
7. Enter the **Local BGP Autonomous System Number** (example: 65001).

Layer 3 Handoff

Border to

Rest of Company (Internal) i

Outside World (External) i

Anywhere (Internal & External) i

Local Autonomous Number

65001

8. Under **Select IP Address Pool**, select the previously defined pool to be used for the Layer-3 handoff (example: BGP-SITE-01).
9. Under **Transit**, select the previously created IP transit (example: DATA CENTER).
10. Click **Add**.
The transit is added.

Layer 3 Handoff

Border to

Rest of Company (Internal) *i*

Outside World (External) *i*

Anywhere (Internal & External) *i*

Local Autonomous Number

65001

Select IP Address Pool

* BGP-SITE-01 (172.16.111.0/24) *i*

Is this site connected to Internet?

Transits

IP: DATA CENTER *i* Add

11. Click the IP transit name to expand it

> DATA CENTER *i*

12. Click + Add Interface.

DATA CENTER *i*

External Interface *i* + Add Interface

13. In the new slide-in pane, select the **External Interface** connected to the first fusion router (example: TenGigabitEthernet1/11).
14. The **BGP Remote AS Number** is automatically filled in due to the creation of the IP-based transit. Click ^ to expand the **Virtual Network** section.
15. Select all **Virtual Networks** that will require access to shared services.
16. Click **Save**.

6832-01

[< Back](#)

External Interface

* TenGigabitEthernet1/11 ▼

Remote AS Number

65333

▼ Virtual Network i

- DEFAULT_VN
- INFRA_VN
- CAMPUS
- RESEARCH
- GUEST

Cancel Save

17. Verify that the desired **External Interfaces** and **Number of VNs** are displayed.
Click **Add** at the bottom of the slide-in pane.

 Tech Tip

The INFRA_VN is described in the next process. It is associated with the global routing table – it is not a VRF definition – and is used by Access Points and Extended Nodes. If these devices require DHCP, DNS, and other shared services, the INFRA_VN should be selected under **Virtual Network**.

6832-01

▼ Layer 3 Handoff

Border to

Rest of Company (Internal) ⓘ

Outside World (External) ⓘ

Anywhere (Internal & External) ⓘ

Local Autonomous Number

65001

ⓘ

Select IP Address Pool

* BGP-SITE-01 (172.16.111.0/24) ▼

ⓘ

Is this site connected to Internet?

▼ Transits

IP: DATA CENTER ▼ Add

▼ DATA CENTER 🗑️

External Interface ⓘ + Add Interface

Interface	Number of VN	
TenGigabitEthernet1/11	4	Remove

Cancel Add

18. Click **OK** in the **Warning** pop-up.
19. Repeat these steps for the redundant internal border node.
20. Once complete, click **Save** in the bottom right of the fabric topology screen.
21. In the **Modify Fabric Domain** slide-in pane, leave the default selection of **Now**.
22. Click **Apply**.

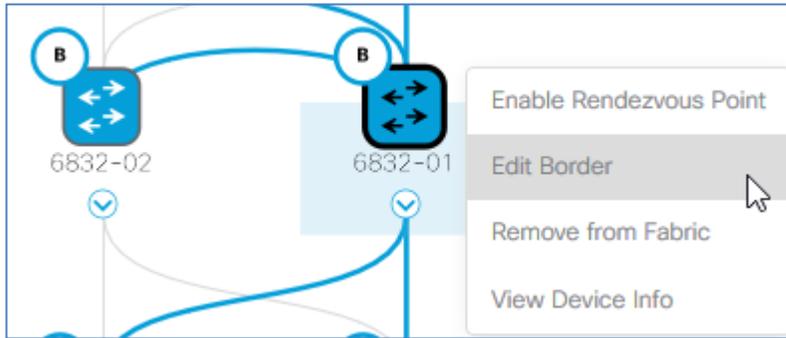
Procedure 3: Provisioning Redundant Layer-3 Handoffs

As a [redundancy](#) strategy, critical network devices are deployed in pairs. Fusion routers are no exception. These too are commonly deployed as a pair of devices to avoid single points of failure in hardware. To avoid single points of failure in uplinks to the fusion routers, fabric border nodes are connected to each of the fusion routers through separate interfaces. The concept is to [build triangles, not squares](#) to produce the highest resiliency and deterministic convergence.

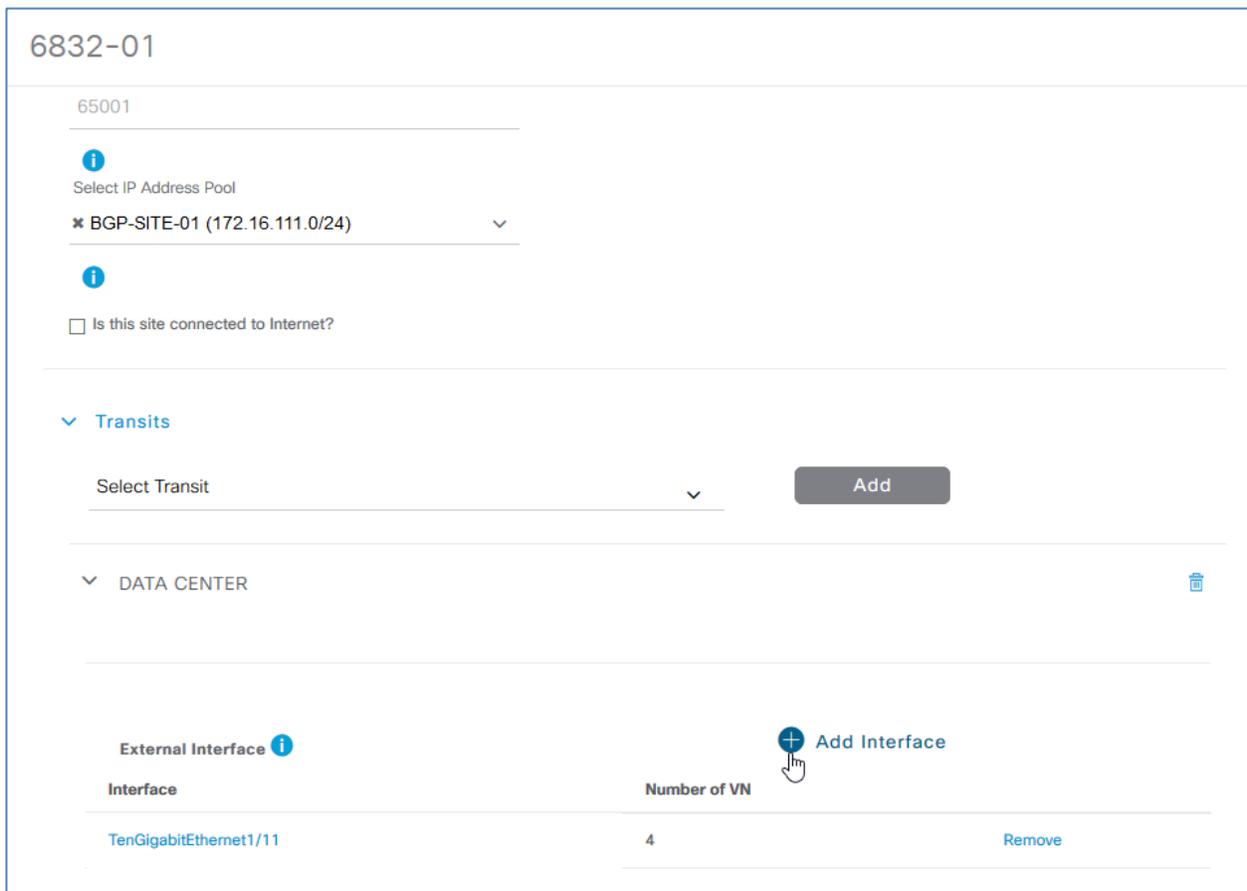
In the previous procedure, a single interface was provisioned for the Layer-3 handoff to a single fusion router on each internal border node. The following steps show the procedure to provision a second Layer-3 handoff interface on each

border node. This second interface is connected to the redundant upstream fusion router.

1. From the fabric topology map for Site-1, click the previously provisioned internal border node.
2. Select **Edit Border** in the pop-up menu.



3. Under **Transits**, click > to expand the previously defined IP transit.
4. The previously defined **External Interface** and **Number of VNs** is displayed. Click **+ Add Interface**.



5. In the new slide-in pane, select the **External Interface** connected to the second fusion router (example: TenGigabitEthernet1/12).
6. The **BGP Remote AS Number** is automatically filled in due to the creation of the IP-based transit. Click ^ to expand the **Virtual Network** selection.
7. Select all **Virtual Networks** that will require access to shared services.
8. Click **Save**.

6832-01

< Back

External Interface

* TenGigabitEthernet1/12

Remote AS Number

65333

Virtual Network 

DEFAULT_VN

INFRA_VN

CAMPUS

RESEARCH

GUEST

Cancel

Save

9. Verify that the second **External Interface** and **Number of VNs** are displayed along with the previously provisioned **External Interface**.
Click **Add**.

6832-01

Border to

Rest of Company (Internal) i
 Outside World (External) i
 Anywhere (Internal & External) i

Local Autonomous Number
65001

i
Select IP Address Pool
* BGP-SITE-01 (172.16.111.0/24) v

i
 Is this site connected to Internet?

v **Transits**

Select Transit v Add

v DATA CENTER 🗑️

External Interface i + Add Interface

Interface	Number of VN	
TenGigabitEthernet1/11	4	Remove
TenGigabitEthernet1/12	4	Remove

Cancel Add

10. Click **OK** in the **Warning** pop-up.
11. Repeat these steps for the second internal border node's interface connected to the second fusion router.
12. Once complete, click **Save** in the bottom right of the fabric topology screen.
13. In the **Modify Fabric Domain** slide-in pane, leave the default selection of **Now**.
14. Click **Apply**.

Process 8: Configuring Host Onboarding – Fabric Provisioning Part 4

Host onboarding is the culmination of all the previous steps. It binds the reserved IP address pools from the Design application with the VN configured in the Policy application, and provisions the remainder of the fabric configuration down to the devices operating in a fabric role.

Host onboarding is comprised of four distinct steps – all located under the **Provision > Fabric > Host Onboarding** tab for a fabric site.

1. Define authentication template.
2. Create host pools.
3. Define SSID address pool.
4. Assign Access Point ports.

The first step is to select the authentication template. These templates are predefined in Cisco DNA Center and are pushed

down to all devices that are operating as edge nodes within a site. It is mandatory to complete this step first – an authentication template must be defined before host pool creation.

The second step is to bind the IP address pools to the Virtual Networks (VNs). Once bound, these components are referred to collectively as *host pools*. Multiple IP address pools can be associated with the same VN. However, an IP address pool must not be associated with more than one VN. Doing so would allow communication between the VNs and break macro-segmentation in SD-Access.

When deploying SD-Access Wireless, the third step is required. This step binds an SSID to an IP address pool. The last step is to selectively modify the authentication template at the port level so that the Access Points can connect to the network. This step is required if using 802.1x in the deployment.

About Authentication Templates

Authentication templates deploy certain interface-level commands on the downstream (access) ports of all edge nodes. By default, access ports are considered all copper ports operating as switchports (as opposed to routed ports) on the edge nodes. These interface-level commands are port-based Network Access Control configurations that validate a connected endpoint before it can connect to a network.

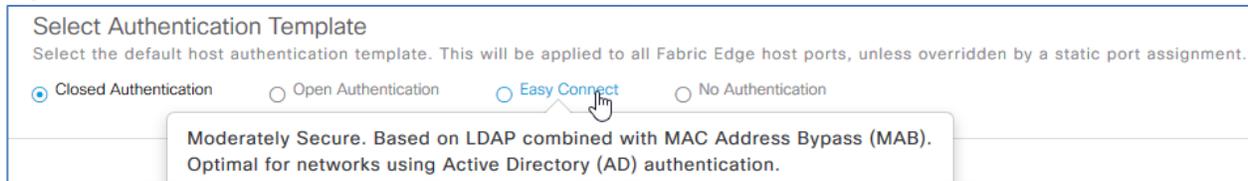
There are four supported built-in authentication templates:

1. Closed Authentication
2. Open Authentication
3. Easy Connect
4. No Authentication

These templates are based on the AAA [Phased Deployment Implementation Strategy](#) of High Security mode, Low Impact mode, Monitor mode, and No Authentication mode.

Hovering over each option in Cisco DNA Center provides information about the authentication template.

Figure 15 Cisco DNA Center Authentication Templates



Tech Tip

While Cisco Access Points contain an 802.1x supplicant and can connect to the network using 802.1x, this connectivity is not currently supported in SD-Access. Therefore, if one of the first three authentication templates is used, the ports on the edge nodes that connect to the Access Points must be modified to use No Authentication. This will become clearer in later steps.

Procedure 1: Assign Authentication Template and Create Wired Host Pool – Host Onboarding Part 1

When a **host pool** is created, a subnet in the form of a reserved IP address pool is bound to a VN. From the perspective of device configuration, Cisco DNA Center creates the VRF definition on the fabric nodes, creates an SVI or loopback interface (on switches and routers, respectively), defines these interfaces to forward for the VRF, and gives it the IP address defined as the gateway for the reserved pool.

When creating the host pools, there are multiple options including Traffic Type, Layer-2 Extension, Layer-2 Flooding, Groups, Critical Pool, and Auth Policy. Many are beyond the scope of this document although are listed here for reference.

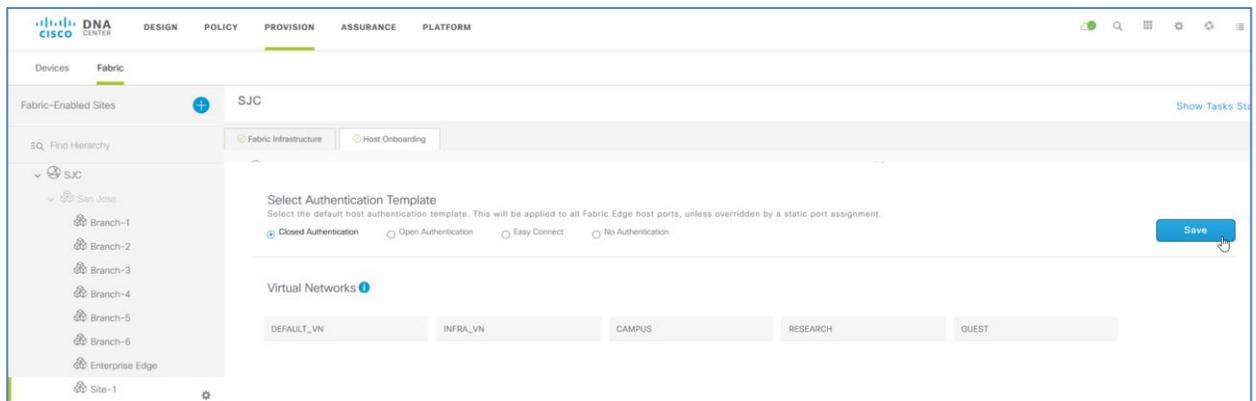
Figure 16 Host Pool Configuration Options

IP Pool Name	Traffic Type	Address Pool	Layer-2 Extension	Layer-2 Flooding	Groups	Critical Pool	Auth Policy
--------------	--------------	--------------	-------------------	------------------	--------	---------------	-------------

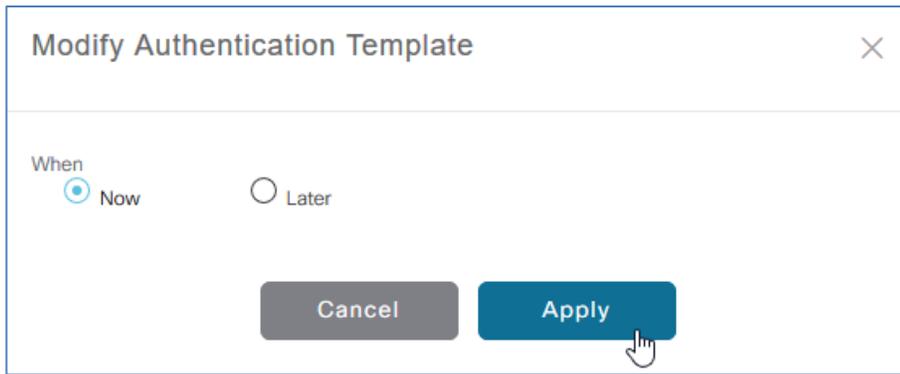
Traffic Type is defined as either **Voice** or **Data** which applies a QoS configuration. **Layer-2 Extension** should **always** be enabled for any post >1.1.8 deployment. For additional details on the remaining options please see the Software-Defined Access CVDs and additional references in [Appendix B](#).

The following steps will define the **Closed Authentication Template** for the Site-1 fabric and create host pools for the wired VNs.

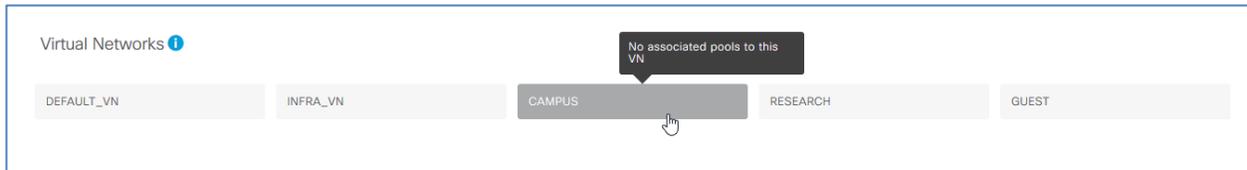
1. In the Cisco DNA Center dashboard, navigate to **PROVISION > Fabric**.
2. Select the **Fabric Domain (SJC)**.
3. Select the **Fabric Site (Site-1)** from the **Fabric-Enabled Sites** on the right.
4. Click the **Host Onboarding** tab.
5. Under **Select Authentication Template**, select **Closed Authentication** and click the **Save** button on the right.



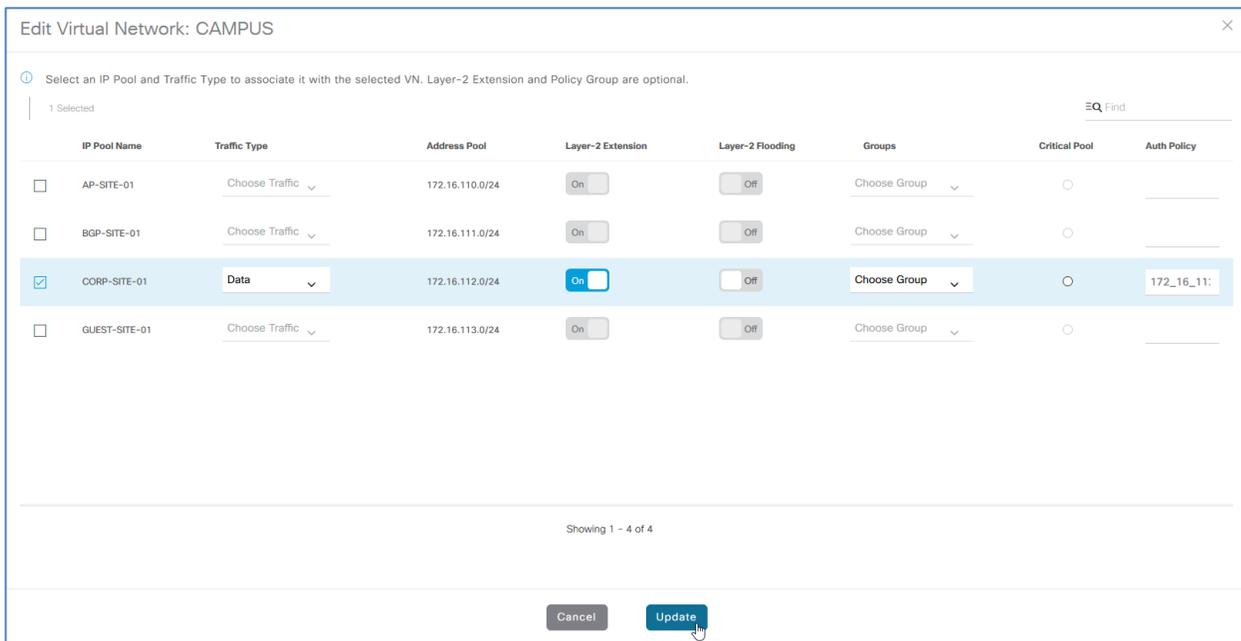
6. In the **Modify Authentication Template** slide-in pane, leave the default selection of **Now**.
7. Click **Apply** to demonstrate the intent to use the desired template.
The configuration will actually be provisioned at the end of this procedure.



- Under **Virtual Networks**, select a VN to be used for wired clients (example: CAMPUS).



- In the **Edit Virtual Network: CAMPUS** slide-in pane, select the **IP Pool Name(s)** to associate it to the VN (example: CORP-SITE-01).
- Select a **Traffic Type** of **Data**.
- Verify that Layer 2 Extension is **On**.
- Click **Update**.



- In the **Modify Virtual Network** slide-in pane, leave the default selection of **Now**.
- Click **Apply**.
- Repeat these steps with the remaining user-defined VNs and IP address pools in the site to create wired host pools.

Procedure 2: Create *INFRA_VN* Host Pool – Host Onboarding Part 2

Access Points are a special case in the fabric. They are connected to edge nodes like an endpoint, although they are actually part of the fabric infrastructure. Because of this, their traffic pattern is unique. Access Points receive a DHCP address via the overlay network and associate with the WLC via the underlay network. Once associated with the WLC, they are registered

with Cisco DNA Center by the WLC through the overlay network. To accommodate this traffic flow, the Access Point subnet – which is in the Global Routing Table (GRT) – is associated with the overlay network. Cisco DNA Center GUI calls this special overlay network associated with the GRT the *INFRA_VN*.

INFRA_VN stands for *Infrastructure Virtual Network* and is intended to represent devices that are part of the network infrastructure, but associate and connect to the network in a similar method to endpoints - directly connected to the downstream ports of an edge node. Both Access Points and Extended Nodes (SD-Access Extension for IoT) are part of the *INFRA_VN*.



INFRA_VN, while labeled as a *Virtual Network* in the GUI because it is part of the overlay network, is not provisioned as a VRF definition in the network. It is not truly a VRF and routes for the *INFRA_VN* are in the Global Routing Table. This is an important, although subtle distinction, particularly with devices such as the Catalyst 9200 that have a smaller number of supported VRFs.

The distinction is also important because of the unique and asymmetric traffic flow. Additional information on AP traffic flow can be found in [SD-Access Wireless Design and Deployment Guide](#).

1. In the Cisco DNA Center dashboard, navigate to **PROVISION > Fabric**.
2. Select the **Fabric Domain (SJC)**.
3. Select the **Fabric Site (Site-1)** from the **Fabric-Enabled Sites** on the right.
4. Click the **Host Onboarding** tab.
5. Under Virtual Networks, select *INFRA_VN* to be used as the Access Point host pool.
6. In the **Edit Virtual Network: INFRA_VN** slide-in pane, select the **IP Pool Name(s)** to associate it to the VN (example: AP-SITE-01).
7. Select **AP** as the **Pool Type**.
8. Verify that Layer-2 Extension is **On**.
This **must** be **On** for the *INFRA_VN*.
9. Click **Update**.

1 Selected

Select an IP Pool and Traffic Type to associate it with the selected VN. Layer-2 Extension and Policy Group are optional.

IP Pool Name	Address Pool	Pool Type	Layer-2 Extension	Layer-2 Flooding
<input checked="" type="checkbox"/> AP-SITE-01	172.16.110.0/24	<input checked="" type="radio"/> AP <input type="radio"/> EXTENDED (BETA)	<input type="checkbox"/> On <input type="checkbox"/> Off	<input type="checkbox"/> On <input checked="" type="checkbox"/> Off
<input type="checkbox"/> BGP-SITE-01	172.16.111.0/24	<input checked="" type="radio"/> AP <input type="radio"/> EXTENDED (BETA)	<input type="checkbox"/> On <input type="checkbox"/> Off	<input type="checkbox"/> On <input checked="" type="checkbox"/> Off
<input type="checkbox"/> CORP-SITE-01	172.16.112.0/24	<input checked="" type="radio"/> AP <input type="radio"/> EXTENDED (BETA)	<input type="checkbox"/> On <input type="checkbox"/> Off	<input type="checkbox"/> On <input checked="" type="checkbox"/> Off
<input type="checkbox"/> GUEST-SITE-01	172.16.113.0/24	<input checked="" type="radio"/> AP <input type="radio"/> EXTENDED (BETA)	<input type="checkbox"/> On <input type="checkbox"/> Off	<input type="checkbox"/> On <input checked="" type="checkbox"/> Off
<input type="checkbox"/> RESEARCH-SITE-01	172.16.114.0/24	<input checked="" type="radio"/> AP <input type="radio"/> EXTENDED (BETA)	<input type="checkbox"/> On <input type="checkbox"/> Off	<input type="checkbox"/> On <input checked="" type="checkbox"/> Off

Showing 1 - 5 of 5

10. In the **Modify Virtual Network** slide-in pane, leave the default selection of **Now**.

11. Click **Apply**.

Procedure 3: Assign SSID Subnet and Assign Access Point Ports- Host Onboarding Part 3

Each SSID for a fabric site must be assigned an IP address pool so that wireless hosts are associated with the correct subnet when connecting to the wireless network. This places them in the client overlay network, as the *INFRA_VN* is the AP overlay network.

For the Access Points themselves to connect to the network, the authentication template must be modified at the port level. Cisco DNA Center enables automatic onboarding of APs by provisioning a Cisco Discovery Protocol (CDP-based) macro at the fabric edge nodes when the authentication template is set to **No Authentication**. If a different authentication template is used globally, for example **Closed Authentication**, then the downstream switchport configurations on the edge nodes must be changed in Cisco DNA Center.

The following steps will assign an IP address pool for the Enterprise and Guest SSIDs and modify the authentication template used at the edge node port level to support Access Port connectivity.

1. In the Cisco DNA Center dashboard, navigate to **PROVISION > Fabric**.
2. Select the **Fabric Domain (SJC)**.
3. Select the **Fabric Site (Site-1)** from the **Fabric-Enabled Sites** on the right.
4. Click the **Host Onboarding** tab.
5. Under **Wireless SSID's**, select the **IP Address Pool** for the Guest SSID (example: GUEST:172.16.113.0).
6. Repeat this step for the Enterprise SSID (example: CAMPUS:172.16.112.0).
7. Click the **Save** button on the right.

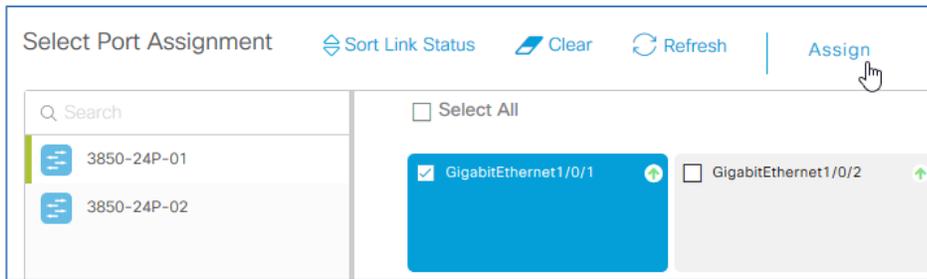
SSID Name	Type	Security	Traffic Type	Address Pool	Scalable Group
GUEST	Guest	Web Auth	Voice + Data	GUEST:172.16.113.0	
CORP	Enterprise	WPA2 Enterprise	Voice + Data	CAMPUS:172.16.112.0	

8. In the **Modify SSID Table** slide-in pane, leave the default selection of **Now**.
9. Click **Apply**.

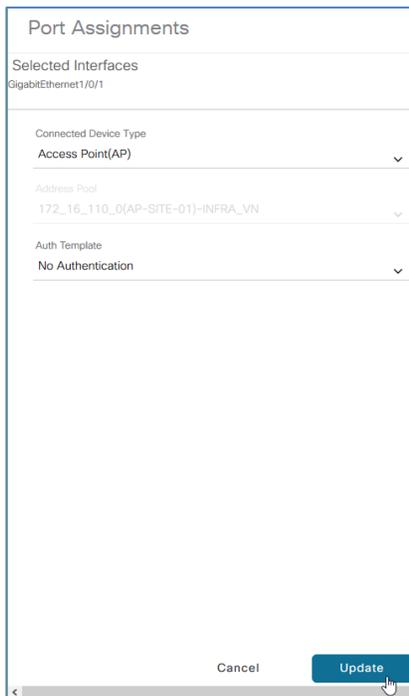
When
 Now Later

Cancel Apply

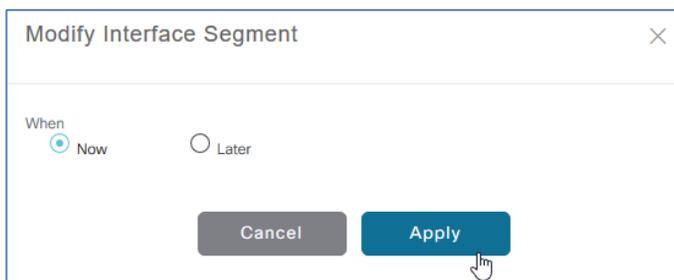
10. Under **Select Port Assignment**, click an edge node (example: 3850-24P-01).
11. Select the check box(es) of the interface(s) to which an Access Point is connected (example: GigabitEthernet1/0/1).
12. Click **Assign**.



13. In the **Port Assignments** slide-in pane, under **Connected Device Type**, select **Access Point(AP)**.
14. Under **Auth Template**, select **No Authentication**.
The **Address Pool** will automatically be filled in based on the *INFRA_VN* host pool.
15. Click **Update**.



16. In the slide-in pane for **Modify Interface Segment**, leave the default selection of **Now**.
17. Click **Apply**.



When modifying the port assignment on multiple edge nodes, the save button must be clicked between selecting each switch.

Select Port Assignment Sort Link Status Clear Refresh Assign Save

Q Search

- 3850-24P-02.dna.local
- 3850-24P-01.dna.local

Select All

<input type="checkbox"/> GigabitEthernet1/0/1 Device-Type: ACCESS_POINT Segment: 1T2_1R_110_0-INFRA_VN Authentication: No Authentication	<input type="checkbox"/> GigabitEthernet1/0/2	<input type="checkbox"/> GigabitEthernet1/0/3	<input type="checkbox"/> GigabitEthernet1/0/4	<input type="checkbox"/> GigabitEthernet1/0/5
---	---	---	---	---

Operate

The Operate section provides a comprehensive overview on providing shared services and Internet access for the SD-Access fabric domain. Much of this configuration is performed manually on the CLI of the fusion routers and Internet edge routers. While shared services are a necessary part of the deployment, they are, strictly speaking, outside of the fabric. Located outside of the fabric, these devices will not receive automated provisioning from Cisco DNA Center.

The Operate section begins with a discussion on access to shared services using a fusion router along with corresponding information on fusion router requirements and configuration. Once access to shared services is provided to the devices downstream from the edge nodes, Access Points can receive an IP address via DHCP with [DHCP Option 43](#) pointing to their site-local wireless LAN controller. The Access Points will be registered (brought into the inventory) in Cisco DNA Center by the WLC. Once in the inventory, they can be provisioned with the specific wireless configuration created during the [Design application](#) steps.

Redundant border nodes have special considerations when it comes to the interaction between BGP and LISP. To assist with failover scenarios, IBGP sessions are manually configured between these redundant devices. As the configuration has variations on routers and switches due to subinterface support, both options are discussed. This section then provides a prescriptive configuration for providing Internet access to the fabric domain and supplies details and considerations regarding [connected-to-Internet](#) option in the border provisioning steps.

Process 1: Providing Access to Shared Services

When creating the fabric overlay through the end-to-end workflow, VNs are defined and created. When these are provisioned to the fabric devices, they are configured as VRF definitions. In the later host onboarding steps, each VRF is associated with one or more subnets defined during the Design application steps. The result is that the endpoints in the fabric overlay are associated and route packets through VRF tables. This is also the first layer of segmentation.

Shared services such as DHCP and DNS are generally accessible through the Global Routing Table (GRT) or could be in their own dedicated VRF, depending on the design. The challenge is to provide a method to advertise these shared services routes from the GRT/VRF to the VN routing tables on the border nodes so that endpoints in the fabric can access them. This is accomplished today using a fusion router.

About Fusion Routers

The generic term *fusion router* comes from [MPLS Layer-3 VPN](#). The basic concept is that the fusion router is aware of the prefixes available inside each VPN (VRF), either because of static routing configuration or through route peering, and can therefore fuse these routes together. A generic fusion router's responsibilities are to route traffic between separate VRFs (VRF leaking) or to route traffic to and from a VRF to a shared pool of resources in such as DHCP and DNS servers in the global routing table (route leaking). Both responsibilities involve moving routes from one routing table into a separate VRF routing table.

A fusion router in SD-Access has several technological requirements. It must support:

- Multiple VRFs
- 802.1q tagging (VLAN tagging)
- Subinterfaces (when using a router) or switched virtual interfaces (SVI) (when using a Layer-3 switch)
- BGPv4 and specifically the MP-BGP extensions ([RFC 4760](#) and [RFC 7606](#)) for extended communities attributes



While it is feasible to use a switch or a router as a fusion router, switches have additional considerations. Only certain high-end models support subinterfaces. Therefore, on a fixed configuration model such as a Catalyst 9300/3850, an SVI must be created on the switches and added to the VRF forwarding definition, and then Layer-2 trunks are used to connect to the border nodes.



Because the fusion router is outside the SD-Access fabric, it is not specifically managed (for automation) by Cisco DNA Center. Therefore, the configuration of a fusion router is manual. Future releases and developments may reduce or eliminate the need for a fusion router.

About Route Leaking

In an SD-Access deployment, the fusion router has a single responsibility: to provide access to shared services for the endpoints in the fabric. There are two primary ways to accomplish this task depending on how the shared services are deployed.

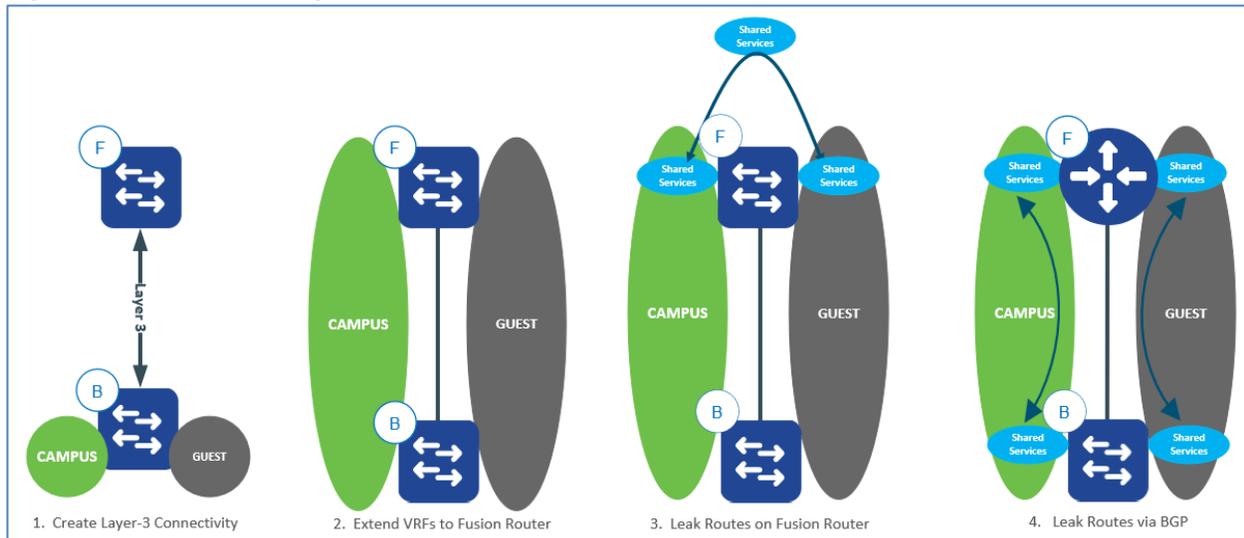
The first option is used when the shared services routes are in the GRT. IP prefix lists are used to match the shared services routes, route-maps reference the IP prefix lists, and the VRF configurations reference the route-maps to ensure only the specifically matched routes are leaked. This option, along with the reason it was not selected for this prescriptive guide, is discussed in [Appendix D](#). The second option is to place shared services in a dedicated VRF. With shared services in a VRF and the fabric endpoints in other VRFs, route-targets are used to leak between them.

Each option is viable and valid. For simplicity, this prescriptive deployment guide places the shared services in a dedicated VRF. However, production deployments may use one option or the other.

Access to shared services is a multi-step workflow performed primarily on the command-line interface of the fusion routers.

1. Create the Layer-3 connectivity between border nodes and fusion routers.
2. Use BGP to extend the VRFs from the border nodes and fusion routers.
3. Use route leaking or VRF leaking to share routes between the routing tables on the fusion router.
4. Distribute the leaked routes back to the border nodes via BGP.

Figure 17 Route Leaking Workflow

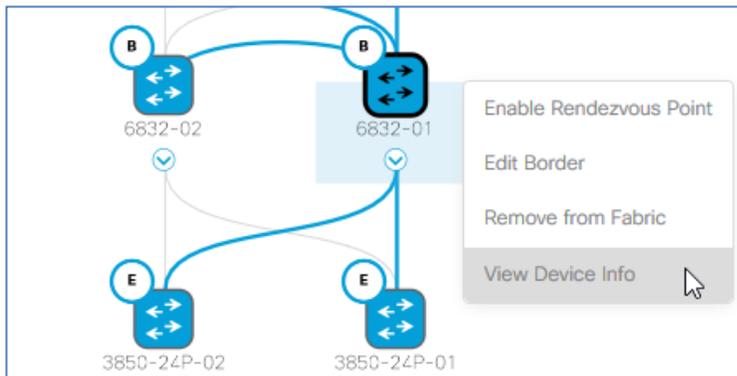


Procedure 1: Use the Provision Application for BGP and VRF-Lite information – Part 1

Cisco DNA Center can be used to determine the Layer-3 handoff configuration that was provisioned in previous steps. This information is needed to determine the corresponding configuration needed on the fusion routers.

1. In the Cisco DNA Center dashboard, navigate to **PROVISION > Fabric**.
2. Select the **Fabric Domain** (SJC).
3. Select the **Fabric Site** (Site-1) from the **Fabric-Enabled Sites** on the right. The topology map loads showing the devices provisioned to that site.
4. Click the internal border node previously provisioned with the Layer-3 VRF-lite handoff.

- In the pop-up menu, select **View Device Info**.



- In the slide-in pane, scroll to the listed **Interfaces** and click > to expand the information.
- The Layer-3 handoff provisioning information is displayed along with the **Local IPs** and necessary **Remote IPs**. Note this information as it will be used in the next procedures.
- Repeat these steps to collect the information on the redundant border node.

6832-01

GENERAL INFORMATION

Device Type: Cisco Catalyst C6832-X-LE Switch

Family: Switches and Hubs

Role: DISTRIBUTION

IP: 192.168.10.3

Software Version: 15.5(1)SY2

Border Type: INTERNAL

Border Handoff

Internal Domain Protocol Number: 65001

External Connectivity IP Pool: BGP-SITE-01

> TenGigabitEthernet1/12

▼ TenGigabitEthernet1/11

Layer3

External Domain Protocol: 65333

Virtual Network	Vlan	Local IP	Remote IP
RESEARCH-Global/San_Jose/Site-1	3007	172.16.111.25/30	172.16.111.26/30
GUEST-Global/San_Jose/Site-1	3008	172.16.111.29/30	172.16.111.30/30
INFRA_VN-Global/San_Jose/Site-1	3005	172.16.111.17/30	172.16.111.18/30
CAMPUS-Global/San_Jose/Site-1	3006	172.16.111.21/30	172.16.111.22/30

6832-01

GENERAL INFORMATION

Device Type: Cisco Catalyst C6832-X-LE Switch

Family: Switches and Hubs

Role: DISTRIBUTION

IP: 192.168.10.3

Software Version: 15.5(1)SY2

Border Type: INTERNAL

Border Handoff

Internal Domain Protocol Number: 65001

External Connectivity IP Pool: BGP-SITE-01

▼ TenGigabitEthernet1/12

Layer3

External Domain Protocol: 65333

Virtual Network	Vlan	Local IP	Remote IP
CAMPUS-Global/San_Jose/Site-1	3014	172.16.111.53/30	172.16.111.54/30
INFRA_VN-Global/San_Jose/Site-1	3013	172.16.111.49/30	172.16.111.50/30
RESEARCH-Global/San_Jose/Site-1	3015	172.16.111.57/30	172.16.111.58/30
GUEST-Global/San_Jose/Site-1	3016	172.16.111.61/30	172.16.111.62/30



During the Layer-3 border handoff automation, Cisco DNA Center uses VLSM on the defined IP address pool to create multiple /30 subnets.

Each subnet is associated with a VLAN beginning at 3001. Cisco DNA Center does not currently support the reuse of VLANs when a device is provisioned and un-provisioned. The VLAN number will continue to advance as demonstrated in the screen captures.

Procedure 2: Configure VRF Definitions on the Fusion Routers – Command Runner Tool

The next task is to create IP connectivity between the fusion routers and the border nodes. Cisco DNA Center has automated the configuration on the border nodes in previous steps. On each fusion router, IP connectivity must be created for each virtual network that requires connectivity to shared services, and this must be done for each interface between the fusion routers and redundant border nodes.

Like the chicken and egg dilemma, to configure an interface to be associated with a VRF instance, the VRF must first be created. As a best practice, the VRF configuration on the border nodes and fusion routers should be the same, including using the same route-target and route-distinguishers on both device sets.

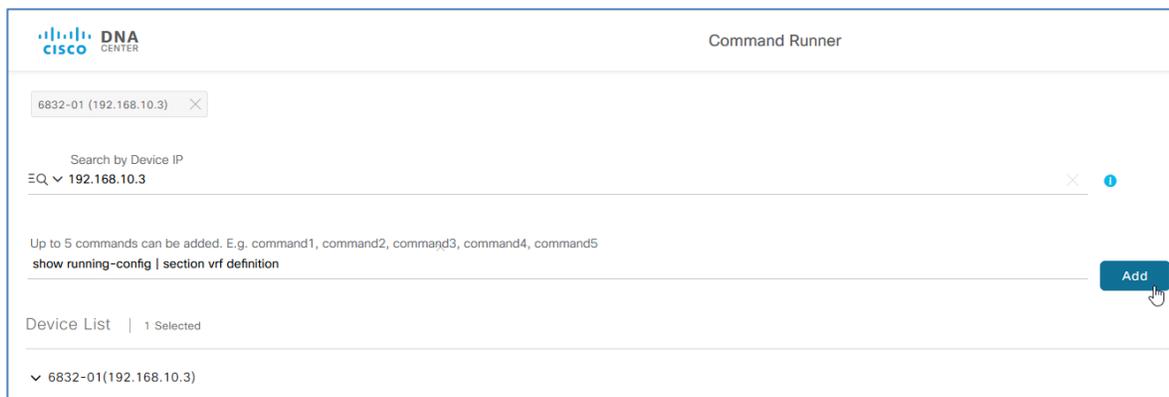


Often, for ease of administration, the route-target (RT) and route-distinguisher (RD) are configured as the same value, although this is not a requirement. It is simply a configuration convention that reduces an administrative burden and provides greater simplicity. This convention is used in the configurations provisioned by Cisco DNA Center. The RD and RT will also match the LISP Instance-ID.

The VRF configuration on the border nodes can be retrieved using the device's CLI or through the Command Runner tool in Cisco DNA Center.

To use Command Runner to display the border node VRF configuration:

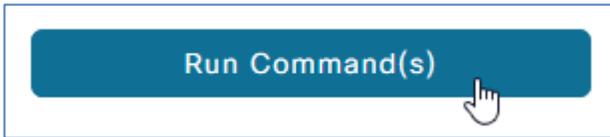
1. At the top-right of the Cisco DNA Center dashboard, navigate to **TOOLS > Command Runner**.
2. Select one of the borders with the Layer-3 IP handoff (example: 6832-01 (192.168.10.3)).
3. The device is added to the device list in the center of the screen.
In the command bar, type the command *show running-config | section vrf definition*.
4. Click the **Add** button.



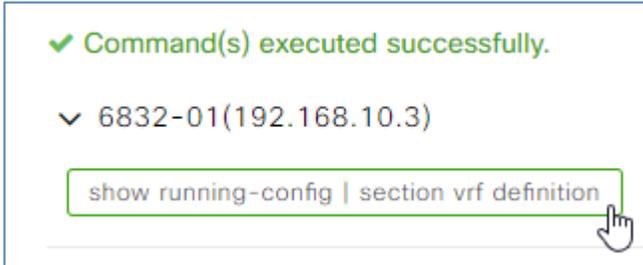
5. Verify that the command appears below the device in the **Device List**.



- Click the **Run Command(s)** button at the bottom of the screen.



- After a few moments the command is executed, and the results are returned to Cisco DNA Center. Click the command (shown with a green border to indicate that it was executed successfully) to view the result.



- Copy the displayed VRF configuration to be used for the next steps.



- Access the CLI of the fusion routers and enter global configuration mode.

```
configure terminal
```

- Paste the VRF configuration so that the RDs, RTs, and VRF names match exactly. Repeat these copy-and-paste steps for the redundant fusion router.



There are reserved values for the route-targets and route-distinguishers provisioned by Cisco DNA Center.

- **INFRA_VN** reserves 1:4097.
- **DEFAULT_VN** reserves 1:4098.

DEFAULT_VN is a predefined VN present from Cisco DNA Center 1.0 and is used primarily in the lab or if the original deployment used this VN and subsequently cloud updated to a modern code version. For this reason, it is not covered in this guide.

RT and RD values continue from 1:4099 and increase for each of the VRFs created and bonded into host pools.

Procedure 3: Create Layer-3 Connectivity Between Border Nodes and Fusion Routers – Fusion Routers Part I

Cisco DNA Center has provisioned a VRF-lite configuration on the internal border nodes. The following section will demonstrate how to create the corresponding VRF-lite configuration on the fusion router.



In this prescriptive deployment, an existing VRF SHARED_SERVICES is configured on the fusion routers. The interfaces facing the shared services are already configured to forward for the VRF definition.

```
vrf definition SHARED_SERVICES
  rd 1:4097
  address-family ipv4
    route-target export 1:4097
    route-target import 1:4097

interface Vlan51
  vrf forwarding SHARED_SERVICES
  ip address 198.51.100.11 255.255.255.0
```

Fusion routers can be either true routers or can be switching platforms. Both platforms support the role of fabric border as well. All modern Cisco Layer-3 switches support SVIs and trunks, although not all of these switches support subinterfaces. Routers must use subinterfaces as they do not support trunk ports. This results in an inherent difference in the possible configurations for VRF-lite – subinterfaces or SVIs and trunk ports.

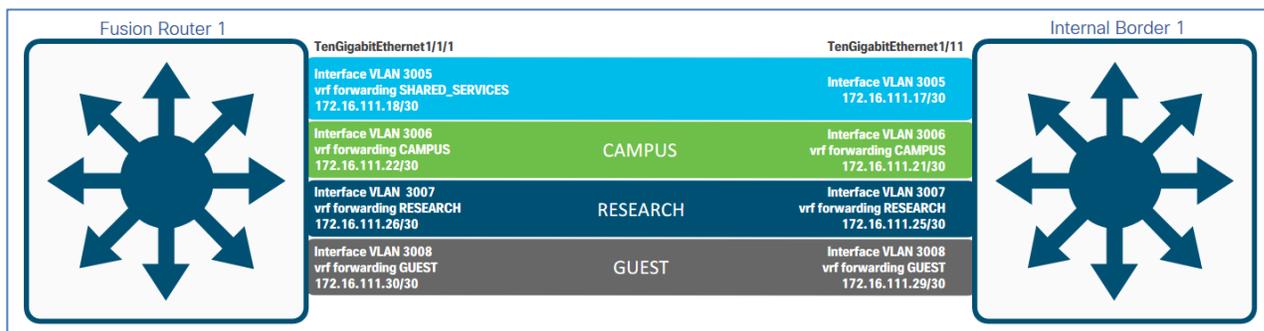
This deployment guide uses Layer-3 switches as the fusion routers. The internal border nodes in Site-1, as part of the distribution layer, are also Layer-3 switches. The VRF-lite configuration will therefore utilize SVIs and trunk ports to create the connectivity, extend the VRFs to the fusion routers, and exchange routing information.

The previous steps created the VRFs on the fusion routers. The next steps will use the IP address and VLAN information from the border node's [View Device Info](#) screen from Cisco DNA Center.

When the fusion router is a switching platform, the VLANs must be created first. Once created, the corresponding SVIs are created, configured to forward for a VRF, and given an IP address. Finally, a trunk port is created that forwards for the appropriate VLANs creating connectivity between the border node and fusion router.

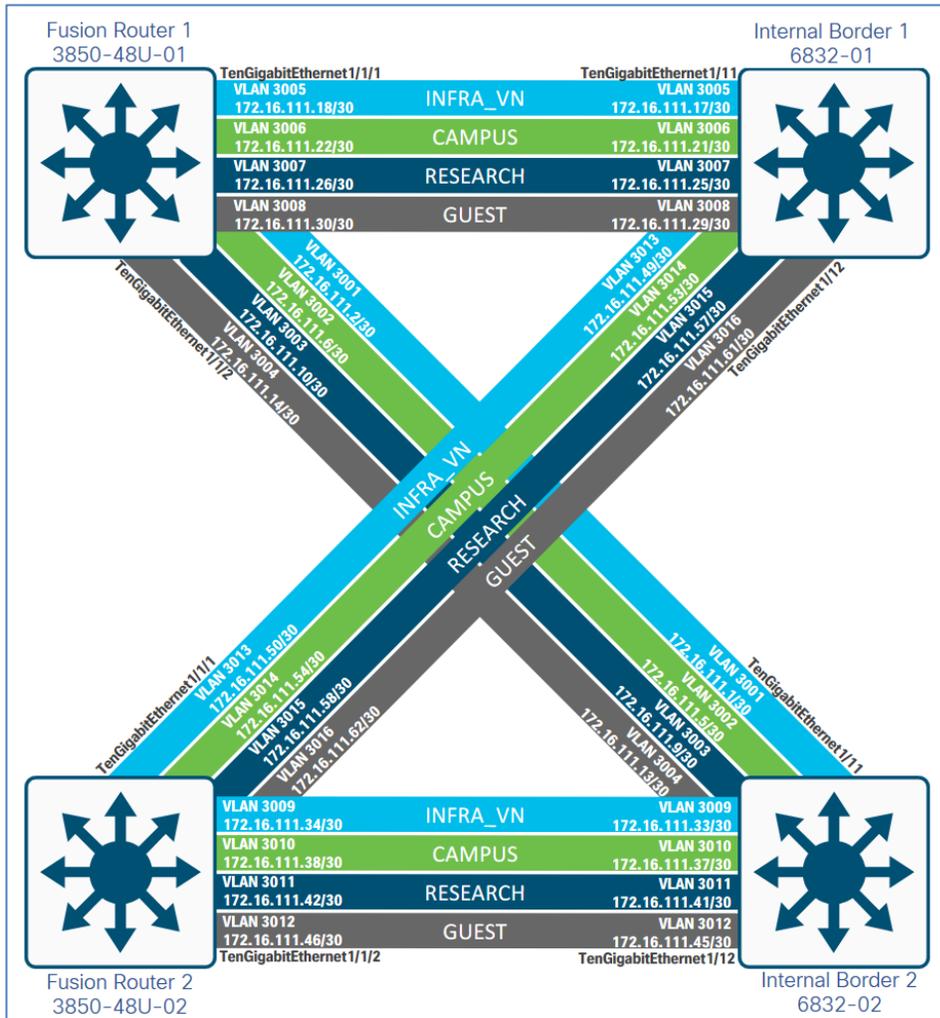


Cisco DNA Center has provisioned a VLAN and SVI on the border node that is forwarding for the GRT called *INFRA_VN* in the GUI. On the fusion router, the corresponding interface will be forwarding for predefined *SHARED_SERVICES* VRF.



Using the information gleaned from the *View Device Info* screen, the following connectivity will be created on the fusion routers in the next steps.

Figure 18 Fusion Router and Internal Border – VRF-lite



1. In global configuration on the first fusion router, create the VLANs.


```
vlan 3001-3008
exit
```
2. Configure each SVI to forward for a VRF and to have the applicable IP address based on the information from the *View Device Info* screen.

Remember to use *no shutdown* to bring the SVI to the up/up state.

```
interface vlan 3001
vrf forwarding SHARED_SERVICES
ip address 172.16.111.2 255.255.255.2
no shutdown

interface vlan 3002
vrf forwarding CAMPUS
ip address 172.16.111.6 255.255.255.252
no shutdown

interface vlan 3003
vrf forwarding RESEARCH
ip address 172.16.111.10 255.255.255.252
no shutdown
```

```
interface vlan 3004
  vrf forwarding GUEST
  ip address 172.16.111.14 255.255.255.252
  no shutdown
```

3. Configure a trunk interface that allows the defined VLANs.

```
interface TenGigabitEthernet1/1/2
  description Connected to Internal Border 6832-02 | TenGig 1/11
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 3001-3004
  no shutdown
```

4. Repeat these steps for the second interface on the first fusion router.

```
interface vlan 3005
  vrf forwarding SHARED_SERVICES
  ip address 172.16.111.18 255.255.255.252
  no shutdown
```

```
interface vlan 3006
  vrf forwarding CAMPUS
  ip address 172.16.111.22 255.255.255.252
  no shutdown
```

```
interface vlan 3007
  vrf forwarding RESEARCH
  ip address 172.16.111.26 255.255.255.252
  no shutdown
```

```
interface vlan 3008
  vrf forwarding GUEST
  ip address 172.16.111.30 255.255.255.252
  no shutdown
```

```
interface TenGigabitEthernet1/1/2
  description Connected to Internal Border 6832-01 | TenGig 1/11
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 3005-3008
  no shutdown
```

5. Verify IP connectivity between the first fusion router and the border nodes using ping commands.

```
ping vrf SHARED_SERVICES 172.16.111.1
ping vrf CAMPUS 172.16.111.5
ping vrf RESEARCH 172.16.111.9
ping vrf GUEST 172.16.111.13
```

```
ping vrf SHARED_SERVICES 172.16.111.17
ping vrf CAMPUS 172.16.111.21
ping vrf RESEARCH 172.16.111.25
ping vrf GUEST 172.16.111.29
```

6. Repeat these steps on the second fusion router.

```
! First Interface and VLAN Set
```

```
vlan 3009-3016
  exit
```

```
interface vlan 3009
  vrf forwarding SHARED_SERVICES
  ip address 172.16.111.34 255.255.255.2
  no shutdown
```

```

interface vlan 3010
  vrf forwarding CAMPUS
  ip address 172.16.111.38 255.255.255.252
  no shutdown

interface vlan 3011
  vrf forwarding RESEARCH
  ip address 172.16.111.42 255.255.255.252
  no shutdown

interface vlan 3012
  vrf forwarding GUEST
  ip address 172.16.111.46 255.255.255.252
  no shutdown

interface TenGigabitEthernet1/1/2
  description Connected to Internal Border 6832-02 | TenGig 1/12
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 3009-3012
  no shutdown

! Second Interface and VLAN Set
interface vlan 3013
  vrf forwarding SHARED_SERVICES
  ip address 172.16.111.50 255.255.255.252
  no shutdown

interface vlan 3014
  vrf forwarding CAMPUS
  ip address 172.16.111.54 255.255.255.252
  no shutdown

interface vlan 3015
  vrf forwarding RESEARCH
  ip address 172.16.111.58 255.255.255.252
  no shutdown

interface vlan 3016
  vrf forwarding GUEST
  ip address 172.16.111.62 255.255.255.252
  no shutdown

interface TenGigabitEthernet1/1/2
  description Connected to Internal Border 6832-01 | TenGig 1/12
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 3013-3016
  no shutdown

```

7. Verify IP connectivity between the second fusion router and the border nodes using ping commands.

```

ping vrf SHARED_SERVICES 172.16.111.33
ping vrf CAMPUS 172.16.111.37
ping vrf RESEARCH 172.16.111.41
ping vrf GUEST 172.16.111.45

ping vrf SHARED_SERVICES 172.16.111.49
ping vrf CAMPUS 172.16.111.53
ping vrf RESEARCH 172.16.111.57
ping vrf GUEST 172.16.111.61

```

Procedure 4: Establish BGP Adjacencies Between Fusion Routers and Border Nodes – Fusion Routers Part 2

Now that IP connectivity has been established and verified, BGP peering can be created between the fusion routers and the border nodes.

On the border nodes, Cisco DNA Center has created SVIs forwarding for a VRF and a BGP configuration listening for peering adjacencies using those SVIs as the update source. A corresponding configuration will be applied on the fusion routers with one deviation. Because the shared services are in a VRF on the fusion routers, they will advertise the SHARED_SERVICES prefixes using *address-family ipv4 vrf* in BGP configuration and not the default *address-family ipv4* which is associated with the GRT.



The deviation above is due to the design choice of placing shared services in a dedicated VRF rather than in the GRT. These steps and the design modality will allow routes to be leaked between VRFs in later steps using route-target import and exports.

Please see [Appendix D](#) for further discussion on route leaking options.

1. Create the BGP process on the first fusion router.
2. Use the corresponding Autonomous-System defined in the IP-based transit from early steps. As a recommended practice, the Loopback 0 interface is used as the BGP router ID.

```
router bgp 65333
  bgp router-id interface Loopback0
```

3. Enter IPv4 Address-Family for vrf **SHARED_SERVICES** and complete the following steps:
 - a. Define the first border node neighbor and its corresponding AS Number. (Remember that this corresponds with the GRT on the border nodes).
 - b. Define the update source to use the applicable SVI.
 - c. Activate the exchange of NLRI with the first border node.

```
address-family ipv4 vrf SHARED_SERVICES
  neighbor 172.16.111.1 remote-as 65001
  neighbor 172.16.111.1 update-source Vlan3001
  neighbor 172.16.111.1 activate
```

4. Repeat these steps for the adjacency with the redundant border node.

```
address-family ipv4 vrf SHARED_SERVICES
  neighbor 172.16.111.17 remote-as 65001
  neighbor 172.16.111.17 update-source Vlan3005
  neighbor 172.16.111.17 activate
```

5. Configure the adjacencies for the remaining VRFs using the applicable SVIs as the update source for the peering.

```
address-family ipv4 vrf CAMPUS
  neighbor 172.16.111.5 remote-as 65001
  neighbor 172.16.111.5 update-source Vlan3002
  neighbor 172.16.111.5 activate

  neighbor 172.16.111.21 remote-as 65001
  neighbor 172.16.111.21 update-source Vlan3006
  neighbor 172.16.111.21 activate
```

```
address-family ipv4 vrf RESEARCH
  neighbor 172.16.111.9 remote-as 65001
  neighbor 172.16.111.9 update-source Vlan3003
  neighbor 172.16.111.9 activate
```

```
  neighbor 172.16.111.25 remote-as 65001
  neighbor 172.16.111.25 update-source Vlan3007
  neighbor 172.16.111.25 activate
```

```
address-family ipv4 vrf GUEST
  neighbor 172.16.111.13 remote-as 65001
  neighbor 172.16.111.13 update-source Vlan3004
  neighbor 172.16.111.13 activate
```

```
  neighbor 172.16.111.29 remote-as 65001
  neighbor 172.16.111.29 update-source Vlan3008
  neighbor 172.16.111.29 activate
```

6. Once this configuration is finished on the first fusion router, repeat on the second fusion router.

```
router bgp 65333
  bgp router-id interface Loopback0

address-family ipv4 vrf SHARED_SERVICES
  neighbor 172.16.111.33 remote-as 65001
  neighbor 172.16.111.33 update-source Vlan3009
  neighbor 172.16.111.33 activate

  neighbor 172.16.111.49 remote-as 65001
  neighbor 172.16.111.49 update-source Vlan3013
  neighbor 172.16.111.49 activate

address-family ipv4 vrf CAMPUS
  neighbor 172.16.111.37 remote-as 65001
  neighbor 172.16.111.37 update-source Vlan3010
  neighbor 172.16.111.37 activate

  neighbor 172.16.111.53 remote-as 65001
  neighbor 172.16.111.53 update-source Vlan3014
  neighbor 172.16.111.53 activate

address-family ipv4 vrf RESEARCH
  neighbor 172.16.111.42 remote-as 65001
  neighbor 172.16.111.42 update-source Vlan3011
  neighbor 172.16.111.42 activate

  neighbor 172.16.111.57 remote-as 65001
  neighbor 172.16.111.57 update-source Vlan3015
  neighbor 172.16.111.57 activate

address-family ipv4 vrf GUEST
  neighbor 172.16.111.45 remote-as 65001
  neighbor 172.16.111.45 update-source Vlan3012
  neighbor 172.16.111.45 activate

  neighbor 172.16.111.61 remote-as 65001
  neighbor 172.16.111.61 update-source Vlan3016
  neighbor 172.16.111.61 activate
```

Procedure 5: Use Route-Targets to Leak Shared Services Routes – Fusion Routers Part 3

With the shared services routes in a VRF on the fusion routers and BGP adjacencies formed, route-targets can be used to leak routes between the VRFs. This leaking will occur internally in the routing tables of the fusion routers. Once leaked, these routes will be advertised to the fabric border nodes and be accessible to endpoints in the fabric needing these services.



Route-targets (RT) are a BGP Extended Community Attribute. It is an optional, transitive attribute that is preserved as the information crosses autonomous system boundaries. The route-target identifies the VPN membership (which VRF) and is used to transfer routes between VRFs. Additional information can be found in [RFC 4360](#). It is similar in structure to the route-distinguisher (RD), although these two elements should not be confused. A route-distinguisher is appended to a route (when advertised through MP-BGP) to make it unique.

The VRF configuration has already been copied from the border nodes and configured on the fusion routers. Each VRF is already importing and exporting its own route-targets. This simply means the VRF is exporting its own routes using the defined route-targets and importing routes from any other VRF with the defined route-targets.

This procedure will complete the following steps and achieve these results:

- Each VRF on the fusion routers will import the route-target associated with the SHARED_SERVICES VRF (1:4097).
- By importing this route-target (1:4097), each VRF will import the shared services subnets into the routing table.
- The SHARED_SERVICES VRF will import the route-targets associated with the other VRFs (1:4099, 1:4100, 1:4101).
- These actions fuse the routes between the various VRF tables.

The resulting configuration will create the following outcome:

Figure 19 Route-Target and Prefix Import and Export



1. On both fusion routers, import the SHARED_SERVICES VRF route-target into the CAMPUS VRF.

```
vrf definition CAMPUS
  address-family ipv4
    route-target import 1:4097
    route-target import 1:4099
  exit-address-family
```

2. Repeat this SHARED_SERVICES VRF route-target import step for the remaining VRFs.

```
vrf definition GUEST
  address-family ipv4
```

```
route-target import 1:4097
route-target import 1:4100
exit-address-family
```

```
vrf definition RESEARCH
address-family ipv4
route-target import 1:4097
route-target import 1:4100
exit-address-family
```

3. Under the SHARED_SERVICES VRF, import the route-targets for the fabric VRFs.

```
vrf definition SHARED_SERVICES
address-family ipv4
route-target import 1:4099
route-target import 1:4100
route-target import 1:4101
exit-address-family
```

Process 2: Provisioning and Verifying Access Points

Now that the edge nodes ports connecting to Access Points have been provisioned, the internal border nodes and fusion routers have been provisioned and configured, and route leaking of the shared services prefixes is complete, Access Points are able to access both the WLC and the DHCP server in order to receive an IP address, associate with the correct WLC through [DHCP Option 43](#), and be registered to the Cisco DNA Center inventory by the WLC.

Once in the inventory, APs must be provisioned in order to receive the correct RF profile configuration and to join the overlay network in the SD-Access Wireless Access Point role. The following steps will provision the Access Points to a floor in a building (site) and provision them with an RF profile, allowing them to operate in the fabric role.

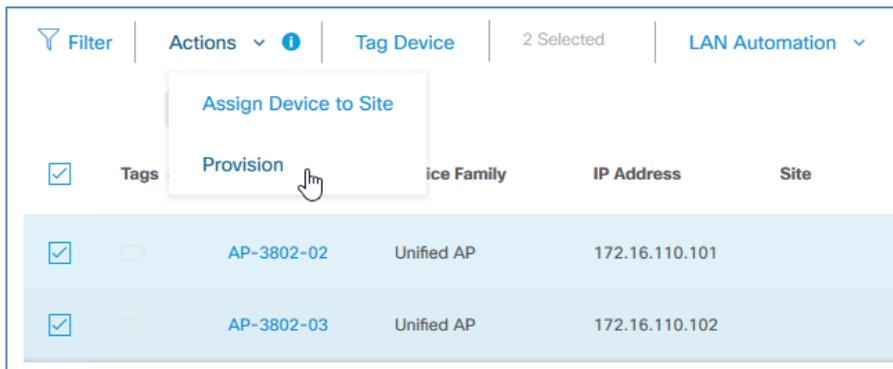


Tech tip

It may take a few minutes for the WLC to register the Access Points to the inventory of Cisco DNA Center. To manually speed up this process, select the WLC in **Inventory**, click **Actions**, and select **Resync**. If the Access Points have been associated with the WLC, this process will manually trigger the WLC to register them with Cisco DNA Center.

1. In the Cisco DNA Center dashboard, navigate to **PROVISION > Devices > Inventory**.
2. Select the Access Points.
3. Click **Actions > Provision**.

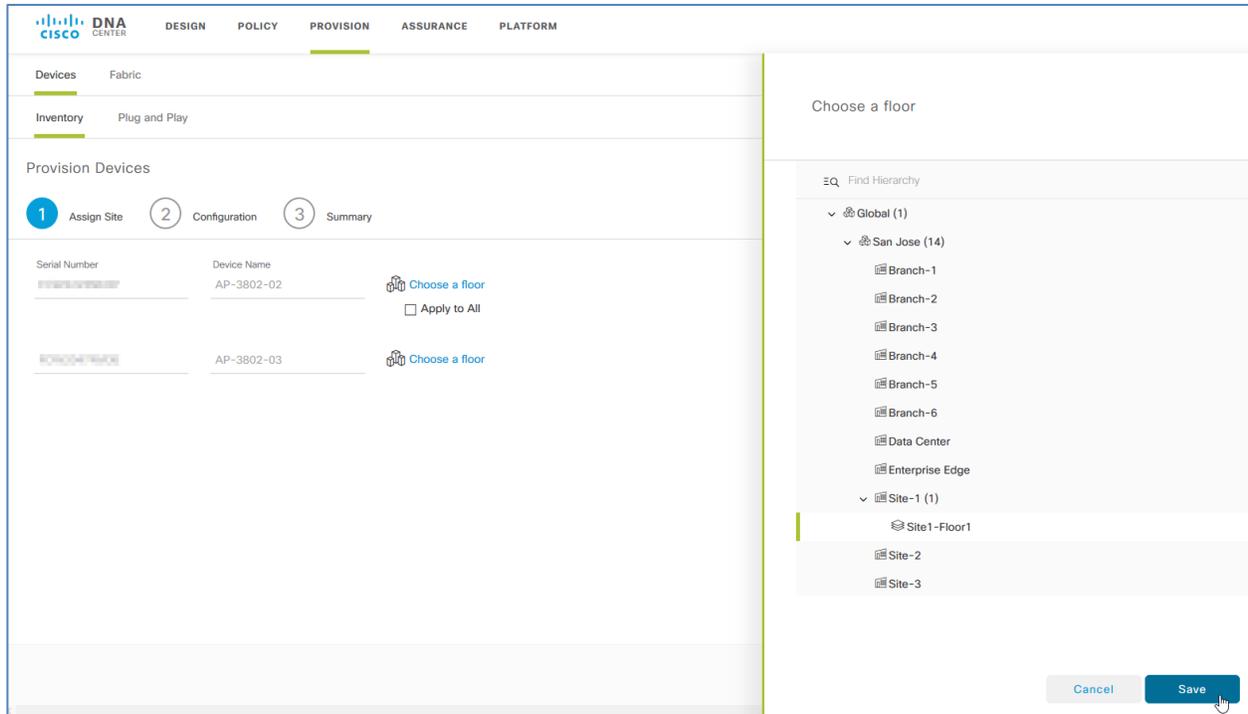
A Provision Devices wizard screen appears.



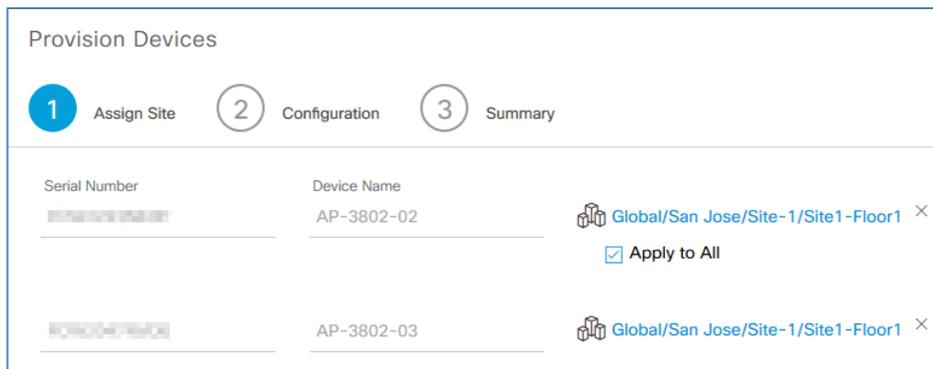


Access points must be assigned to a floor. If floors have not been added to a building, please add them using the steps described in the [Design application](#) portion of the workflow.

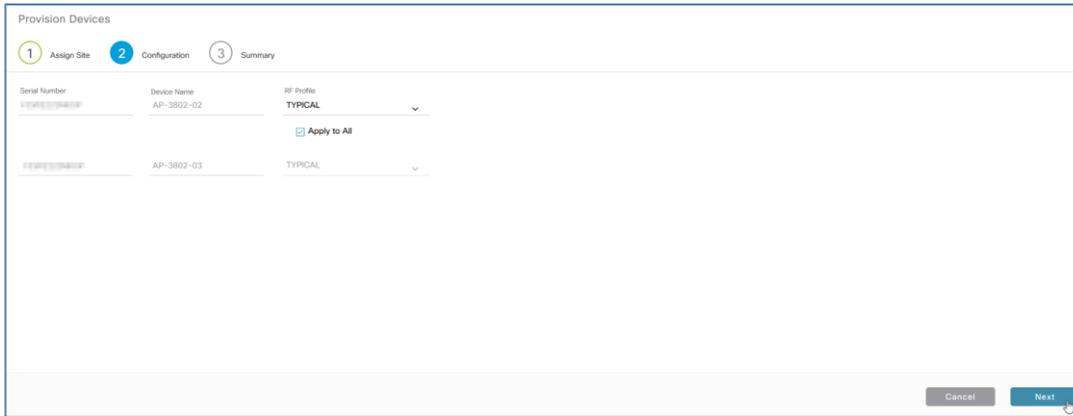
4. Within the first wizard screen, **Assign Site**, click **Choose a floor**.
5. From the slide-in pane on the right, select the **Floor Assignment (Choose a floor)** for the devices.
6. Click **Save**.



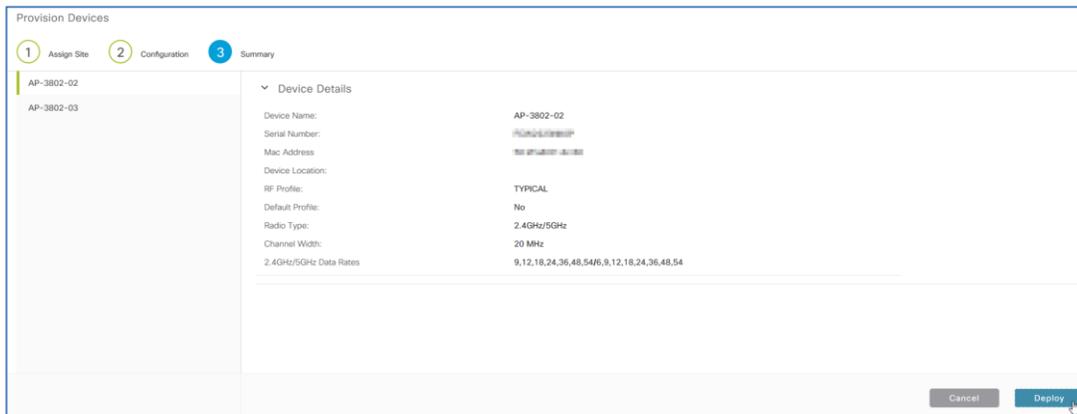
7. If the APs are all in the same site, use the **Apply to All** check box. Different site locations can be selected for each device where applicable.



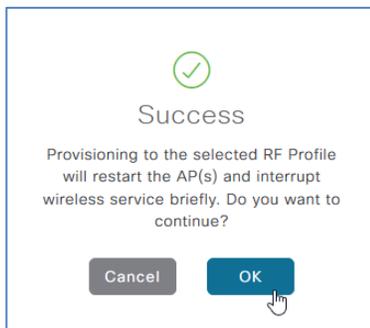
8. On the **Configuration** page of the Wizard, select the desired **RF Profile** (example: Typical).
- Apply to All** can be used to apply the same RF profile to multiple Access Points.
9. Click **Next**.



10. Verify that the **Summary** screen displays the RF profile information which includes the **Radio Type, Channel Width, and Data Rates**.
11. Click **Deploy**.



12. In the slide-in pane for **Provision Device**, leave the default selection of **Now**.
13. Click **Apply**.
14. A **Warning** message appears indicating that Access Points will be rebooted. Click **OK**.



15. Configuration of the Access Points begins, and status messages appear as each device is provisioned successfully. The Device Inventory screen updates with **Last Provisioned Time** and **Provision Status**.

Use the **Refresh** button to see the final status of device provisioning.

[Wireless Radio Frequency Profiles \(RF Profiles\)](#) define the Dynamic Bandwidth Selection (Channel Width), Dynamic Channel Assignment, Supported Data Rates, Transmit Power Configuration and Thresholds, and Receive Start of Packet Detection Threshold. **Low**, **Medium (Typical)**, and **High** are the pre-canned RF profiles in Cisco DNA Center. The Design application also supports the creation of custom RF profiles.

Procedure 1: Access Point Verification

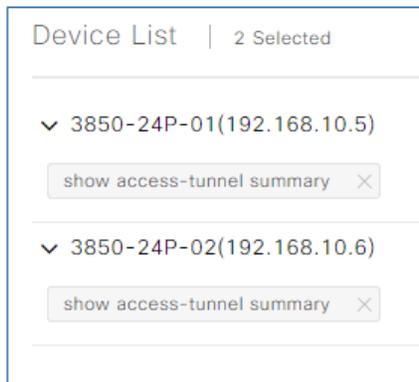
Traditionally, Access Points take the 802.11 traffic from the endpoint, convert it to 802.3, and then either put in on the local wired network or use CAPWAP to tunnel the traffic to the WLC which then places the traffic on the wired network. When an SSID is enabled for the fabric, it tells the Access Points to treat the packets differently. In the fabric, an AP converts 802.11 traffic to 802.3, but also encapsulates it into VXLAN to encode the VN and SGT information of the client. This VXLAN tunnel is terminated at the first-hop switch (the edge node).

The Command Runner tool can be used to validate that the Access Points have joined the fabric overlay and are communicating with their connected edge node by verifying this VXLAN Tunnel.

1. At the top-right of the Cisco DNA Center dashboard, navigate to **TOOLS > Command Runner**.
2. Select an edge node with attached Access Points (example: 3850-24p-01 (192.168.10.5)).
3. Verify that the device is added to the device list in the center of the screen.
4. In the command bar, type the command `show access-tunnel summary`.
5. Click the **Add** button to add the command.



6. Verify that the command appears below the device in the device list.



7. Click the **Run Command(s)** button at the bottom of the screen.



8. After a few moments, the command is executed, and the results are returned to Cisco DNA Center. Click the command (shown with a green text to indicate that it was executed successfully) to view to result.

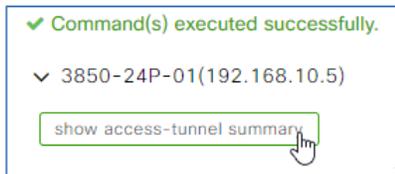
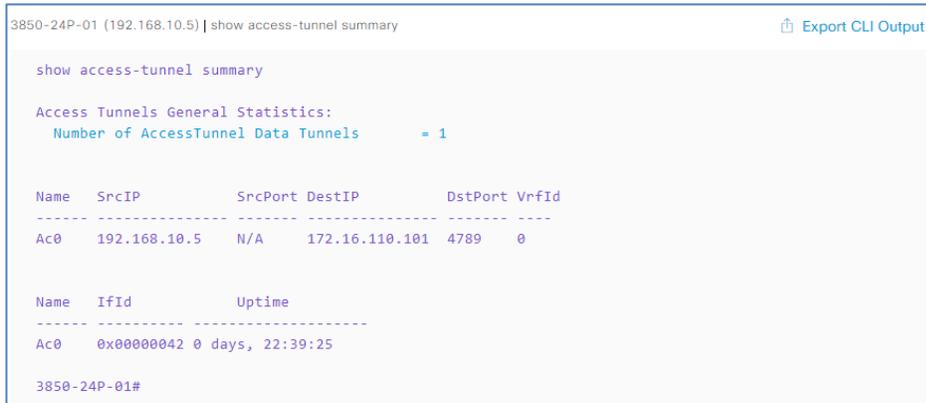


Figure 20 Successful VXLAN Tunnel Between Access Points and Edge Node



Process 3: Creating IBGP Adjacencies Between Redundant External Border Nodes

The fabric overlay works by tunneling traffic between routing locators (RLOC). The RLOCs in SD-Access are always the Loopback 0 interface of the device operating in the fabric role.

When border nodes are configured with the *Outside World (External)* option, there is no redistribution of external network into SD-Access fabric – this is a function performed by internal border nodes.

Any traffic with a destination outside of the fabric site must traverse the border nodes. If an uplink interface or upstream device fails for a border node, the Loopback 0 interface (the RLOC) of that border node is still reachable by other devices in the fabric. This leads to a potential black-holing of packets.

There is no built-in method in the Cisco DNA Center provisioning or in LISP to mitigate this issue. The preferred method to address this potential problem is to form IBGP adjacencies between redundant border nodes. This is accomplished using a cross-link between redundant devices. From a design and recommended practice perspective, and to allow uninterrupted traffic, IBGP adjacencies should be formed between all redundant border types and redundant fusion routers. These adjacencies need to be formed for the global routing table and all VRFs.

Once the IBGP peering sessions are established, a redirection path is provided in the event of a failed interface or upstream device. It is an additional network traffic path that helps avoid traffic black-holing in basic failure situations.

Border nodes and fusion routers may be either Layer-3 switches or a true routing platform. Among Layer-3 switches, the device may support subinterfaces – such as with the Catalyst 6800 Series – or may only support Switched Virtual Interfaces (SVIs). This results in many deployment-specific variations for the IBGP configuration. This prescriptive deployment guide uses routing platforms as the fabric border nodes in most of the fabric-enabled sites. Therefore, the IBGP configuration will be shown using routers.

Layer-3 switches are also used for both the internal border nodes and the fusion routers. For an example IBGP configuration using a switching platform as a border node or fusion router in a redundant pair, please see [Appendix E](#).



Crosslinks between redundant devices are a best practice in network design. In the topology, crosslinks exist between all redundant borders – both internal and external – redundant Fusion devices, and redundant Internet edge routers. IBGP is configured between each of these pairs to provide a redirection path in the event of failed interface or upstream device failure.

Procedure 1: Configure IBGP – Routing Platforms

Routing platforms such as the Integrated Services Routers (ISR) and Aggregation Services Routers (ASR) support subinterfaces (dot1q VLAN subinterface). Subinterfaces on routers should be used to create the IBGP adjacency between redundant devices.

This procedure will first create the subinterfaces, assign them to an existing VRF and give them an IP address. Once IP connectivity is established and verified, the BGP adjacencies will be created using the subinterfaces as the update source.



There are multiple reserved VLANs in Cisco DNA Center. VLANs below 1000 are *generally* safe to use for creating the IBGP adjacencies. Because the VLANs are locally significant between the redundant devices, the same VLANs could be used for multiple redundant pairs.

1. In the CLI of the first external border node router, bring up the physical interface.

```
interface GigabitEthernet0/0/0
no shutdown
```

2. Define a subinterface that will forward for the Global Routing Table.
3. Define the VLAN number and the IP address.

```
interface GigabitEthernet0/0/0.100
description Connected to ISR-4451-08 | GRT
encapsulation dot1q 100
ip address 10.10.10.1 255.255.255.252
```

4. Define subinterfaces, VLANs, and IP addresses for the remaining VRFs.

```
interface GigabitEthernet0/0/0.101
description Connected to ISR-4451-08 | VRF CAMPUS
vrf forwarding CAMPUS
encapsulation dot1q 101
ip address 10.10.10.5 255.255.255.252
```

```
interface GigabitEthernet0/0/0.102
description Connected to ISR-4451-08 | VRF GUEST
encapsulation dot1q 102
vrf forwarding GUEST
ip address 10.10.10.9 255.255.255.252
```

```
interface GigabitEthernet0/0/0.103
description Connected to ISR-4451-08 | VRF RESEARCH
encapsulation dot1q 103
vrf forwarding RESEARCH
ip address 10.10.10.13 255.255.255.252
```

5. In the CLI of the second external border node router, bring up the physical interface.

```
interface GigabitEthernet0/0/0
no shutdown
```

6. Define a subinterface that will forward for the Global Routing Table.
7. Define the corresponding VLAN number and IP address.

```
interface GigabitEthernet0/0/0.100
description Connected to ISR-4451-07 | GRT
encapsulation dot1q 100
ip address 10.10.10.2 255.255.255.252
```

8. Define the corresponding subinterfaces, VLANs, and IP addresses for the remaining VRFs.

```
interface GigabitEthernet0/0/0.101
```

```

description Connected to ISR-4451-07 | VRF CAMPUS
vrf forwarding CAMPUS
encapsulation dot1Q 101
ip address 10.10.10.6 255.255.255.252

interface GigabitEthernet0/0/0.102
description Connected to ISR-4451-07 | VRF GUEST
encapsulation dot1Q 102
vrf forwarding GUEST
ip address 10.10.10.10 255.255.255.252

interface GigabitEthernet0/0/0.103
description Connected to ISR-4451-07 | VRF RESEARCH
encapsulation dot1Q 103
vrf forwarding RESEARCH
ip address 10.10.10.14 255.255.255.252

```

- Verify connectivity between the subinterfaces on the border nodes.

```

ping 10.10.10.2
ping vrf CAMPUS 10.10.10.6
ping vrf RESEARCH 10.10.10.10
ping vrf GUEST 10.10.10.14

```

- On the first border node, enable IBGP the redundant peer for the Global Routing Table. Use the subinterfaces as the update source.

```

router bgp 65001
  neighbor 10.10.10.2 remote-as 65001
  neighbor 10.10.10.2 update-source GigabitEthernet0/0/0.100

address-family ipv4
  neighbor 10.10.10.2 activate

```

- Enable IBGP between the redundant peer for the remainder of the VRFs. Use the subinterfaces as the update source.

```

address-family ipv4 vrf CAMPUS
  neighbor 10.10.10.6 remote-as 65001
  neighbor 10.10.10.6 update-source GigabitEthernet0/0/0.101
  neighbor 10.10.10.6 activate

address-family ipv4 vrf GUEST
  neighbor 10.10.10.10 remote-as 65001
  neighbor 10.10.10.10 update-source GigabitEthernet0/0/0.102
  neighbor 10.10.10.10 activate

address-family ipv4 vrf RESEARCH
  neighbor 10.10.10.14 remote-as 65001
  neighbor 10.10.10.14 update-source GigabitEthernet0/0/0.103
  neighbor 10.10.10.14 activate

```

- Complete the corresponding configuration on the redundant border node.

```

router bgp 65001
  neighbor 10.10.10.1 remote-as 65001
  neighbor 10.10.10.1 update-source GigabitEthernet0/0/0.100

address-family ipv4
  neighbor 10.10.10.1 activate

address-family ipv4 vrf CAMPUS
  neighbor 10.10.10.5 remote-as 65001
  neighbor 10.10.10.5 update-source GigabitEthernet0/0/0.101
  neighbor 10.10.10.5 activate

address-family ipv4 vrf GUEST
  neighbor 10.10.10.9 remote-as 65001
  neighbor 10.10.10.9 update-source GigabitEthernet0/0/0.102
  neighbor 10.10.10.9 activate

```

```

address-family ipv4 vrf RESEARCH
neighbor 10.10.10.13 remote-as 65001
neighbor 10.10.10.13 update-source GigabitEthernet0/0/0.103
neighbor 10.10.10.13 activate

```

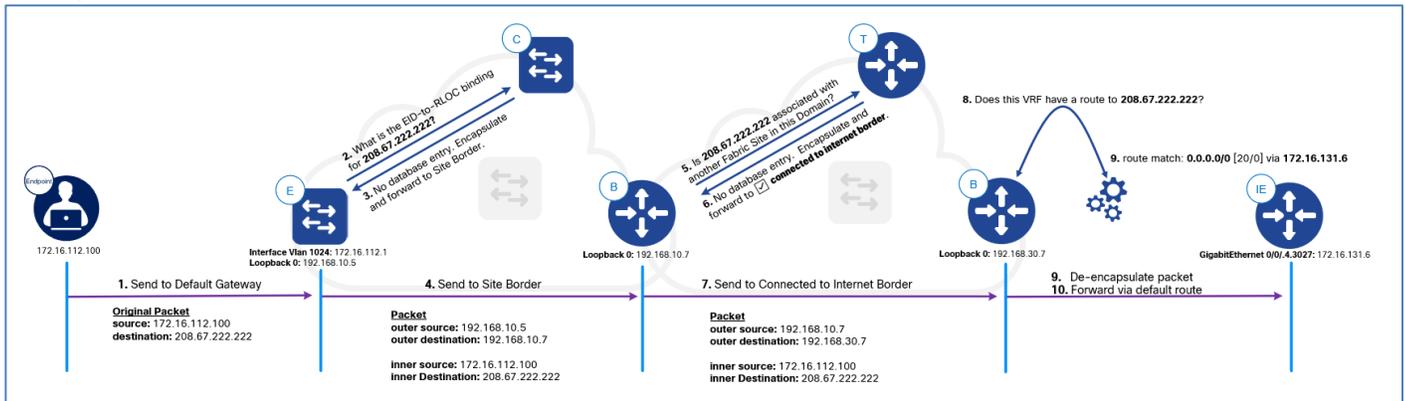
Process 4: Providing Internet Access for SD-Access for Distributed Campus

Internet access for SD-Access for Distributed Campus operates differently from Internet access in a single-site SD-Access deployment, particularly with border node provisioning. In a single site, the egress point for the network to the Internet is the external border nodes defined in network provisioning. In an SD-Access for Distributed Campus deployment, which by its definition has multiple fabric sites, many sites in the deployment may have direct Internet access (DIA) or all Internet traffic may traverse the SD-Access transit to a single-site location in order to egress the fabric domain. Both designs are valid for a deployment, although the latter is the most common in a metro-area with traffic tunneled across the MAN to an HQ location.

Within an individual site or a single-site deployment, the External or Anywhere border node is considered the “fabric site gateway of last resort.” From a traffic flow perspective, the site-local control plane node is queried by the fabric infrastructure for where to send packets. If the site-local control plane node does not have an entry in its database for the destination, it instructs the requesting device to send the traffic to the external (or anywhere) border node.

In SD-Access for Distributed Campus, the site border nodes are the fabric site gateway of last resort, although the control plane traffic differs from a single-site deployment. Within the site, the same control plane requests and replies occur as described above. However, the site border nodes may receive packets for which they have no forwarding path. In this case, they query the transit control plane nodes to determine if the traffic should be sent to another fabric site or if the traffic should be sent towards the Internet. If traffic is bound for the Internet – if there is no entry in the database – the transit control plane node instructs the site borders to send traffic to specific devices in the deployment.

Figure 21 Unknown or Internet Destination Packet Flow



In the Cisco DNA Center GUI, these specific devices are border nodes flagged with the Is this site connected to Internet? option when provisioning. Selecting this option enables this border node as a “fabric domain gateway of last resort.”

When a site border queries the transit control plane node to reach an unknown destination and there is no entry in the control plane database, the site border is instructed to send the traffic to the border configured with the **connected-to-Internet** option. Traffic is sent to each border node with this flag in a round-robin fashion. This is an important consideration if more than one site in the deployment has Internet access.



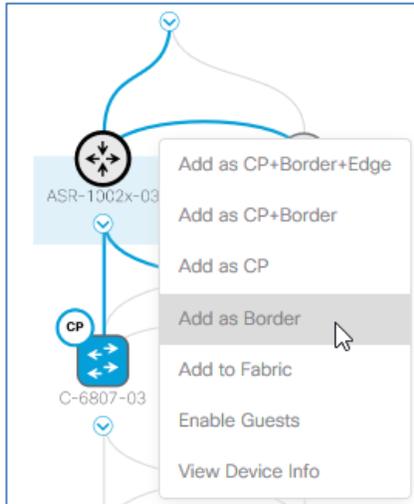
The **Is this site connected to Internet?** option should only be used if this border node has an upstream hop such as the enterprise edge or service provider edge router. Generally, the **connected-to-Internet** device should be the device in the deployment with a default route (0.0.0.0/o) to a next hop that ultimately provides access to the Internet. Border nodes connected **only** to the SDA transit should not use this option.

Within the device itself, this flag tells the border nodes to use the RIB and not LISP for the default route.

Procedure 1: Provision the **Connected-to-Internet** Border Node

In the deployment topology, the border node for Site-3 is connected to both the SD-Access transit and to the upstream Internet edge router. The workflow will demonstrate provisioning both the SD-Access transit and IP-based transit with a Layer-3 handoff using VRF-lite on a **connected-to-Internet** border node.

1. In the Cisco DNA Center dashboard, navigate to **PROVISION > Fabric**.
2. Select the **Fabric Domain (SJC)**.
3. Select the **Fabric Site (Site-3)** from the **Fabric-Enabled Sites** on the right. The topology map loads showing the devices provisioned to that site.
4. Click a device to perform the fabric border node role.
5. For provisioning the IP-based transit, this device should also be connected to the upstream Internet edge router.
6. Select **Add as Border** (or **Add as CP+Border** where applicable) from the popup menu. A slide-in pane appears.



7. As the border is connected to an upstream provider and the metro-E transit, select **Outside World (External)** in the slide-in pane.
8. Supply the **BGP Local AS Number** (example: 65003).
9. Under **Select IP Address Pool**, chose the IP address pool to be used for the IP-based handoff (example: BGP-SITE-03 (172.16.131.0/24)).

ASR-1002x-03

Border to

Rest of Company (Internal) ⓘ

Outside World (External) ⓘ

Anywhere (Internal & External) ⓘ

Local Autonomous Number

65003

ⓘ

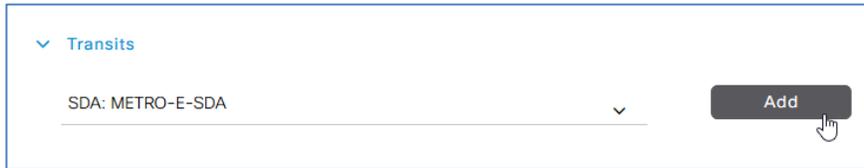
Select IP Address Pool

✖ BGP-SITE-03 (172.16.131.0/24) ▾

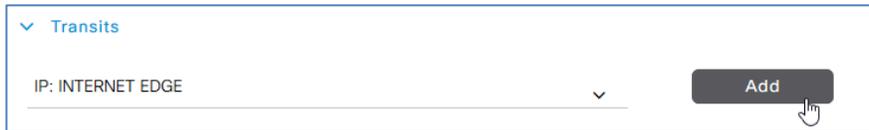
10. Because this device is the egress point to the Internet and has a default route pointing to an upstream device, select the option **Is this site connected to Internet?**

Is this site connected to Internet?

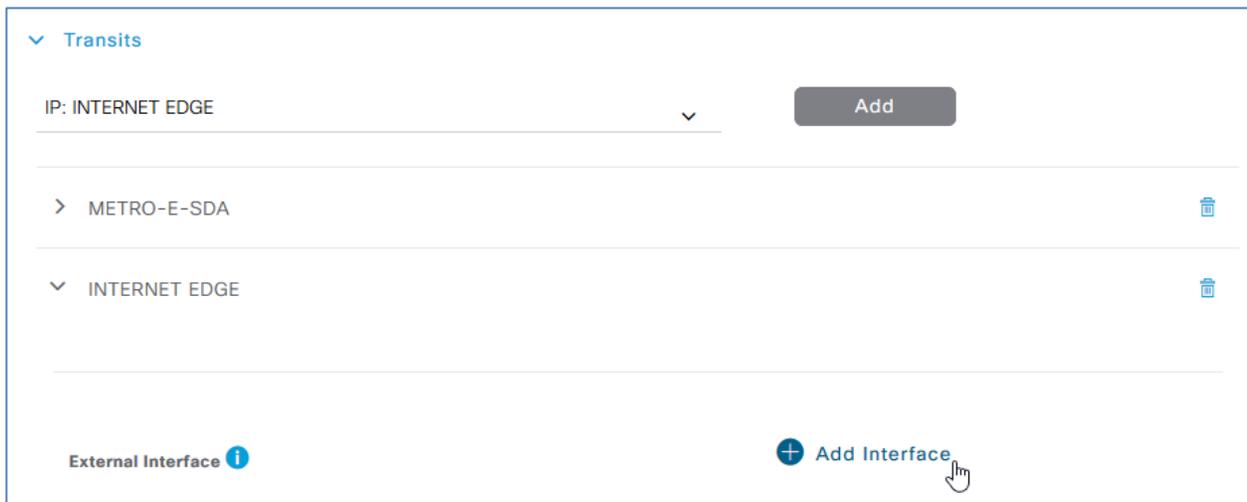
11. Under **Transits**, select the previously-created SDA transit (example: SDA: METRO-E-SDA).
12. Click the **gray Add Button**.
The SDA transit is added.



13. Under **Transits** again, select the previously-created IP-based transit (example: IP: INTERNET EDGE).
14. Click the **gray Add Button**.
The IP-based transit is added.



15. Click the IP transit name to expand it.
16. Click **+ Add Interface**.



17. In the new slide-in pane, select the **External Interface** connected to the first Internet edge router (example: GigabitEthernet0/1/2).
18. The BGP Remote AS Number is automatically filled in due to previous steps in the workflow. Click **^** to expand the **Virtual Network** selection.
19. Select all VNs that will require access to the Internet.
Click **Save**.

ASR-1002x-03

< Back

External Interface
* GigabitEthernet0/1/2

Remote AS Number
65535

Virtual Network ⓘ

- GUEST
- CAMPUS
- RESEARCH
- DEFAULT_VN
- INFRA_VN

Cancel Save

20. Verify that the defined **External Interface** and **Number of VNs** are displayed.

21. Click **+ Add Interface**.

ASR-1002x-03.dna.local

Local Autonomous Number
65003

Select IP Address Pool
* BGP-SITE-03 (172.16.131.0/24)

Is this site connected to Internet?

Transits

Select Transit Add

METRO-E-SDA

INTERNET EDGE

External Interface ⓘ

Interface	Number of VN	
GigabitEthernet0/1/3	4	Remove

+ Add Interface

22. In the new slide-in pane, select the **External Interface** connected to the second Internet edge router (example: GigabitEthernet0/1/3).

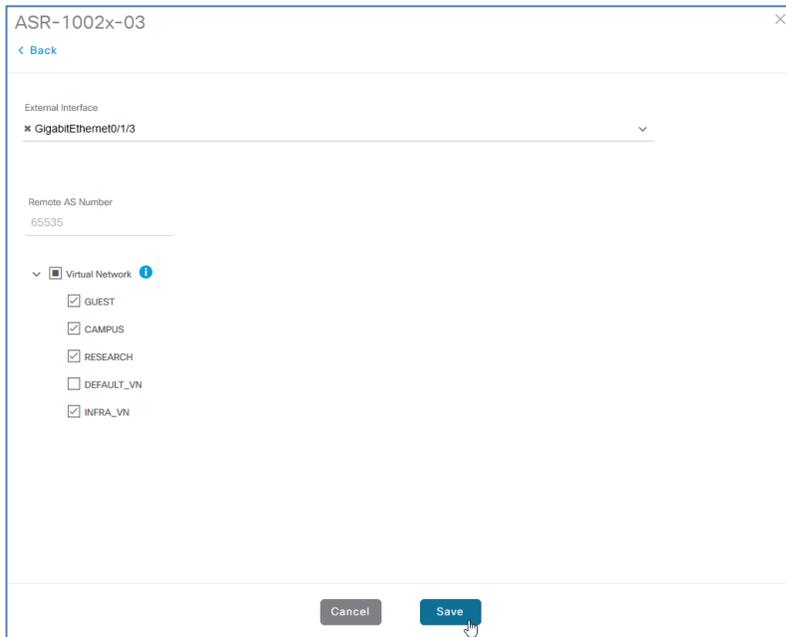
This should be a different interface than the interface used in the previous steps.

ASR-1002x-03.dna.local

< Back

External Interface
* GigabitEthernet0/1/3

23. The BGP Remote AS Number is automatically filled in due to previous steps in the workflow. Click **^** to expand the **Virtual Network** selection.
24. Select all VNs that will require access to the Internet. These should be the same VNs as selected for the other interface.
25. Click **Save**.



The screenshot shows a configuration window titled "ASR-1002x-03" with a close button in the top right corner. Below the title is a "< Back" link. The main content area is divided into sections: "External Interface" with a dropdown menu showing "GigabitEthernet0/1/3"; "Remote AS Number" with the value "65535"; and "Virtual Network" which is expanded to show a list of virtual networks with checkboxes: "GUEST" (checked), "CAMPUS" (checked), "RESEARCH" (checked), "DEFAULT_VN" (unchecked), and "INFRA_VN" (checked). At the bottom of the window are "Cancel" and "Save" buttons, with a mouse cursor hovering over the "Save" button.

26. Click the blue **Add** button at the bottom of the border node slide-in pane.
27. A **Warning** pop-up provides additional information regarding fabric borders. Click **OK**.

ASR-1002x-03

Border to

Rest of Company (Internal) ⓘ
 Outside World (External) ⓘ
 Anywhere (Internal & External) ⓘ

Local Autonomous Number
65003

Select IP Address Pool
* BGP-SITE-03 (172.16.131.0/24) ▼

Is this site connected to Internet?

▼ Transits

IP: INTERNET EDGE ▼ Add

> METRO-E-SDA 🗑️

▼ INTERNET EDGE 🗑️

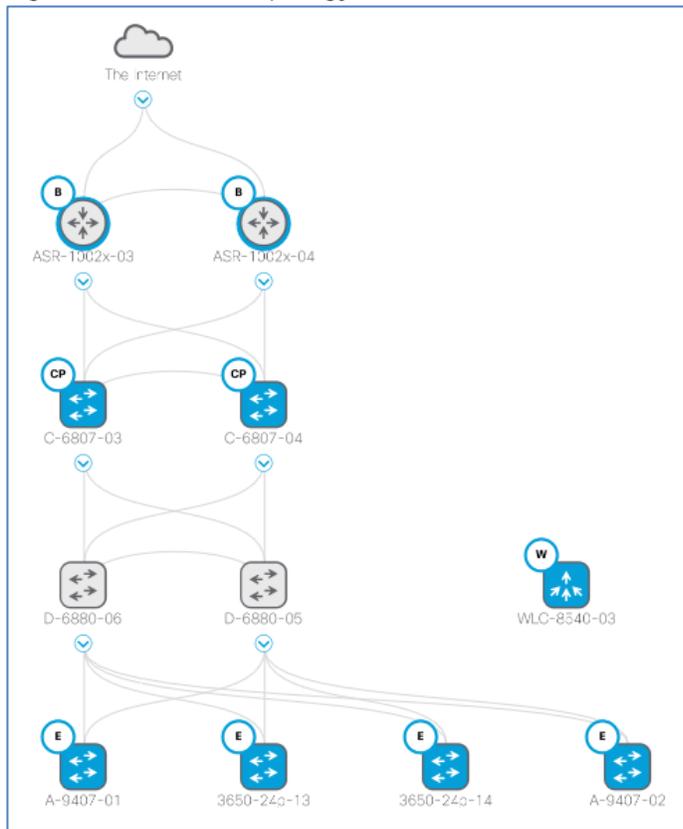
External Interface ⓘ + Add Interface

Interface	Number of VN	
GigabitEthernet0/1/2	4	Remove
GigabitEthernet0/1/3	4	Remove

Cancel Add

28. Repeat these steps for the second External-Only border node connected to the Internet edge routers.
29. Once complete, click **Save** in the bottom right of the fabric topology screen.
30. In the **Provision Device** slide-in pane, leave the default selection of **Now**.
31. Click **Apply**.

Figure 22 Site-3 Topology – Fabric Role Intent

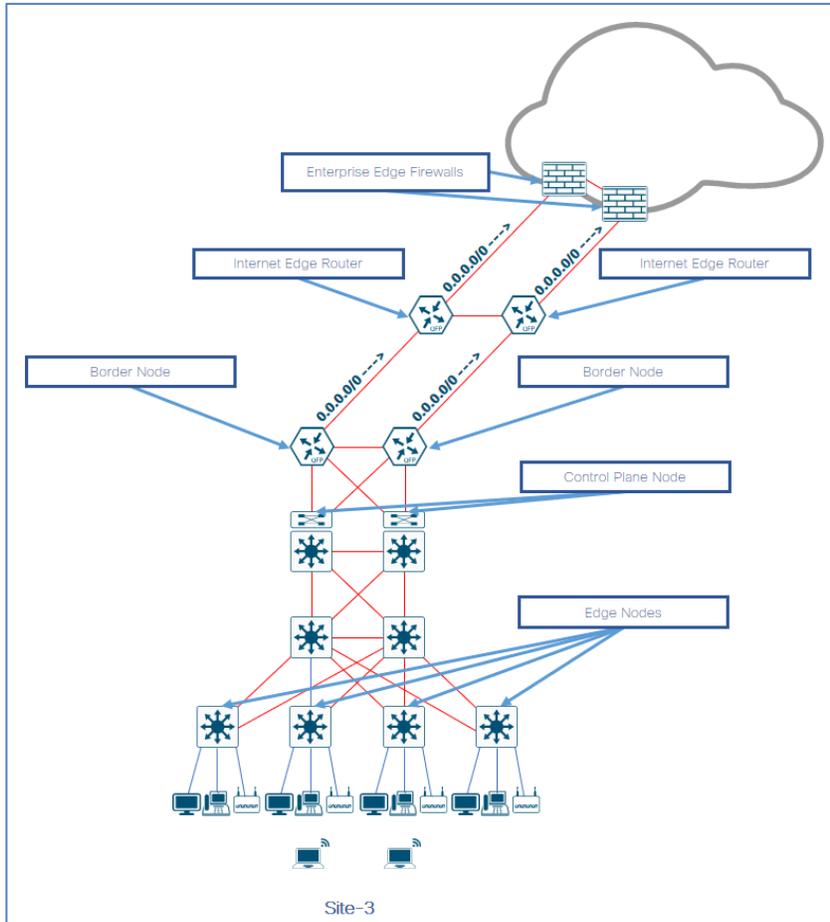


Internet Edge Topology Overview

External (Internet) connectivity for the fabric domain has a significant number of possible variations, and these variations are based on underlying network design. The common similarity among these variations is that the **connected-to-Internet** border node will be connected to a next-hop device that will be traversed to access the Internet due to its routing information. This device could be an actual Internet edge router, a fusion router, the ISP router, or some other next-hop device.

In this prescriptive deployment, the Site-3 border nodes are connected to a pair of Internet edge routers as their external next-hop. These Internet edge routers are connected to upstream Enterprise edge firewalls that are configured to use NAT and are ultimately connected to the ISP. The Site-3 border nodes have a static default route to Internet edge routers via the Global Routing Table. The Internet edge routers have a static default route to the Enterprise edge firewalls via their Global Routing Table.

Figure 23 Default Route Overview – Site-3



Tech Tip

For topology diagram simplicity, redundant links from devices to the Internet edge routers and from these routers to the Enterprise edge firewalls are not shown although they are present in the deployment.

Configuring Internet Connectivity

Regardless of how the rest of the network itself is designed or deployed outside of the fabric, a few things are going to be in common in deployments due to the configuration provisioned by Cisco DNA Center. A *connected-to-Internet* border node will have the SD-Access prefixes in its VRF routing tables. As a prerequisite for being *connected-to-Internet*, it will also have a default route to its next hop in its Global Routing Table. This default route must somehow be advertised from the GRT to the VRFs. This allows packets to egress the fabric domain towards the Internet. In addition, the SD-Access prefixes in the VRF tables on the border nodes must be advertised to the external domain – outside of the fabric domain – to draw (attract) packets back in.

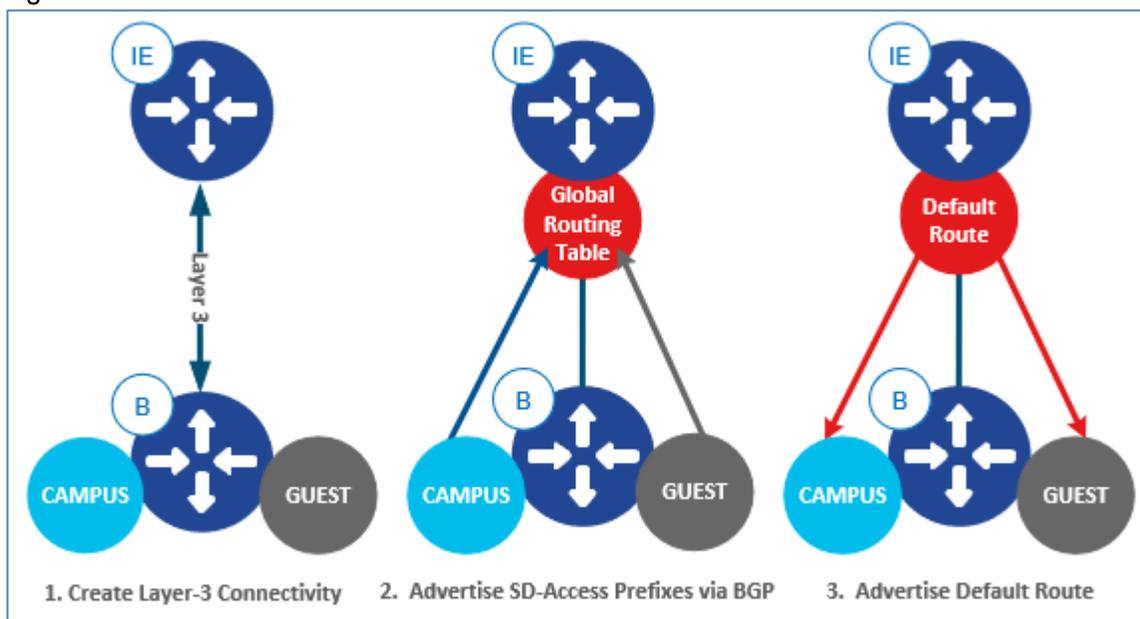


The solution used in the validation is one of many possible solutions to achieve the goal of providing Internet access to the endpoints in the fabric domain. The solution provided represents the recommended practices based on the fabric domain topology, the upstream connectivity, and the specific needs of the deployment.

Providing Internet access to the SDA fabric (leaking between the global and VRF routing tables) is a multi-step process performed in the CLI of the Internet edge routers.

1. Create Layer-3 connectivity between the border nodes and Internet edge routers.
2. Use BGP to create adjacencies and receive VRF routes.
3. Advertise the default route back to the VRF tables of the Site-3 border.

Figure 24 Default Route Advertisement Workflow



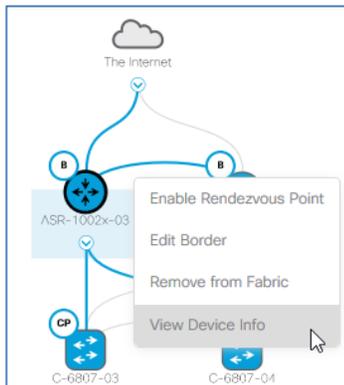
On the Site-3 border nodes, the VRF and BGP configurations have been provisioned by Cisco DNA Center along with the Layer-3 handoff. For Internet access in this prescriptive guide, their upstream peers will not become VRF-aware. All fabric domain prefixes will be learned in the Global Routing Table of the Internet edge routers, and the default route from these devices will be advertised back to Site-3 border nodes via BGP.

Procedure 2: Use the Provision Application for BGP and VRF-Lite information – Part 2

Cisco DNA Center is used to determine the Layer-3 handoff configuration that was provisioned in previous steps. This is needed for the corresponding configuration on the Internet edge routers.

1. In the Cisco DNA Center dashboard, navigate to **PROVISION > Fabric**.
2. Select the **Fabric Domain (SJC)**.

3. Select the **Fabric Site (Site-3)** from the **Fabric-Enabled Sites** on the right. The topology map loads showing the devices provisioned to that site.
4. Click the **Connected-to-Internet Border** provisioned in the previous steps.
5. In the pop-up menu, select **View Device Info**. A new slide-in pane appears.



6. Scroll to the interfaces listed in blue, and click > to expand the section. The Layer-3 handoff provisioning information is displayed. Note the information as it will be used in the next steps.

ASR-1002x-03			
Border Type	EXTERNAL		
Border Handoff			
Internal Domain Protocol Number	65003		
External Connectivity IP Pool	BGP-SITE-03		
GigabitEthernet0/1/3			
Layer3			
External Domain Protocol	65535		
Virtual Network	Vlan	Local IP	Remote IP
GUEST-Global/San_Jose /Site-3	3032	172.16.131.25/30	172.16.131.26/30
RESEARCH-Global/San_Jose /Site-3	3033	172.16.131.29/30	172.16.131.30/30
CAMPUS-Global/San_Jose /Site-3	3031	172.16.131.21/30	172.16.131.22/30
INFRA_VN-Global/San_Jose /Site-3	3030	172.16.131.17/30	172.16.131.18/30
GigabitEthernet0/1/2			
Layer3			
External Domain Protocol	65535		
Virtual Network	Vlan	Local IP	Remote IP
INFRA_VN-Global/San_Jose /Site-3	3026	172.16.131.1/30	172.16.131.2/30
GUEST-Global/San_Jose /Site-3	3028	172.16.131.9/30	172.16.131.10/30
RESEARCH-Global/San_Jose /Site-3	3029	172.16.131.13/30	172.16.131.14/30
CAMPUS-Global/San_Jose /Site-3	3027	172.16.131.5/30	172.16.131.6/30

ASR-1002x-04			
Border Type	EXTERNAL		
Border Handoff			
Internal Domain Protocol Number	65003		
External Connectivity IP Pool	BGP-SITE-03		
GigabitEthernet0/1/3			
Layer3			
External Domain Protocol	65535		
Virtual Network	Vlan	Local IP	Remote IP
CAMPUS-Global/San_Jose /Site-3	3035	172.16.131.37/30	172.16.131.38/30
RESEARCH-Global/San_Jose /Site-3	3037	172.16.131.45/30	172.16.131.46/30
GUEST-Global/San_Jose /Site-3	3036	172.16.131.41/30	172.16.131.42/30
INFRA_VN-Global/San_Jose /Site-3	3034	172.16.131.33/30	172.16.131.34/30
GigabitEthernet0/1/2			
Layer3			
External Domain Protocol	65535		
Virtual Network	Vlan	Local IP	Remote IP
GUEST-Global/San_Jose /Site-3	3040	172.16.131.57/30	172.16.131.58/30
RESEARCH-Global/San_Jose /Site-3	3041	172.16.131.61/30	172.16.131.62/30
CAMPUS-Global/San_Jose /Site-3	3039	172.16.131.53/30	172.16.131.54/30
INFRA_VN-Global/San_Jose /Site-3	3038	172.16.131.49/30	172.16.131.50/30

7. Repeat these steps to collect the information on the second border node.

Procedure 3: Create Layer-3 Connectivity Between Border Nodes and Internet Edge Routers

The next task is to create IP connectivity from the border nodes to Internet edge routers. This must be done for each virtual network that requires connectivity to the Internet, on each Internet edge router, and for each interface between the routers

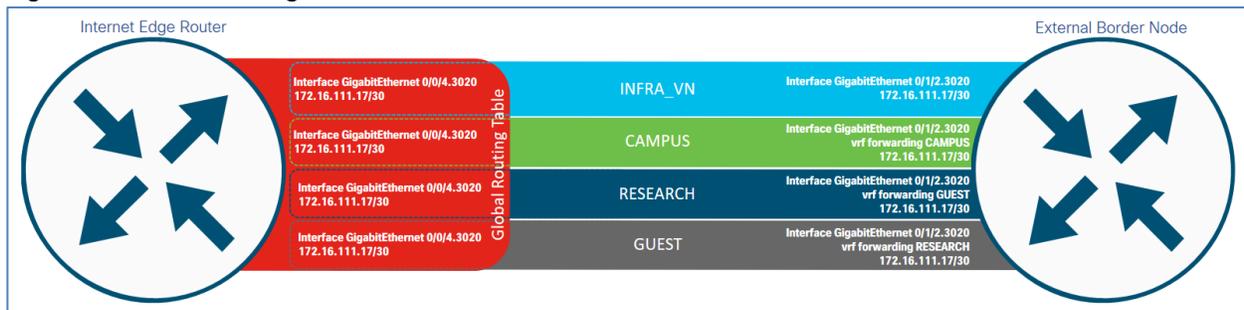
and redundant borders.

Cisco DNA Center has configured the border nodes in previous steps. This configuration will remain unchanged. The Internet edge routers will not become VRF-aware. All SD-Access prefixes (the EID-space) will be learned in their Global Routing Table, and the default route will be advertised back to border nodes via BGP. Using the method in this guide, the BGP configuration on Internet edge routers is not required to use multi-protocol BGP (`address-family vrf` OR `address-family vpv4`) or to form adjacencies using VRFs.

Why Will This Work?

VRF's are locally significant. The automated BGP configuration on the Site-3 border nodes is expecting BGP adjacencies via VRF routing tables. This means that routes learned from the neighbor will be installed into the VRF tables instead of the Global Routing Table. On the other side of the physical link – the Internet edge routers – the BGP neighbor relationship does not need to be formed using VRFs. When a neighbor is defined under [BGP configuration](#), it simply means that the router is going to exchange routes from the VRF or the address family with the defined neighbor.

Figure 25 Internet Edge Router to Border Node – VRF to GRT



The following steps will create subinterfaces on the Internet edge routers. These are used for connectivity and BGP peering, but will have no VRF-awareness.

1. On the first Internet edge router, create a subinterface on the link connected to the first border node.
2. Use the corresponding VLAN and IP address for the `INFRA_VN` from the `View Device Info` from earlier steps.

```
interface GigabitEthernet0/0/4
  no shutdown

interface GigabitEthernet0/0/4.3026
  description Connected to ASR-1002x-03 | GRT
  encapsulation dot1q 3026
  ip address 172.16.131.2 255.255.255.252
```

3. Create the remaining subinterfaces defining the appropriate VLAN and IP address. Each subinterface will be in the GRT but will peer with a VRF on the border node.

```
interface GigabitEthernet0/0/4.3027
  description Connected to ASR-1002x-03 | VRF CAMPUS
  encapsulation dot1q 3027
  ip address 172.16.131.6 255.255.255.252

interface GigabitEthernet0/0/4.3028
  description Connected to ASR-1002x-03 | VRF GUEST
  encapsulation dot1q 3028
  ip address 172.16.131.10 255.255.255.252

interface GigabitEthernet0/0/4.3029
  description Connected to ASR-1002x-03 | VRF RESEARCH
  encapsulation dot1q 3029
  ip address 172.16.131.14 255.255.255.252
```

4. Repeat these steps for the second interface on the Internet edge router which is connected to the second border node.

```
interface GigabitEthernet0/0/5
```

```

no shutdown

interface GigabitEthernet0/0/5.3034
description Connected to ASR-1002x-03 | GRT
encapsulation dot1Q 3034
ip address 172.16.131.34 255.255.255.252

interface GigabitEthernet0/0/5.3035
description Connected to ASR-1002x-03 | VRF CAMPUS
encapsulation dot1Q 3035
ip address 172.16.131.38 255.255.255.252

interface GigabitEthernet0/0/5.3036
description Connected to ASR-1002x-03 | VRF GUEST
encapsulation dot1Q 3036
ip address 172.16.131.42 255.255.255.252

interface GigabitEthernet0/0/5.3037
description Connected to ASR-1002x-03 | VRF RESEARCH
encapsulation dot1Q 3037
ip address 172.16.131.46 255.255.255.252

```

5. Test the IP reachability between the first Internet edge router and the two border nodes.

```

ping 172.16.131.1
ping 172.16.131.5
ping 172.16.131.9
ping 172.16.131.13

ping 172.16.131.33
ping 172.16.131.37
ping 172.16.131.41
ping 172.16.131.245

```

6. Complete a similar corresponding configuration on the first interface on second Internet edge router.

```

interface GigabitEthernet0/0/4
no shutdown

interface GigabitEthernet0/0/4.3038
description Connected to ASR-1002x-04 | GRT
encapsulation dot1Q 3038
ip address 172.16.131.50 255.255.255.252

interface GigabitEthernet0/0/4.3039
description Connected to ASR-1002x-04 | VRF CAMPUS
encapsulation dot1Q 3039
ip address 172.16.131.54 255.255.255.252

interface GigabitEthernet0/0/4.3040
description Connected to ASR-1002x-04 | VRF GUEST
encapsulation dot1Q 3040
ip address 172.16.131.58 255.255.255.252

interface GigabitEthernet0/0/4.3041
description Connected to ASR-1002x-04 | VRF RESEARCH
encapsulation dot1Q 3041
ip address 172.16.131.62 255.255.255.252

```

7. Complete a similar corresponding configuration on the second interface on second Internet edge router.

```

interface GigabitEthernet0/0/5
no shutdown

interface GigabitEthernet0/0/5.3030
description Connected to ASR-1002x-03 | GRT
encapsulation dot1Q 3030
ip address 172.16.131.18 255.255.255.252

interface GigabitEthernet0/0/5.3031

```

```

description Connected to ASR-1002x-03 | VRF CAMPUS
encapsulation dot1Q 3031
ip address 172.16.131.22 255.255.255.252

interface GigabitEthernet0/0/5.3032
description Connected to ASR-1002x-03 | VRF GUEST
encapsulation dot1Q 3032
ip address 172.16.131.26 255.255.255.252

interface GigabitEthernet0/0/5.3033
description Connected to ASR-1002x-03 | VRF RESEARCH
encapsulation dot1Q 3033
ip address 172.16.131.30 255.255.255.252

```

8. Test the IP reachability between the second Internet edge router and the two border nodes.

```

ping 172.16.131.49
ping 172.16.131.53
ping 172.16.131.57
ping 172.16.131.61

ping 172.16.131.17
ping 172.16.131.21
ping 172.16.131.25
ping 172.16.131.29

```

Procedure 4: Create BGP Adjacencies Between Border Nodes and Internet Edge Routers

Since IP connectivity has been established and verified, BGP adjacencies can be created between the Internet edge routers and the border nodes.

On the border nodes, a subinterface was provisioned that forwards for a VRF along with a BGP configuration listening for adjacencies using that subinterface as the update source. A corresponding configuration will be applied on the Internet edge routers although the adjacency is through the GRT.

1. Create the BGP process on the Internet edge router.
2. Use the corresponding Autonomous-System defined in the IP-based transit in early steps. As a recommended practice, the Loopback 0 interface is used as the BGP router ID.

```

router bgp 65535
  bgp router-id interface Loopback0

```

3. Define the neighbor and its corresponding AS Number.
4. Define the update source to use the appropriate subinterface.
5. Activate the exchange of NLRI with the first border node.

```

neighbor 172.16.131.1 remote-as 65003
neighbor 172.16.131.1 update-source GigabitEthernet0/0/4.3026
neighbor 172.16.131.1 activate

```

6. Configure the remaining adjacencies using the applicable subinterfaces as the update-source for the peering.

```

neighbor 172.16.131.5 remote-as 65003
neighbor 172.16.131.5 update-source GigabitEthernet0/0/4.3027
neighbor 172.16.131.5 activate

neighbor 172.16.131.9 remote-as 65003
neighbor 172.16.131.9 update-source GigabitEthernet0/0/4.3028
neighbor 172.16.131.9 activate

neighbor 172.16.131.13 remote-as 65003
neighbor 172.16.131.13 update-source GigabitEthernet0/0/4.3029
neighbor 172.16.131.13 activate

```

7. Continue these steps to build the adjacencies and activate the exchange of NLRI with the second border node.

```

neighbor 172.16.131.33 remote-as 65003
neighbor 172.16.131.33 update-source GigabitEthernet0/0/5.3034

```

```

neighbor 172.16.131.33 activate

neighbor 172.16.131.37 remote-as 65003
neighbor 172.16.131.37 update-source GigabitEthernet0/0/5.3035
neighbor 172.16.131.37 activate

neighbor 172.16.131.41 remote-as 65003
neighbor 172.16.131.41 update-source GigabitEthernet0/0/5.3036
neighbor 172.16.131.41 activate

neighbor 172.16.131.45 remote-as 65003
neighbor 172.16.131.45 update-source GigabitEthernet0/0/5.3037
neighbor 172.16.131.45 activate

```

8. Verify that the adjacencies have formed with each border node.

```
ASR-1002x-01# show ip bgp summary
```

```

BGP router identifier 10.254.123.200, local AS number 65535
BGP table version is 11, main routing table version 11
9 network entries using 2232 bytes of memory
25 path entries using 3400 bytes of memory
5/2 BGP path/bestpath attribute entries using 1400 bytes of memory
1 BGP AS-PATH entries using 24 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 7056 total bytes of memory
BGP activity 9/0 prefixes, 25/0 paths, scan interval 60 secs

```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
10.10.10.2	4	65535	7368	7383	11	0	0	4d15h	8
172.16.131.1	4	65003	7367	7383	11	0	0	4d15h	2
172.16.131.5	4	65003	7375	7375	11	0	0	4d15h	2
172.16.131.9	4	65003	7357	7372	11	0	0	4d15h	2
172.16.131.13	4	65003	7362	7375	11	0	0	4d15h	2
172.16.131.33	4	65003	7370	7378	11	0	0	4d15h	2
172.16.131.37	4	65003	7372	7374	11	0	0	4d15h	2
172.16.131.41	4	65003	7369	7369	11	0	0	4d15h	2
172.16.131.45	4	65003	7382	7371	11	0	0	4d15h	2

```
ASR-1002x-01#
```

9. Complete the corresponding configuration on the second Internet edge router.

```

router bgp 65535
  bgp router-id interface Loopback0

  neighbor 172.16.131.49 remote-as 65003
  neighbor 172.16.131.49 update-source GigabitEthernet0/0/4.3038
  neighbor 172.16.131.49 activate

  neighbor 172.16.131.53 remote-as 65003
  neighbor 172.16.131.53 update-source GigabitEthernet0/0/4.3039
  neighbor 172.16.131.53 activate

  neighbor 172.16.131.57 remote-as 65003
  neighbor 172.16.131.57 update-source GigabitEthernet0/0/4.3040
  neighbor 172.16.131.57 activate

  neighbor 172.16.131.61 remote-as 65003
  neighbor 172.16.131.61 update-source GigabitEthernet0/0/4.3041
  neighbor 172.16.131.61 activate

  neighbor 172.16.131.17 remote-as 65003
  neighbor 172.16.131.17 update-source GigabitEthernet0/0/5.3030
  neighbor 172.16.131.17 activate

  neighbor 172.16.131.21 remote-as 65003
  neighbor 172.16.131.21 update-source GigabitEthernet0/0/5.3031
  neighbor 172.16.131.21 activate

```

```

neighbor 172.16.131.25 remote-as 65003
neighbor 172.16.131.25 update-source GigabitEthernet0/0/5.3032
neighbor 172.16.131.25 activate

neighbor 172.16.131.29 remote-as 65003
neighbor 172.16.131.29 update-source GigabitEthernet0/0/5.3033
neighbor 172.16.131.29 activate

```

10. Verify that the adjacencies have formed with each border node.

```

ASR-1002x-02# show ip bgp summary

BGP router identifier 10.254.24.2, local AS number 65535
BGP table version is 11, main routing table version 11
9 network entries using 2232 bytes of memory
25 path entries using 3400 bytes of memory
5/2 BGP path/bestpath attribute entries using 1400 bytes of memory
1 BGP AS-PATH entries using 24 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 7056 total bytes of memory
BGP activity 9/0 prefixes, 25/0 paths, scan interval 60 secs

Neighbor      V      AS MsgRcvd MsgSent  TblVer  InQ  OutQ  Up/Down  State/PfxRcd
10.10.10.1    4      65535   7386   7372    11    0    0  4d15h    8
172.16.131.17 4      65003   7379   7374    11    0    0  4d15h    2
172.16.131.21 4      65003   7370   7372    11    0    0  4d15h    2
172.16.131.25 4      65003   7370   7380    11    0    0  4d15h    2
172.16.131.29 4      65003   7369   7376    11    0    0  4d15h    2
172.16.131.49 4      65003   7368   7370    11    0    0  4d15h    2
172.16.131.53 4      65003   7366   7378    11    0    0  4d15h    2
172.16.131.57 4      65003   7372   7370    11    0    0  4d15h    2
172.16.131.61 4      65003   7381   7381    11    0    0  4d15h    2
ASR-1002x-02#

```



The **10.10.10.1** and **10.10.10.2** entries in the CLI outputs above are a result of the IBGP configuration between the redundant Internet edge routers to provide a packet forwarding path in the event of link failure or upstream device failure, as described in earlier procedures. Because the Internet Edge routers are not VRF aware, only a single peering was needed between the pair.

About Default Route Advertisement Using BGP

There are four different methods to advertise a default route in BGP – each with its own caveats.

The first three methods are very similar and result in the same effect: a default route is injected into the BGP RIB resulting in a general advertisement to all BGP neighbors. The fourth method selectively advertises the default route to a specific neighbor.

1. *network 0.0.0.0*

- This will inject the default route into BGP if there is a default route present in the Global Routing Table.
- The route is then advertised to all configured neighbors.

2. *redistribute <IGP>*

- The default route must currently be in the Global Routing table **and** be learned from that routing protocol that is being redistributed (example: redistributing IS-IS into BGP and the default route learned via IS-IS).
- This will inject the default route into BGP and then advertise it to all configured neighbors.

3. *default-information originate*

- This will cause the default route to be artificially generated and injected into the BGP RIB regardless of whether or not it is present in the Global Routing Table.
- The *default-information originate* command alone is not enough to trigger the default route advertisement to all configured neighbors. In current Cisco software versions, an explicit redistribution of *0.0.0.0/0* is required to trigger the default route to be advertised via BGP when using this method.

4. *neighbor X.X.X.X default-originate*

- This command is different in that it will only advertise the default route to a specified neighbor while each of the previous three methods will advertise the default route to all neighbors.
- The default route is artificially generated and injected into BGP similarly to *default-information originate*.
- The default route will not be present in the BGP RIB of the router where this command is configured. This prevents its advertisement to all neighbors.



The origin (Origin Code) of that route is the key difference between these first three methods. Origin is a [well-known mandatory](#) attribute. This means that this attribute of BGP must be recognized by all BGP routers, present in all BGP updates, and passed on to other BGP routers in advertisements. Origin is also the fifth item in the [BGP Best Path Algorithm](#) implemented by Cisco software. A different origin code will have a subtle influence on the BGP Best Path decision process and therefore must be considered for creating deterministic convergence.

Each approach has its own benefits and considerations. The *neighbor X.X.X.X default-originate* option is particularly useful in cases with stub-autonomous systems – an AS where a full BGP table is not required. It also provides the most granularity and control over the default route advertisement, as any other method will advertise the default route to all neighbors. With the number of possible variations in upstream router configuration and peering, fine-grained granularity is needed, and thus it is the method used in this prescriptive deployment guide.

Procedure 5: Advertise the Default Route From the Internet Edge Routers to Fabric VN

The following steps will use the fourth method discussed above to advertise the default route from the Internet edge routers to the fabric border nodes.



Exhaustive coverage of BGP is beyond the scope of this validation. BGP peering with upstream routers generally uses filter-lists to block routes such as an ill-configured default route advertisement that might create routing loops.

Whichever method is used to advertise a default route using BGP must be carefully considered based on the needs and design of the deployment. Please use the [PPDIOO Lifecycle Approach](#) to the BGP design and implementation.

1. On the first Internet edge router, advertise the default route to the first border node using BGP.

```
router bgp 65535

address-family ipv4
neighbor 172.16.131.1 default-originate
neighbor 172.16.131.5 default-originate
neighbor 172.16.131.9 default-originate
neighbor 172.16.131.13 default-originate
```

2. Advertise the default route to the second border node.

```
neighbor 172.16.131.33 default-originate
neighbor 172.16.131.37 default-originate
neighbor 172.16.131.41 default-originate
neighbor 172.16.131.45 default-originate
```

3. Repeat these steps on the second Internet edge router.

```
router bgp 65535

address-family ipv4
neighbor 172.16.131.49 default-originate
neighbor 172.16.131.53 default-originate
neighbor 172.16.131.57 default-originate
neighbor 172.16.131.61 default-originate

neighbor 172.16.131.17 default-originate
neighbor 172.16.131.21 default-originate
neighbor 172.16.131.25 default-originate
neighbor 172.16.131.29 default-originate
```

Procedure 6: Verify Default Route Advertisement

Once the default route is advertised from the upstream routers to the connected-to-Internet border nodes, it will be present in the VRF tables of those border nodes.

1. On each connected-to-Internet border node, verify the presence of the default route in each VRF.

```
ASR-1002x-03# show ip route vrf CAMPUS

Routing Table: CAMPUS
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       a - application route
       + - replicated route, % - next hop override, p - overrides from PFR

Gateway of last resort is 172.16.131.6 to network 0.0.0.0
```

```

B* 0.0.0.0/0 [20/0] via 172.16.131.6, 3w2d
    172.16.0.0/16 is variably subnetted, 10 subnets, 3 masks
B    172.16.112.0/24 [20/10] via 192.168.200.100, 1w5d
B    172.16.122.0/24 [20/10] via 192.168.200.100, 1w5d
C    172.16.131.4/30 is directly connected, GigabitEthernet0/1/2.3027
L    172.16.131.5/32 is directly connected, GigabitEthernet0/1/2.3027
C    172.16.131.20/30 is directly connected, GigabitEthernet0/1/3.3031
L    172.16.131.21/32 is directly connected, GigabitEthernet0/1/3.3031
B    172.16.132.0/24 [200/0] via 192.168.30.1, 3w6d
C    172.16.132.1/32 is directly connected, Loopback1033
B    172.16.152.0/24 [20/10] via 192.168.200.200, 1w5d
B    172.16.252.105/32 [20/10] via 192.168.200.200, 00:11:14
B    198.51.100.0/24 [20/10] via 192.168.200.200, 1d00h

ASR-1002x-03#

```

```
ASR-1002x-03# show ip route vrf RESEARCH
```

Routing Table: RESEARCH

```

Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       a - application route
       + - replicated route, % - next hop override, p - overrides from PfR

```

Gateway of last resort is 172.16.131.14 to network 0.0.0.0

```

B* 0.0.0.0/0 [20/0] via 172.16.131.14, 3w2d
    172.16.0.0/16 is variably subnetted, 9 subnets, 3 masks
B    172.16.114.0/24 [20/10] via 192.168.200.100, 1w5d
B    172.16.124.0/24 [20/10] via 192.168.200.100, 1w5d
C    172.16.131.12/30 is directly connected, GigabitEthernet0/1/2.3029
L    172.16.131.13/32 is directly connected, GigabitEthernet0/1/2.3029
C    172.16.131.28/30 is directly connected, GigabitEthernet0/1/3.3033
L    172.16.131.29/32 is directly connected, GigabitEthernet0/1/3.3033
B    172.16.134.0/24 [200/0] via 192.168.30.1, 3w6d
C    172.16.134.1/32 is directly connected, Loopback1035
B    172.16.154.0/24 [20/10] via 192.168.200.100, 1w5d
B    198.51.100.0/24 [20/10] via 192.168.200.200, 1d00h

ASR-1002x-03#

```

```
ASR-1002x-03# show ip route vrf GUEST
```

Routing Table: GUEST

```

Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       a - application route
       + - replicated route, % - next hop override, p - overrides from PfR

```

Gateway of last resort is 172.16.131.10 to network 0.0.0.0

```

B* 0.0.0.0/0 [20/0] via 172.16.131.10, 3w2d
    172.16.0.0/16 is variably subnetted, 9 subnets, 3 masks
B    172.16.113.0/24 [20/10] via 192.168.200.100, 1w5d
B    172.16.123.0/24 [20/10] via 192.168.200.100, 1w5d
C    172.16.131.8/30 is directly connected, GigabitEthernet0/1/2.3028

```

```

L      172.16.131.9/32 is directly connected, GigabitEthernet0/1/2.3028
C      172.16.131.24/30 is directly connected, GigabitEthernet0/1/3.3032
L      172.16.131.25/32 is directly connected, GigabitEthernet0/1/3.3032
B      172.16.133.0/24 [200/0] via 192.168.30.1, 3w6d
C      172.16.133.1/32 is directly connected, Loopback1034
B      172.16.153.0/24 [20/10] via 192.168.200.100, 1w5d
B      198.51.100.0/24 [20/10] via 192.168.200.200, 1d00h
ASR-1002x-03#

```

2. Verify the presence of the default route in the Global Routing Table.

```

ASR-1002x-03# show ip route bgp
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       a - application route
       + - replicated route, % - next hop override, p - overrides from PfR

Gateway of last resort is 172.16.131.2 to network 0.0.0.0

B*    0.0.0.0/0 [20/0] via 172.16.131.2, 3w2d
      172.16.0.0/16 is variably subnetted, 7 subnets, 3 masks
B      172.16.130.0/24 [200/0] via 192.168.30.1, 3w6d
ASR-1002x-03#

```

```

ASR-1002x-03# show ip route o.o.o.o
Routing entry for 0.0.0.0/0, supernet
  Known via "bgp 65003", distance 20, metric 0, candidate default path
  Tag 65535, type external
  Last update from 172.16.131.2 3w2d ago
Routing Descriptor Blocks:
  * 172.16.131.2, from 172.16.131.2, 3w2d ago
    Route metric is 0, traffic share count is 1
    AS Hops 1
    Route tag 65535
    MPLS label: none
    MPLS Flags: NSF
ASR-1002x-03#

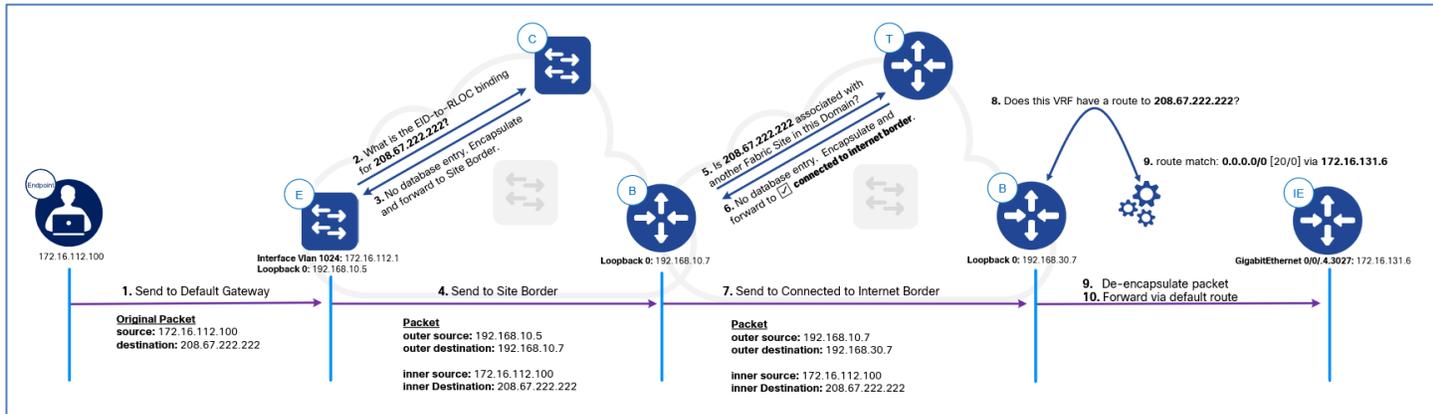
```

Procedure 7: Verify Default Route Via LISP – Advanced and Optional

Troubleshooting and verification of LISP requires a strong understanding of the protocol, its interaction with the RIB and CEF, and a solid understanding and familiarity with the command line interface. Exhaustive study of LISP is beyond the scope of this guide. However, some relatively straightforward **show** commands can be used to verify that the remaining fabric sites know about the Site-3 **connected-to-Internet** border node and will send packets destined for the Internet to this site.

In SD-Access for Distributed Campus, the **Is this site Connected to Internet?** option will advertise that border node to all other sites as a gateway of last resort for the fabric domain. It will also configure the other sites' border nodes to query the transit control plane node for unknown destinations.

Figure 26 Packet Walk Traffic Flow



The following steps will demonstrate how to validate the packet flow in the diagram above hop-by-hop. For the purposes of this demonstration, `LIG` will be used to manually trigger the control plane lookup. Each step in the packet walk will be shown along with commentary on what information to glean from the resulting CLI output.

1. On the edge node, display the current LISP map cache.

```
3850-24P-01# show ip lisp map-cache instance-id 4099

LISP IPv4 Mapping Cache for EID-table vrf CAMPUS (IID 4099), 2 entries
0.0.0.0/0, uptime: 00:00:50, expires: never, via static-send-map-request
Negative cache entry, action: send-map-request
172.16.112.0/24, uptime: 00:00:50, expires: never, via dynamic-EID, send-map-request
Negative cache entry, action: send-map-request

3850-24P-01#
```

There is no map-cache entry for the desired prefix 208.67.222.222, and the only entry covering that prefix is the LISP default entry 0.0.0.0/0.

2. Use `lig` to trigger a control plane lookup.

```
3850-24P-01# lig instance-id 4099 208.67.222.222

Mapping information for EID 208.67.222.222 from 192.168.10.1 with RTT 2 msec
208.0.0.0/4, uptime: 00:00:00, expires: 23:59:59, via map-reply, forward-native
Encapsulating to proxy ETR

3850-24P-01#
```

LISP computes this aggregate (208.0.0.0/4) as the largest possible block that does not contain known EID prefixes but does cover the queried address. This is due to LISP's history as a method to solve the DFZ and TCAM space concerns.

3. Redisplay the LISP map-cache.

```
3850-24P-01# show ip lisp map-cache instance-id 4099

LISP IPv4 Mapping Cache for EID-table vrf CAMPUS (IID 4099), 3 entries
0.0.0.0/0, uptime: 00:13:55, expires: never, via static-send-map-request
Negative cache entry, action: send-map-request
172.16.112.0/24, uptime: 00:13:55, expires: never, via dynamic-EID, send-map-request
Negative cache entry, action: send-map-request
208.0.0.0/4, uptime: 00:12:38, expires: 23:47:21, via map-reply, forward-native
Encapsulating to proxy ETR

3850-24P-01#
```

The map-cache entry has a state of *forward-native* but an action of *Encapsulating to proxy ETR*. What is the result when

the packet is forwarded? To understand the forwarding decision, the LISP configuration and CEF tables must be observed.

4. Verify a proxy-ETR is configured on the edge node for the instance ID.

```
3850-24P-01# show running-config | section router lisp

router lisp
 locator-table default
 locator-set rloc_26ed60b7-cf1f-4310-bfa1-571b6846c75d
  IPv4-interface Loopback0 priority 10 weight 10
  exit-locator-set
 !
 locator default-set rloc_26ed60b7-cf1f-4310-bfa1-571b6846c75d
 service ipv4
  encapsulation vxlan
  itr map-resolver 192.168.10.1
  itr map-resolver 192.168.10.2
  etr map-server 192.168.10.1 key 7 105B0A10
  etr map-server 192.168.10.1 proxy-reply
  etr map-server 192.168.10.2 key 7 02130752
  etr map-server 192.168.10.2 proxy-reply
  etr
  sgt
  no map-cache away-eids send-map-request
  use-petr 192.168.10.7
  use-petr 192.168.10.8
  proxy-itr 192.168.10.5
  exit-service-ipv4
 !
```

The Site-1 border nodes (**192.168.10.7** and **192.168.10.8**) are configured as proxy-ETR entries. This configuration is done at the default *service ipv4* level under *router lisp* and therefore inherited by all instance IDs. With the presence of this configuration, the edge node should forward packets destined for 208.67.222.222 to the Site-1 borders based on the LISP map cache entry.

5. Verify the logical forwarding path for the LISP map cache entry.

```
3850-24P-01# show ip cef vrf CAMPUS 208.67.222.222 detail

208.0.0.0/4, epoch 1, flags [subtree context, check lisp eligibility], per-destination sharing
 SC owned,sourced: LISP remote EID - locator status bits 0x00000000
 LISP remote EID: 2 packets 1152 bytes fwd action encap
 LISP source path list
  nexthop 192.168.10.7 LISP0.4099
  nexthop 192.168.10.8 LISP0.4099
 2 IPL sources [no flags]
 nexthop 192.168.10.7 LISP0.4099
 nexthop 192.168.10.8 LISP0.4099
```

From the perspective of CEF, the Site-1 border node will check to see if the packet meets the LISP eligibility checks. If met, the packet will be forwarded using the LISP virtual interface 4099. What is the physical forwarding path for this prefix?

6. Verify the physical forwarding path for the LISP map cache entry.

```
3850-24P-01# show ip cef vrf CAMPUS 208.67.222.222 internal

208.0.0.0/4, epoch 1, flags [sc, lisp elig], refcnt 6, per-destination sharing
 sources: LISP, IPL
 feature space:
  Broker: linked, distributed at 1st priority
 subblocks:
  SC owned,sourced: LISP remote EID - locator status bits 0x00000000
  LISP remote EID: 2 packets 1152 bytes fwd action encap
  LISP source path list
   path list FF97AB0AA8, 4 locks, per-destination, flags 0x49 [shble, rif, hwn]
   ifnums:
    LISP0.4099(45): 192.168.10.7, 192.168.10.8
  2 paths
   path FF97AB26C0, share 1/1, type attached nexthop, for IPv4
    nexthop 192.168.10.7 LISP0.4099, IP midchain out of LISP0.4099, addr 192.168.10.7 FF9BDF4740
   path FF97AB2A00, share 1/1, type attached nexthop, for IPv4
    nexthop 192.168.10.8 LISP0.4099, IP midchain out of LISP0.4099, addr 192.168.10.8 FF9BDF4500
```

```

1 output chain
  chain[0]: loadinfo FF9714B030, per-session, 2 choices, flags 0003, 5 locks
            flags [Per-session, for-rx-IPv4]
            16 hash buckets
              < 0 > IP midchain out of LISP0.4099, addr 192.168.10.7 FF9BDF4740
                  IP adj out of GigabitEthernet1/1/2, addr 10.10.45.4 FF9A009480
              < 1 > IP midchain out of LISP0.4099, addr 192.168.10.8 FF9BDF4500
                  IP adj out of GigabitEthernet1/1/2, addr 10.10.45.4 FF9A009480

```

From the perspective of CEF, the edge node will encapsulate the packet as it is LISP eligible and send it from the interface **LISP0.4099** with a destination of either **192.168.10.7** or **192.168.10.8**. To reach either of these IP addresses, the **GigabitEthernet 1/1/2** interface is used with a next-hop router of **10.10.45.4**.

7. On the border node, display the current LISP map cache.

```

ISR-4451-07# show ip lisp map-cache instance-id 4099

LISP IPv4 Mapping Cache for EID-table vrf CAMPUS (IID 4099), 3 entries

0.0.0.0/0, uptime: 4d17h, expires: never, via static-send-map-request
  Negative cache entry, action: send-map-request
172.16.112.0/24, uptime: 4d17h, expires: never, via route-import, self, send-map-request
  Negative cache entry, action: send-map-request
198.51.100.0/24, uptime: 03:50:32, expires: never, via route-import, send-map-request
  Negative cache entry, action: send-map-request

ISR-4451-07#

```

There is no map-cache entry for the desired prefix 208.67.222.222, and the only entry covering that prefix is the LISP default entry 0.0.0.0/0.

8. Use lig to trigger a control plane lookup.

```

ISR-4451-07# lig instance-id 4099 208.67.222.222

Mapping information for EID 208.67.222.222 from 192.168.200.100 with RTT 2 msec
208.0.0.0/4, uptime: 00:00:00, expires: 00:14:59, via map-reply, forward-native
  Encapsulating to proxy ETR

ISR-4451-07#

```

LISP computes this aggregate (208.0.0.0/4) as the largest possible block that does not contain known EID prefixes but does cover the queried address. This is due to LISP's history as a method to solve the DFZ and TCAM space concerns.

9. Redisplay the LISP map-cache.

```

ISR-4451-07# show ip lisp map-cache instance-id 4099

LISP IPv4 Mapping Cache for EID-table vrf CAMPUS (IID 4099), 7 entries

0.0.0.0/0, uptime: 5d17h, expires: never, via static-send-map-request
  Negative cache entry, action: send-map-request
172.16.112.0/24, uptime: 5d17h, expires: never, via route-import, self, send-map-request
  Negative cache entry, action: send-map-request
172.16.122.0/24, uptime: 19:01:12, expires: never, via route-import, send-map-request
  Negative cache entry, action: send-map-request
172.16.132.0/24, uptime: 19:01:12, expires: never, via route-import, send-map-request
  Negative cache entry, action: send-map-request
172.16.152.0/24, uptime: 19:01:12, expires: never, via route-import, send-map-request
  Negative cache entry, action: send-map-request
198.51.100.0/24, uptime: 00:38:08, expires: never, via route-import, send-map-request
  Negative cache entry, action: send-map-request
208.0.0.0/4, uptime: 00:00:03, expires: 00:14:56, via map-reply, forward-native
  Encapsulating to proxy ETR

ISR-4451-07#

```

The mapping-source of this map-cache entry can be displayed to confirm that it was sourced from the transit control plane nodes.

10. Verify the source of the map-cache entry.

```
ISR-4451-07# show ip lisp instance-id 4099 map-cache 208.67.222.222

LISP IPv4 Mapping Cache for EID-table vrf CAMPUS (IID 4099), 7 entries

208.0.0.0/4, uptime: 00:01:23, expires: 00:13:36, via map-reply, forward-native
Sources: map-reply
State: forward-native, last modified: 00:01:23, map-source: 192.168.200.200
Active, Packets out: 0(0 bytes)
Encapsulating to proxy ETR
ISR-4451-07#
```

The map-cache entry has a state of **forward-native** but an action of **Encapsulating to proxy ETR**. What is the result when the packet is forwarded? To understand the forwarding decision, the LISP configuration and CEF tables must be observed.

11. Verify the proxy-ETR and use-petr configuration.

```
ISR-4451-07# show running-config | section router lisp

router lisp
 locator-table default
 locator-set rloc_637e1ea3-f9ec-48c8
  IPv4-interface Loopback0 priority 10 weight 10
  auto-discover-rlocs
  exit-locator-set
!
prefix-list Global/San_Jose/Site-1_SJC_list1
 172.16.110.0/24
 172.16.112.0/24
 172.16.113.0/24
 172.16.114.0/24
 exit-prefix-list
!
service ipv4
 encapsulation vxlan
 itr map-resolver 192.168.10.1
 itr map-resolver 192.168.10.2
 etr map-server 192.168.10.1 key 7 0106050D
 etr map-server 192.168.10.1 proxy-reply
 etr map-server 192.168.10.2 key 7 06130C28
 etr map-server 192.168.10.2 proxy-reply
 etr
 sgt
 no map-cache away-eids send-map-request
 proxy-etr
 proxy-itr 192.168.10.7
 exit-service-ipv4
!
service ethernet
 database-mapping limit dynamic 5000
 itr map-resolver 192.168.10.1
 itr map-resolver 192.168.10.2
 itr
 etr map-server 192.168.10.1 key 7 105B0A10
 etr map-server 192.168.10.1 proxy-reply
 etr map-server 192.168.10.2 key 7 06130C28
 etr map-server 192.168.10.2 proxy-reply
 etr
 exit-service-ethernet
!
instance-id 4097
 remote-rloc-probe on-route-change
 service ipv4
  eid-table default
  database-mapping 172.16.110.0/24 locator-set rloc_637e1ea3-f9ec-48c8 proxy
  map-cache 172.16.110.0/24 map-request
  route-import database bgp 65001 route-map site-local-eids locator-set rloc_637e1ea3-f9ec-48c8 proxy
  route-import prefix-list Global/San_Jose/Site-1_SJC_list1 bgp 65001 route-map site-local-eids
  itr map-resolver 192.168.10.1 prefix-list Global/San_Jose/Site-1_SJC_list1
  itr map-resolver 192.168.10.2 prefix-list Global/San_Jose/Site-1_SJC_list1
  itr map-resolver 192.168.200.100
  itr map-resolver 192.168.200.200
```

```

etr map-server 192.168.200.100 key 7 051E0506
etr map-server 192.168.200.100 proxy-reply
etr map-server 192.168.200.200 key 7 111C1A0C
etr map-server 192.168.200.200 proxy-reply
use-petr 192.168.30.7
use-petr 192.168.30.8
exit-service-ipv4

```

The router itself is configured as a proxy ETR. This configuration is done at the default *service ipv4* level under *router lisp* and would otherwise be inherited by all instance IDs if a more specific entry is not added. With the presence of this configuration, the border is a site-local gateway of last resort.

The Site-3 border nodes (**192.168.30.7** and **192.168.30.8**) are configured as proxy-ETR entries under the instance IDs with the *use-petr* command. If the site-local border node receives a negative map-reply from the transit control plane node (a reply indicating there is no entry in the database), traffic will be forwarded to the *connected-to-Internet* border nodes.

12. Verify the logical forwarding path for the LISP map cache entry.

```

ISR-4451-07# show ip cef vrf CAMPUS 208.67.222.222 detail

208.0.0.0/4, epoch 0, flags [subtree context, check lisp eligibility], per-destination sharing
SC owned,sourced: LISP remote EID - locator status bits 0x00000000
LISP remote EID: 0 packets 0 bytes fwd action encap
LISP source path list
  nexthop 192.168.30.7 LISP0.4099
  nexthop 192.168.30.8 LISP0.4099
2 IPL sources [no flags]
  nexthop 192.168.30.7 LISP0.4099
  nexthop 192.168.30.8 LISP0.4099

```

From the perspective of CEF, the Site-1 border node will check to see if the packet meets the LISP eligibility checks. If met, the packet will be forwarded using the LISP virtual interface 4099. What is the physical forwarding path for this prefix?

13. Verify the physical forwarding path for the LISP map cache entry.

```

ISR-4451-07# show ip cef vrf CAMPUS 208.67.222.222 internal

208.0.0.0/4, epoch 0, flags [sc, lisp elig], refcnt 6, per-destination sharing
sources: LISP, IPL
feature space:
  Broker: linked, distributed at 1st priority
subblocks:
  SC owned,sourced: LISP remote EID - locator status bits 0x00000000
  LISP remote EID: 0 packets 0 bytes fwd action encap
  LISP source path list
  path list 7F1D1629AB18, 4 locks, per-destination, flags 0x49 [shble, rif, hwn]
  ifnums:
  LISP0.4099 (20): 192.168.30.7, 192.168.30.8
  2 paths
  path 7F1D1628EB28, share 1/1, type attached nexthop, for IPv4
    nexthop 192.168.30.7 LISP0.4099, IP midchain out of LISP0.4099, addr 192.168.30.7
  path 7F1D1628ECC8, share 1/1, type attached nexthop, for IPv4
    nexthop 192.168.30.8 LISP0.4099, IP midchain out of LISP0.4099, addr 192.168.30.8
  1 output chain
  chain[0]: loadinfo 80007F1D161F2630, per-session, 2 choices, flags 0003, 5 locks
    flags [Per-session, for-rx-IPv4]
    16 hash buckets
      < 0 > IP midchain out of LISP0.4099, addr 192.168.30.7 7F1D161D8B10
        IP adj out of GigabitEthernet0/0/0, addr 10.10.71.254 7F1D161DC7C8
      < 1 > IP midchain out of LISP0.4099, addr 192.168.30.8 7F1D161D88F8
        IP adj out of GigabitEthernet0/0/0, addr 10.10.71.254 7F1D161DC7C8

```

From the perspective of CEF, the Site-1 border node will encapsulate the packet as it is LISP eligible and send it from the interface *LISP0.4099* with a destination of either **192.168.30.7** or **192.168.30.8**. To reach either of these IP addresses, the *GigabitEthernet 0/0/0* interface is used with a next-hop router of **10.10.71.254** which traverses the Metro-E circuit and SD-Access transit. Once reaching the *connected-to-Internet* border nodes in Site-3, the packet will use the default-route learned from BGP to reach the Internet.

Appendix A: Hardware and Software Code Versions

Table 3 Device Platform, Model, and Software Version

Platform	Model (PID)	Software Code Version
Cisco DNA Center	DN1-HW-APL	Cisco DNA Center 1.2.8
Identity Services Engine	R-ISE-VMS-K9	ISE 2.4 Patch 5
Catalyst 6807 Series	C6807-XL	15.5(1)SY2
Catalyst 6880 Series	C6880-X-LE	15.5(1)SY2
Catalyst 6840 Series	C6840-X-LE-40G	15.5(1)SY2
Catalyst 6832 Series	C6832-X-LE	15.5(1)SY2
Catalyst 9500 Series	C9500-24Q	16.9.2s
Catalyst 9400 Series	C9407R	16.9.2s
Catalyst 9300 Series	C9300-48U	16.9.2s
Catalyst 4500E Series	WS-C4503-E	03.10.02.E
Catalyst 4500x Series	WS-C4500X-16	03.10.02.E
Catalyst 3850 Series	WS-C3850-24P	16.9.2s
Catalyst 3650 Series	WS-C3650-24PS	16.9.2s
ASR 1000x Series	ASR1002-X	16.9.2s
ISR 4451x Series	ISR4451-X/K9	16.9.2s
ISR 4431 Series	ISR4431/K9	16.9.2s
ISR 4351 Series	ISR4351/K9	16.9.2s
ISR 4331 Series	ISR4331/K9	16.9.2s
WLC 8540 Series	AIR-CT8540-K9	8.8.100.0
WLC 5520 Series	AIR-CT5520-K9	8.8.100.0
WLC 3504 Series	AIR-CT3504-K9	8.8.100.0
AP 4800 Series	AIR-AP4800-B-K9	8.8.100.0
AP 3800 Series	AIR-AP3802I-B-K9	8.8.100.0
AP 2800 Series	AIR-AP2802I-B-K9	8.8.100.0
AP 1850 Series	AIR-AP1852I-B-K9	8.8.100.0

Table 4 Cisco DNA Center Package Versions

Package Name - CLI	Package Name - GUI	Software Version
application-policy	Application Policy	2.1.28.170014
assurance	Assurance Base	1.2.10.258
automation-core	NCP Services	2.1.28.60244
base-provision-core	Automation Base	2.1.28.60244
command-runner	Command Runner	2.1.28.60244
device-onboarding	Device Onboarding	2.1.28.60244
dnac-platform	DNAC Platform	1.0.8.8
icap-automation	Automation - Intelligent Capture	2.1.28.60244
image-management	Image Management	2.1.28.60244
ncp-system	NCP Base	2.1.28.60244
ndp-base-analytics	Network Data Platform - Base Analytics	1.1.10.659
ndp-platform	Network Data Platform - Core	1.1.10.1012
ndp-ui	Network Data Platform - Manager	1.1.10.705
network-visibility	Network Controller Platform	2.1.28.60244
path-trace	Path Trace	2.1.28.60244
platform-ui	Cisco DNA Center UI	1.2.10.163
sd-access	SD Access	2.1.28.60244
sensor-assurance	Assurance - Sensor	1.2.10.254
sensor-automation	Automation - Sensor	2.1.28.60244
system	System	1.1.0.754

Appendix B: Additional References

[CVD – Software-Defined Access Design Guide – Solution 1.2 \(December 2018\)](#)

[CVD – Software-Defined Access Deployment Guide – Solution 1.2 \(October 2018\)](#)

[CVD – SD-Access Segmentation Design Guide – May 2018](#)

[Cisco Design Zone – Design Guides](#)

[Cisco Validated Design \(CVDs\) – Cisco Community](#)

[SD-Access Resources – Cisco Community](#)

[SD-Access Wireless Design and Deployment Guide](#)

[Cisco Identity Services Engine Installation Guide, Release 2.4 – Distributed ISE Deployments](#)

[Designing ISE for Scale & High Availability - BRKSEC-3699 Cisco Live 2018 \(Orlando\)](#)

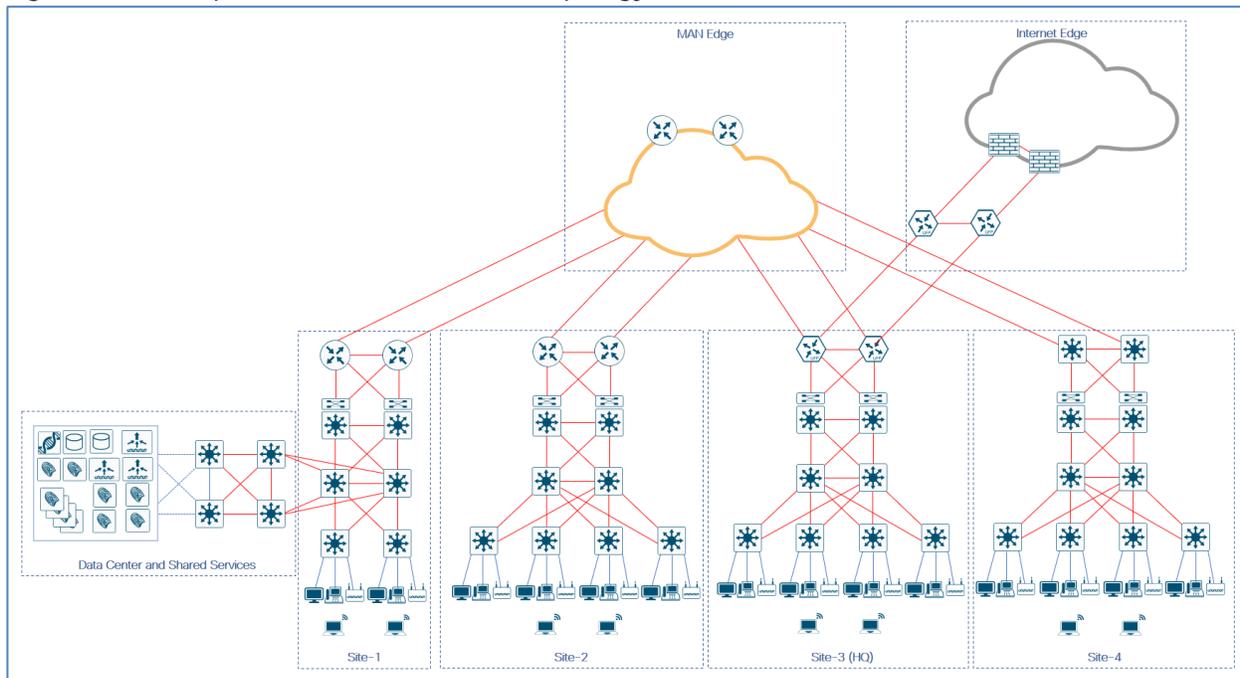
[ISE Security Ecosystem Integration Guides](#)

Appendix C: Advanced Topology Diagrams

High-Level Overview

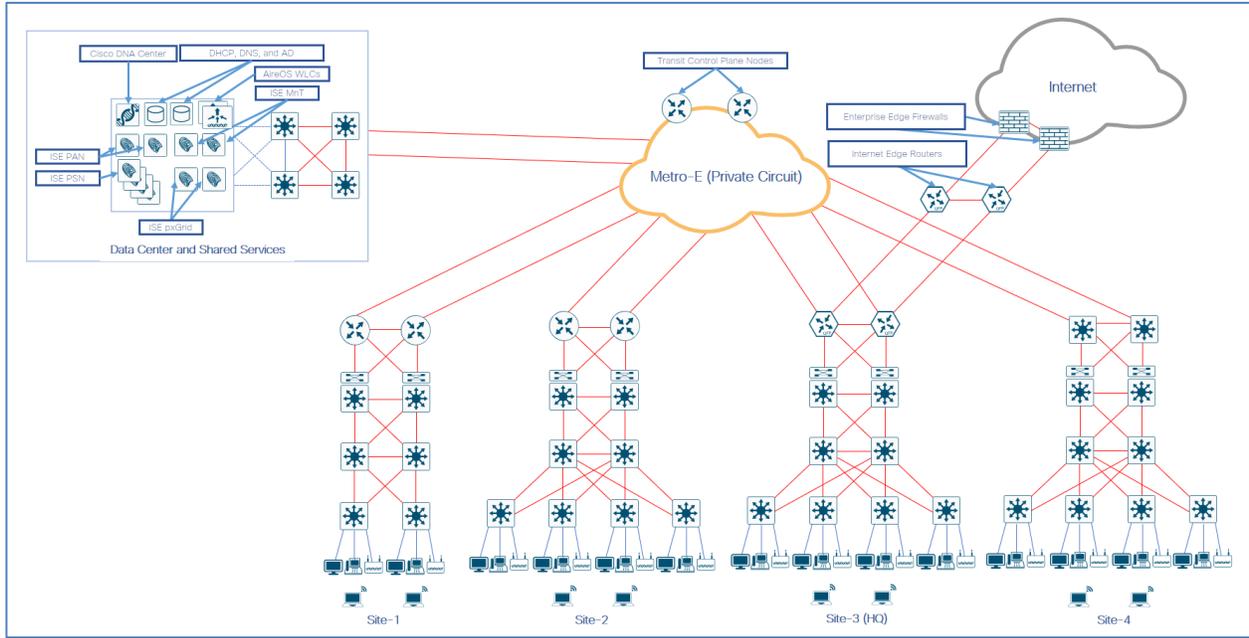
Enterprise Architecture Model Topology

Figure 27 Enterprise Architecture Model Topology



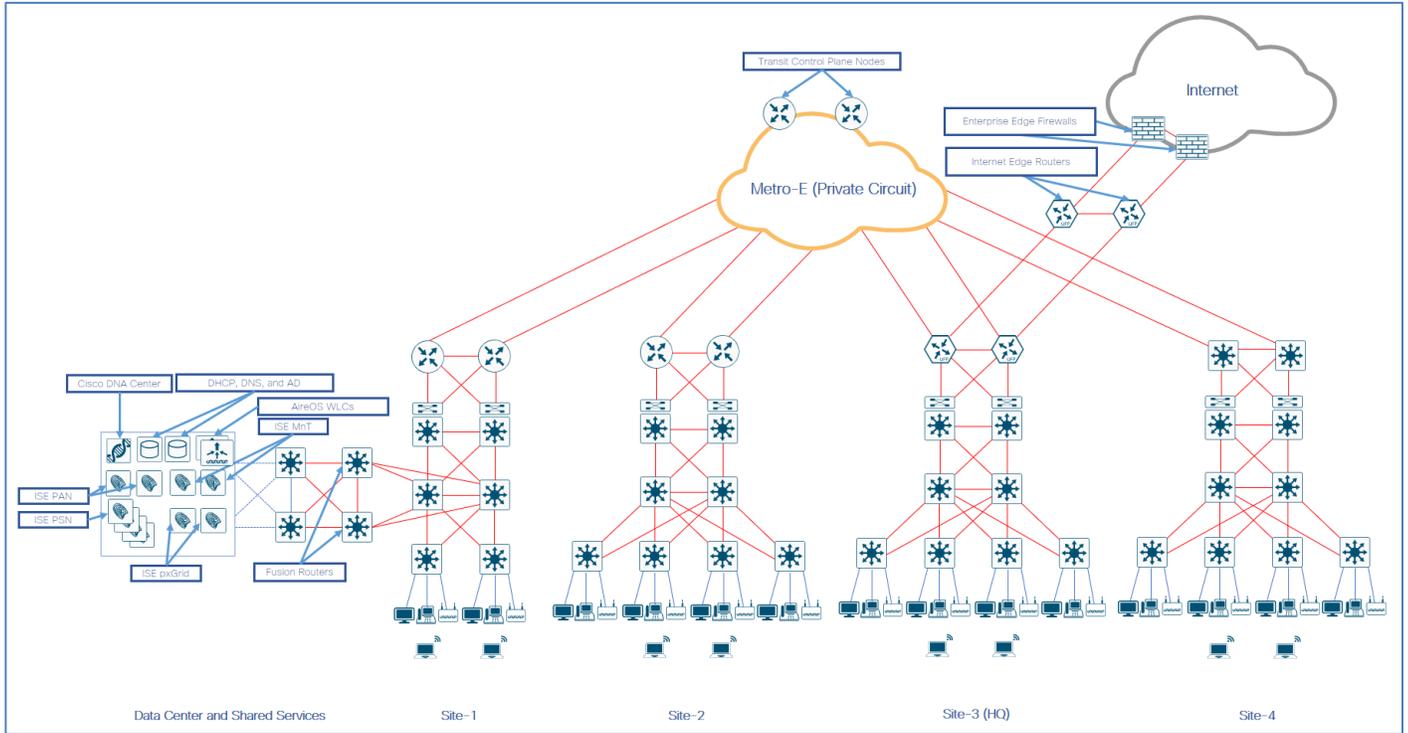
Underlay Connectivity

Figure 28 Underlay Topology



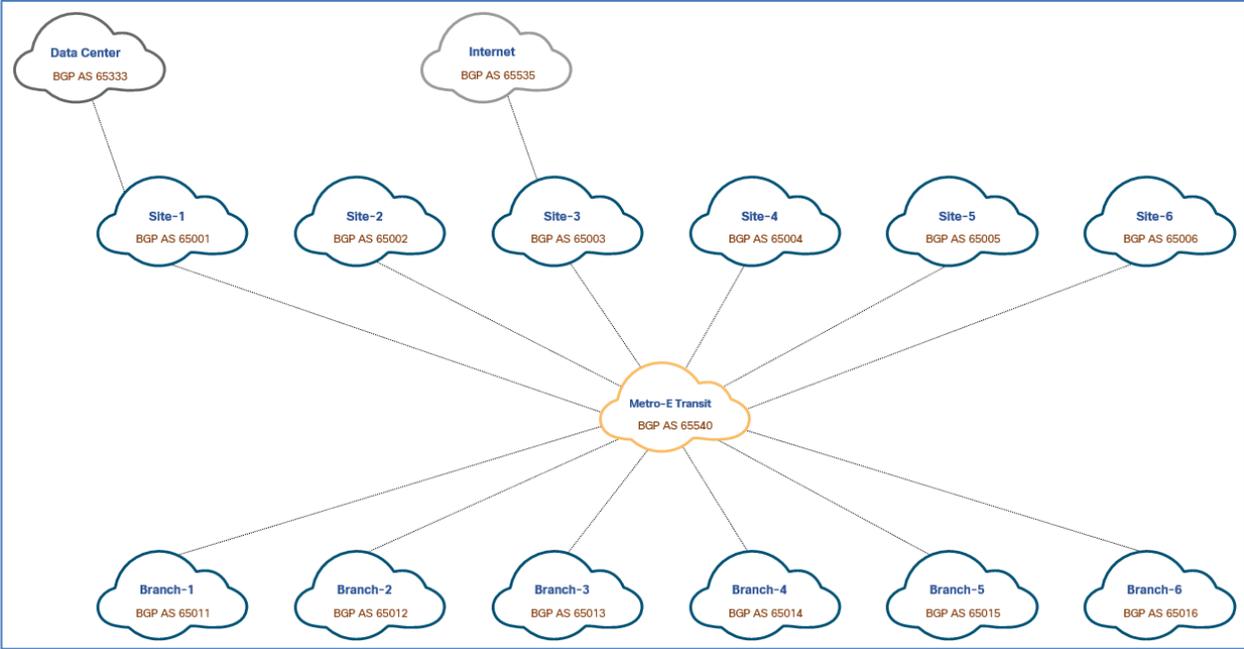
Overlay Connectivity

Figure 29 Overlay Topology



BGP Autonomous System Overview

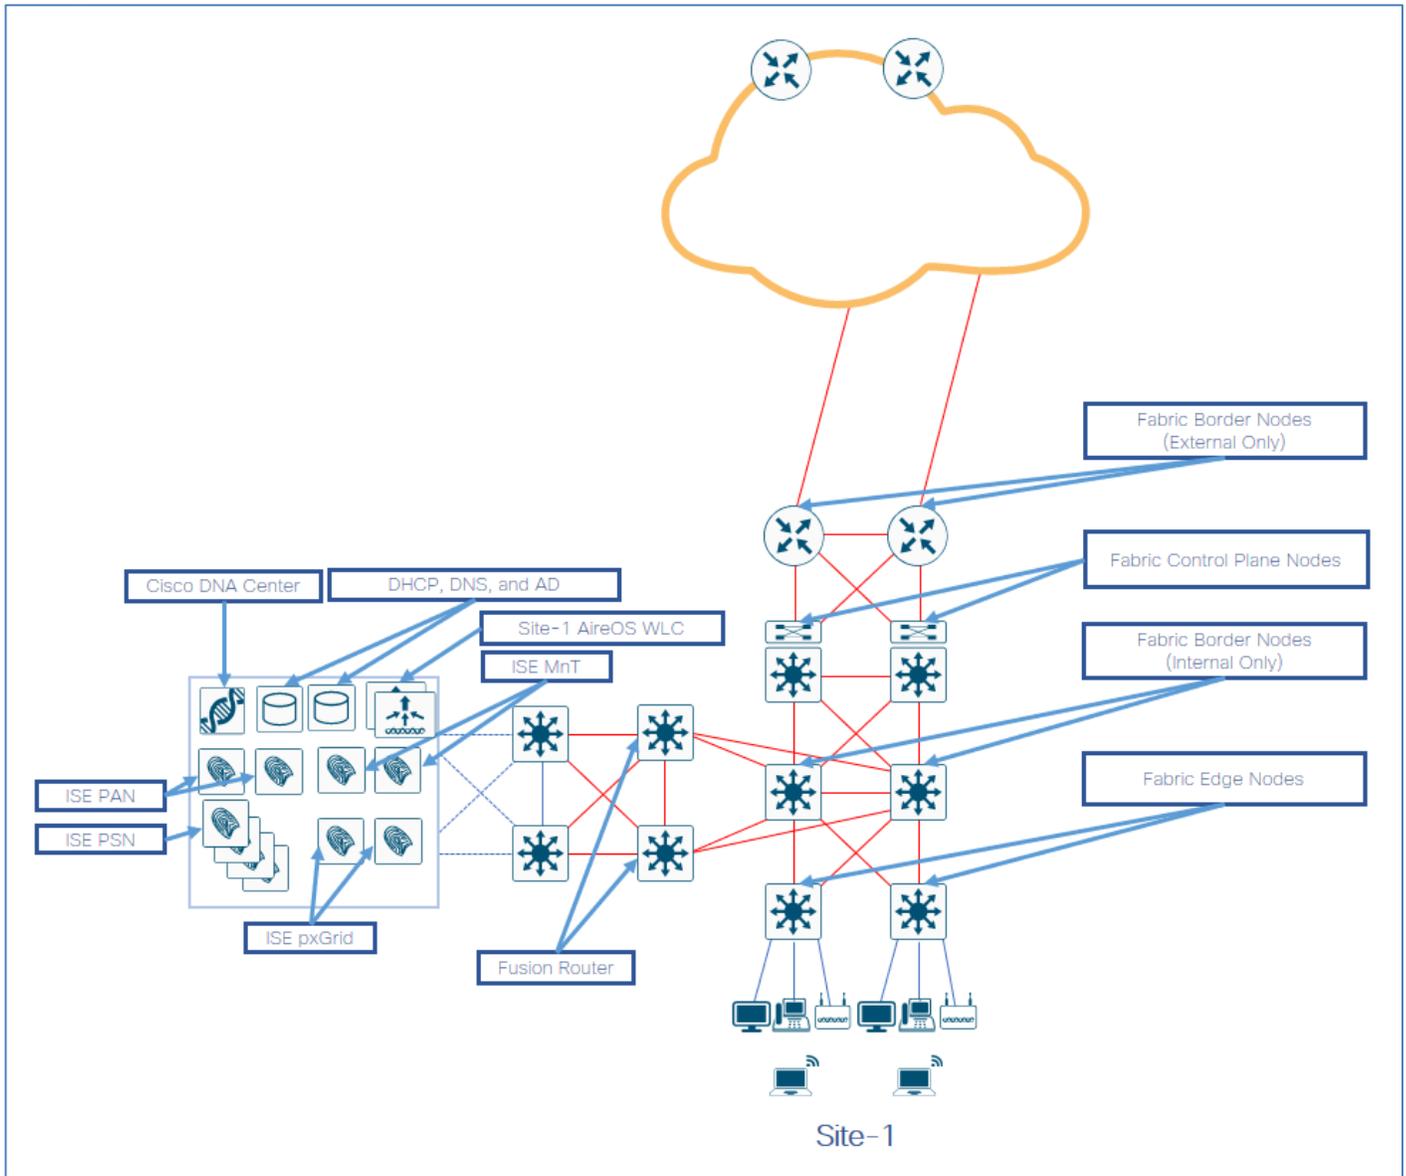
Figure 30 BGP Autonomous System Topology



Fabric Role Overview

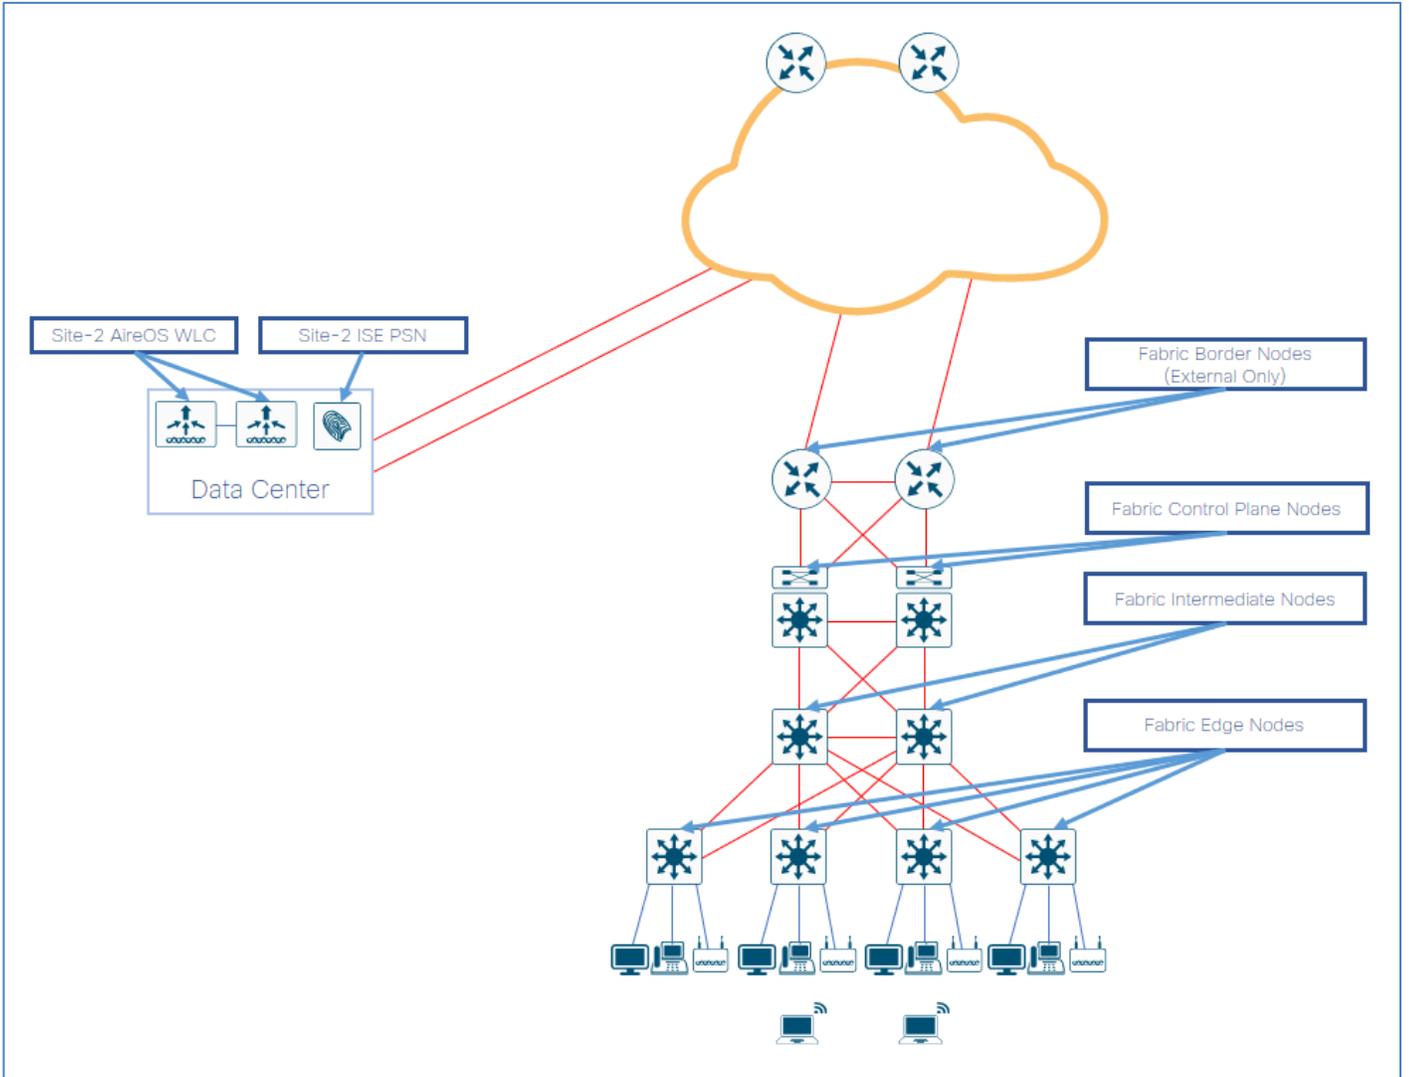
Site-1 Fabric Roles

Figure 31 Fabric Roles – Site-1



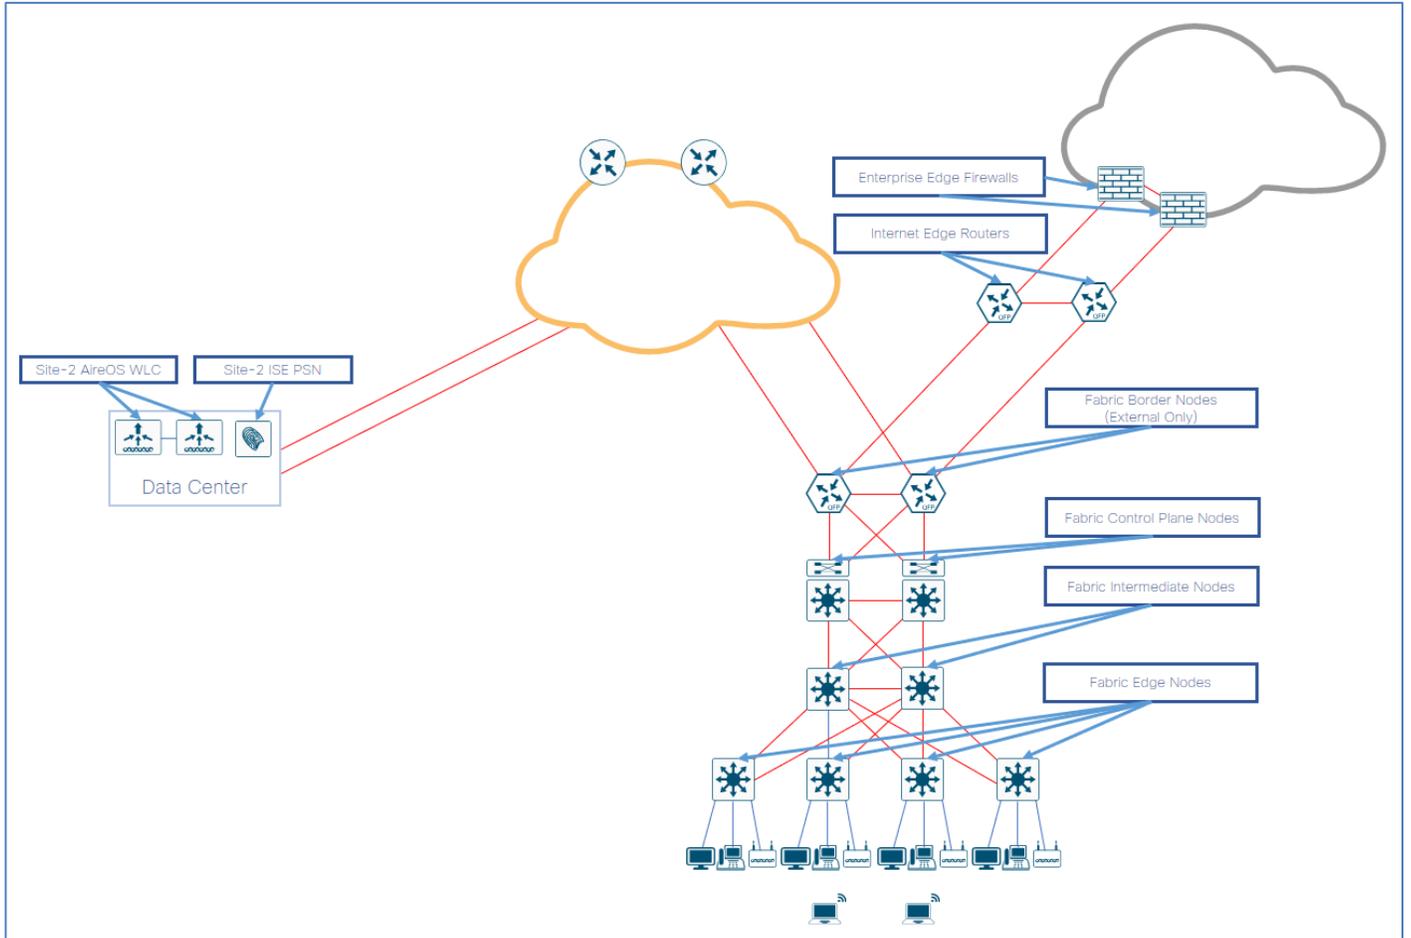
Site-2 Fabric Roles

Figure 32 Fabric Roles – Site-2



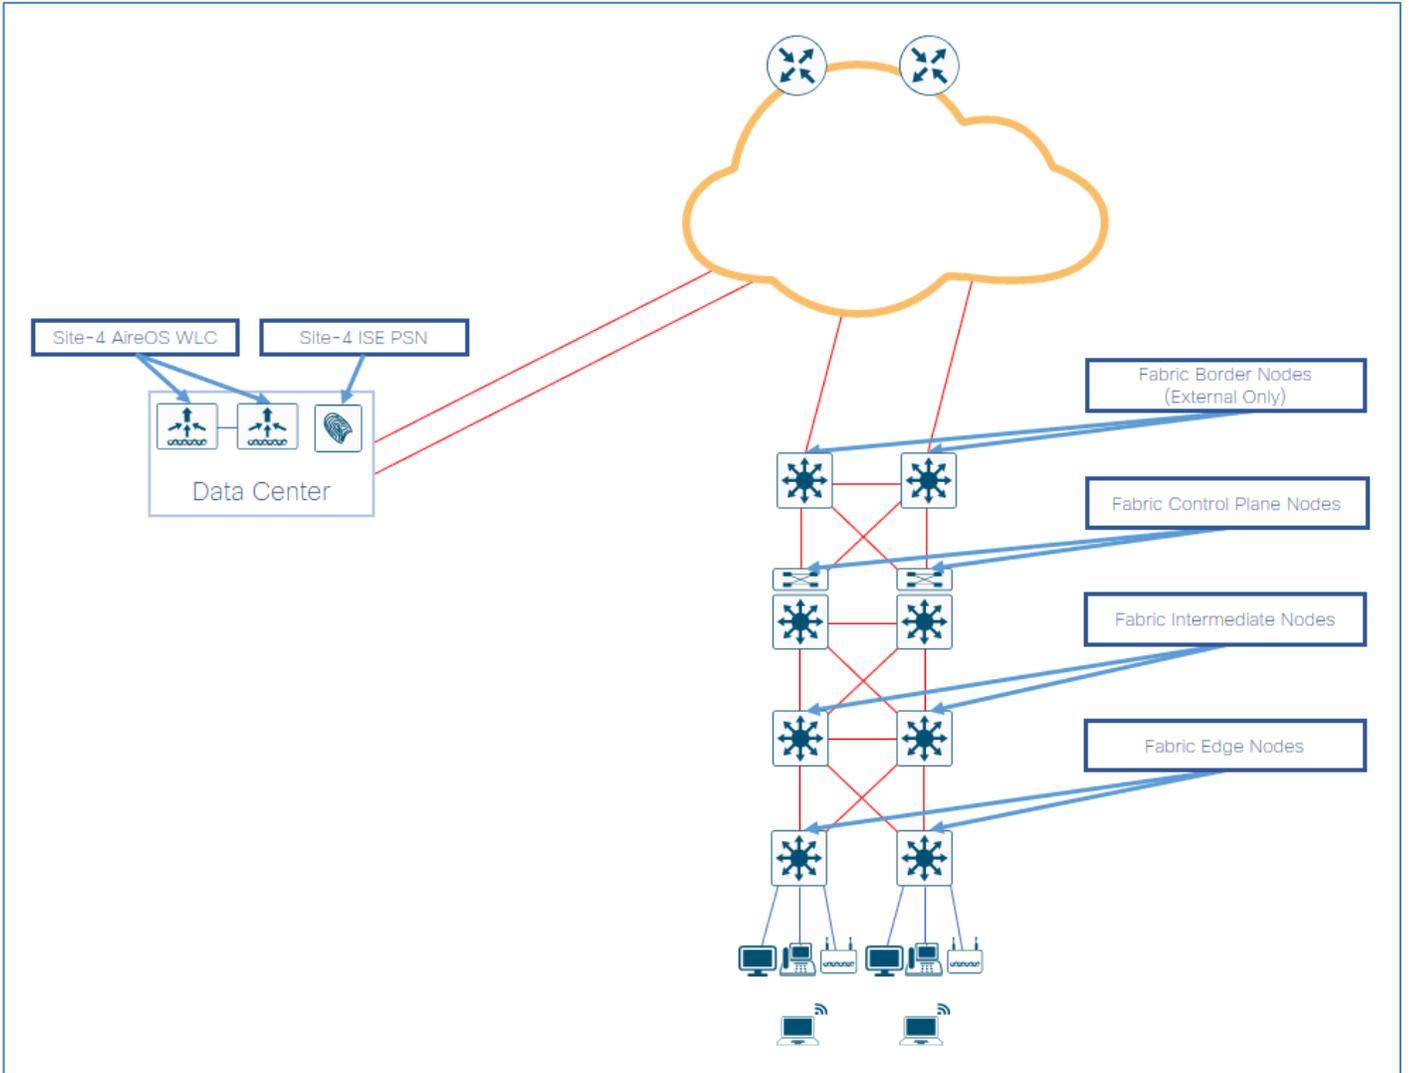
Site-3 Fabric Roles

Figure 33 Fabric Roles – Site-3



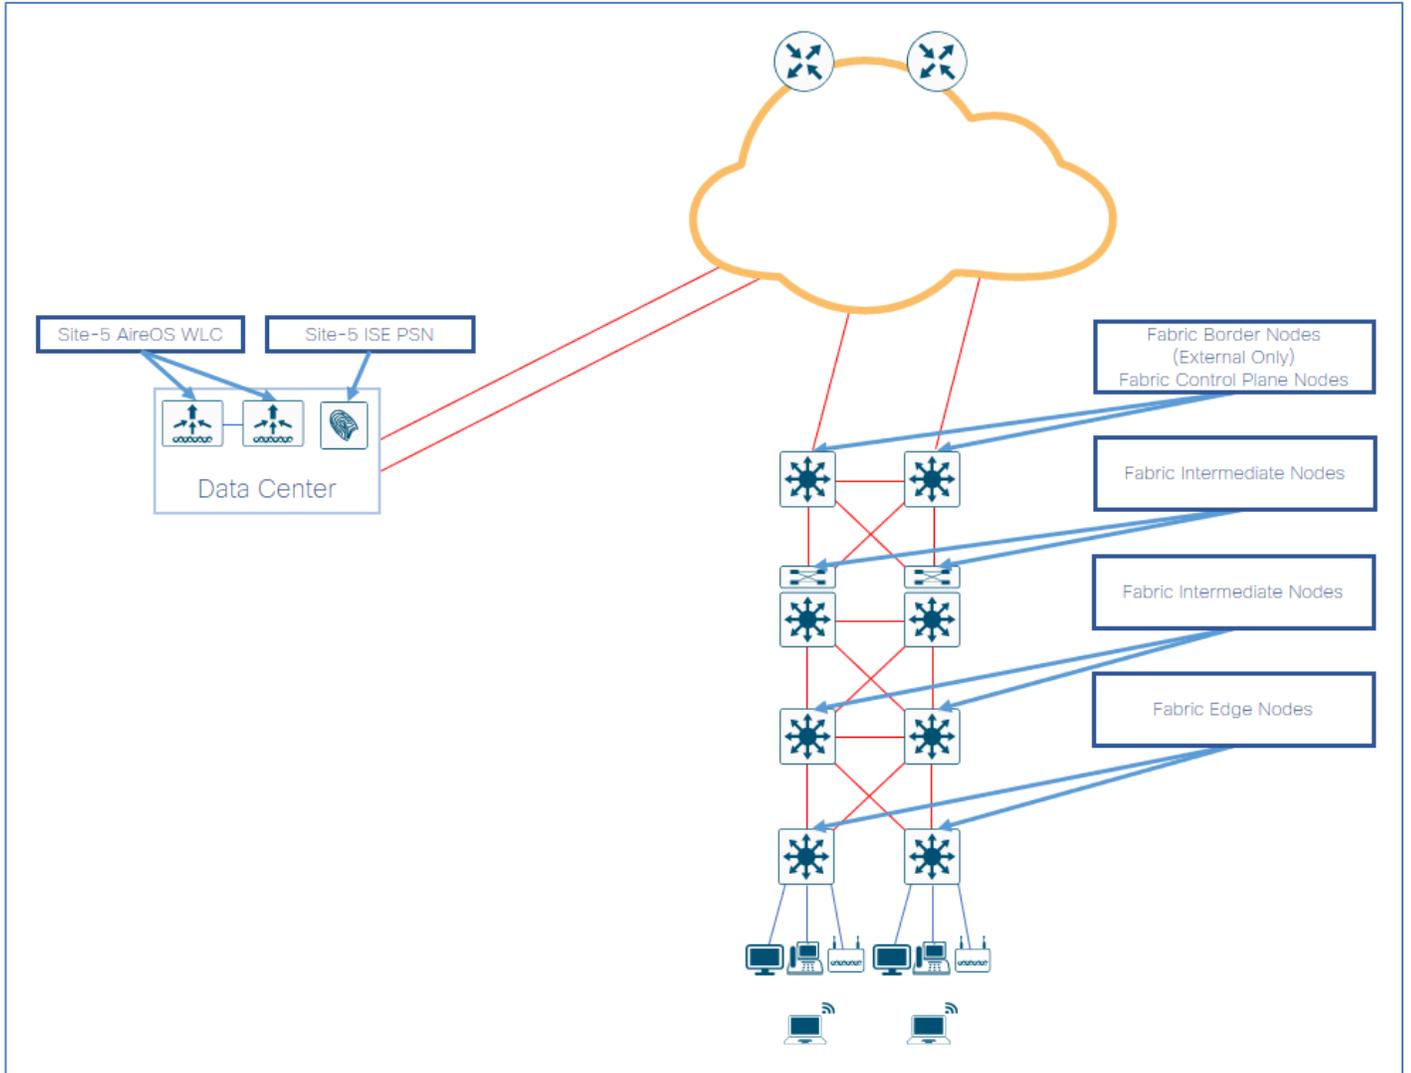
Site-4 Fabric Roles

Figure 34 Fabric Roles – Site-4



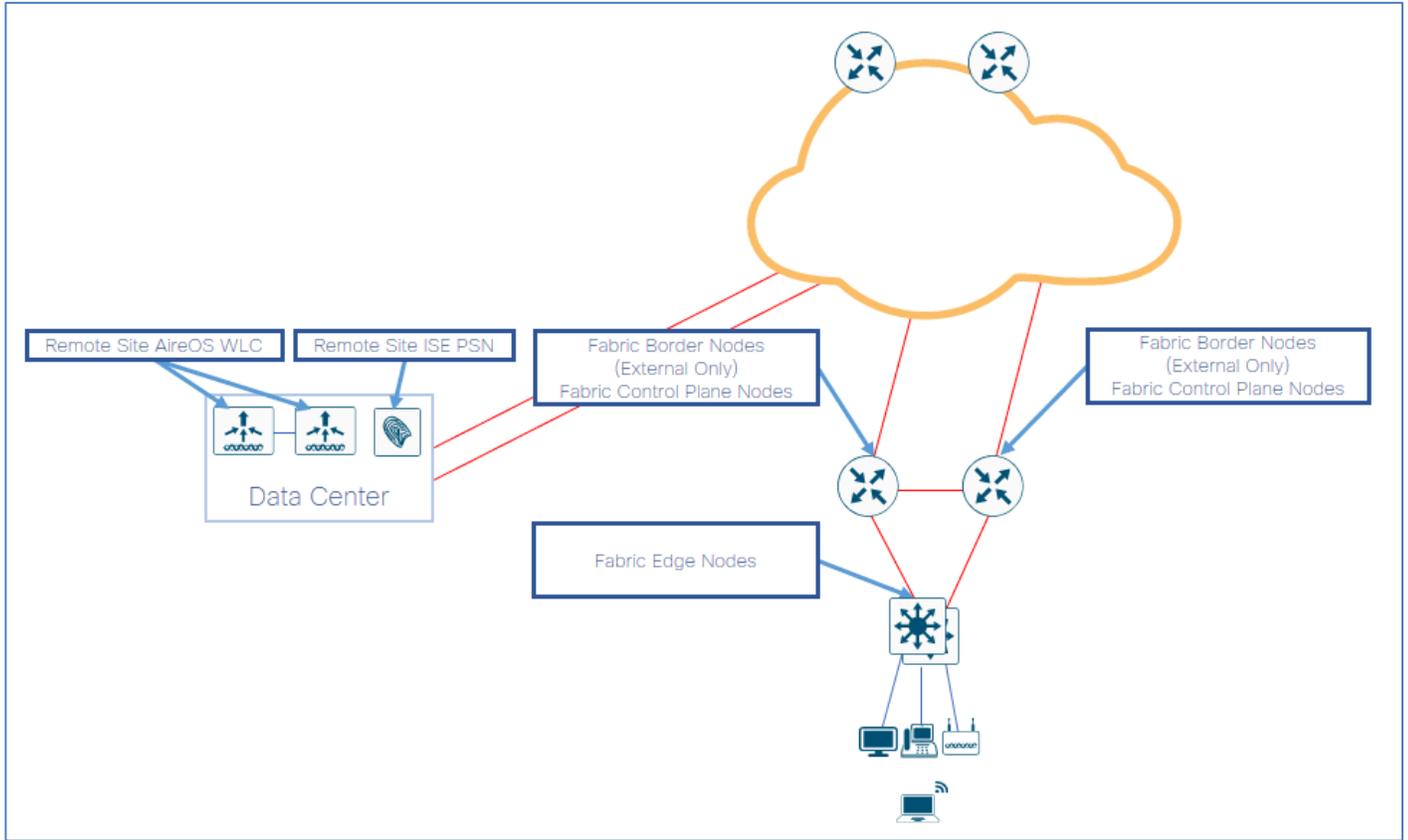
Site-5 Fabric Roles

Figure 35 Fabric Roles – Site-5



Branch Fabric Roles

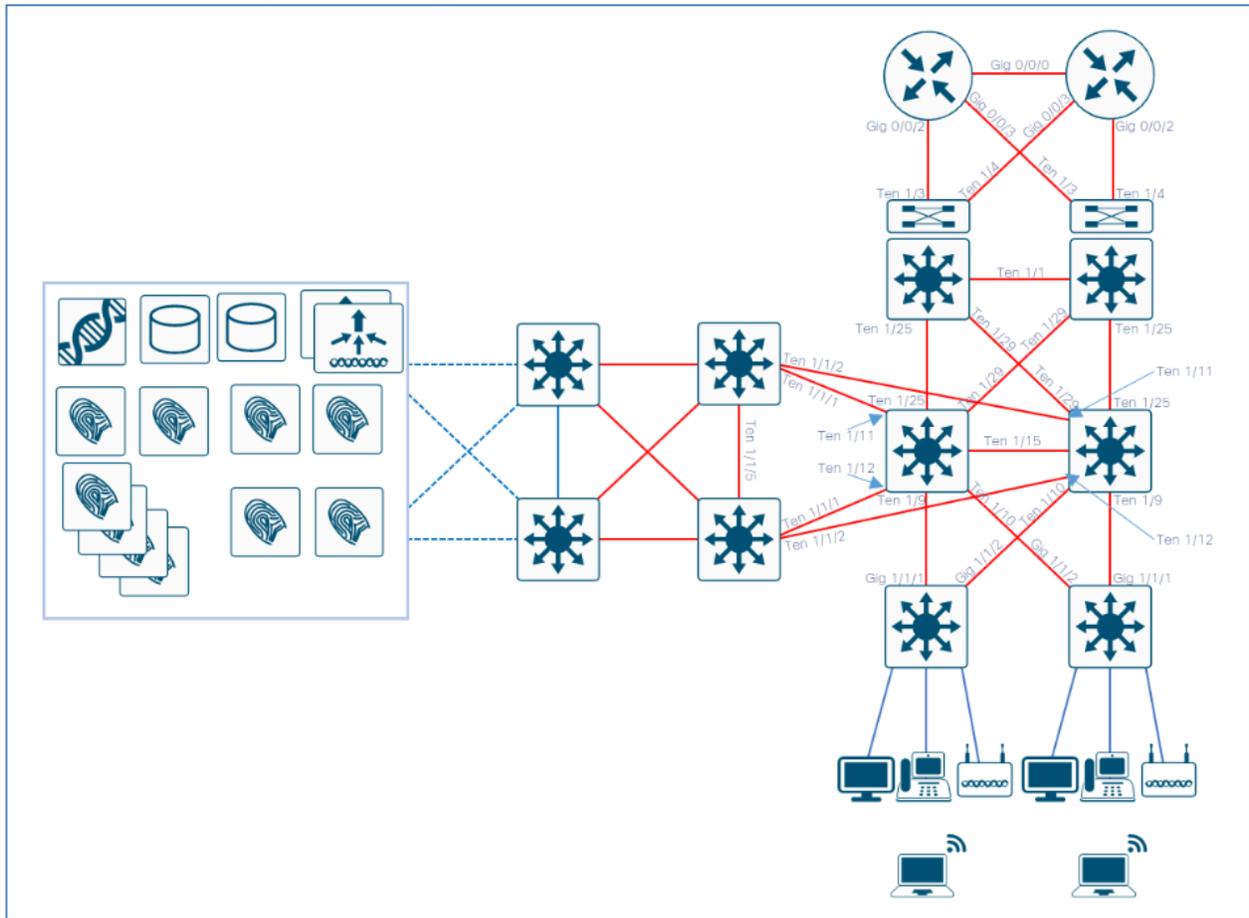
Figure 36 Fabric Role – Branch Sites



Layer-2 Overview

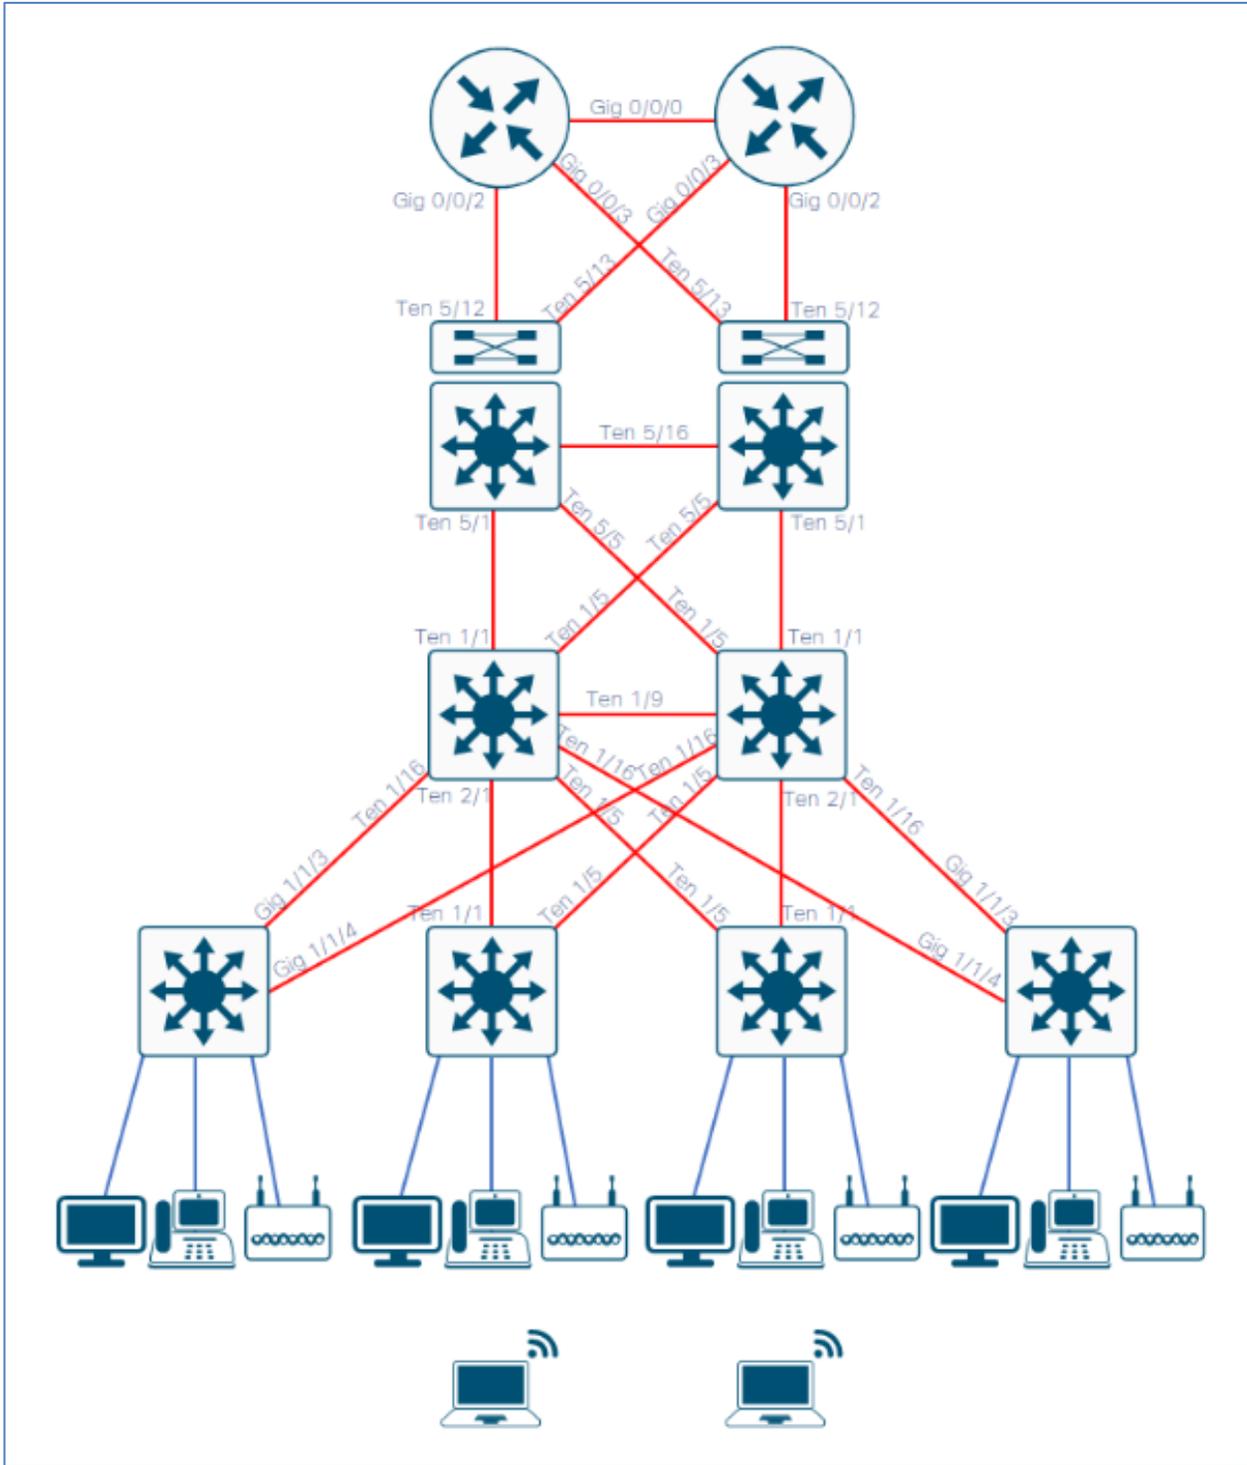
Site-1 Layer-2

Figure 37 Site-1 Layer-2 Topology



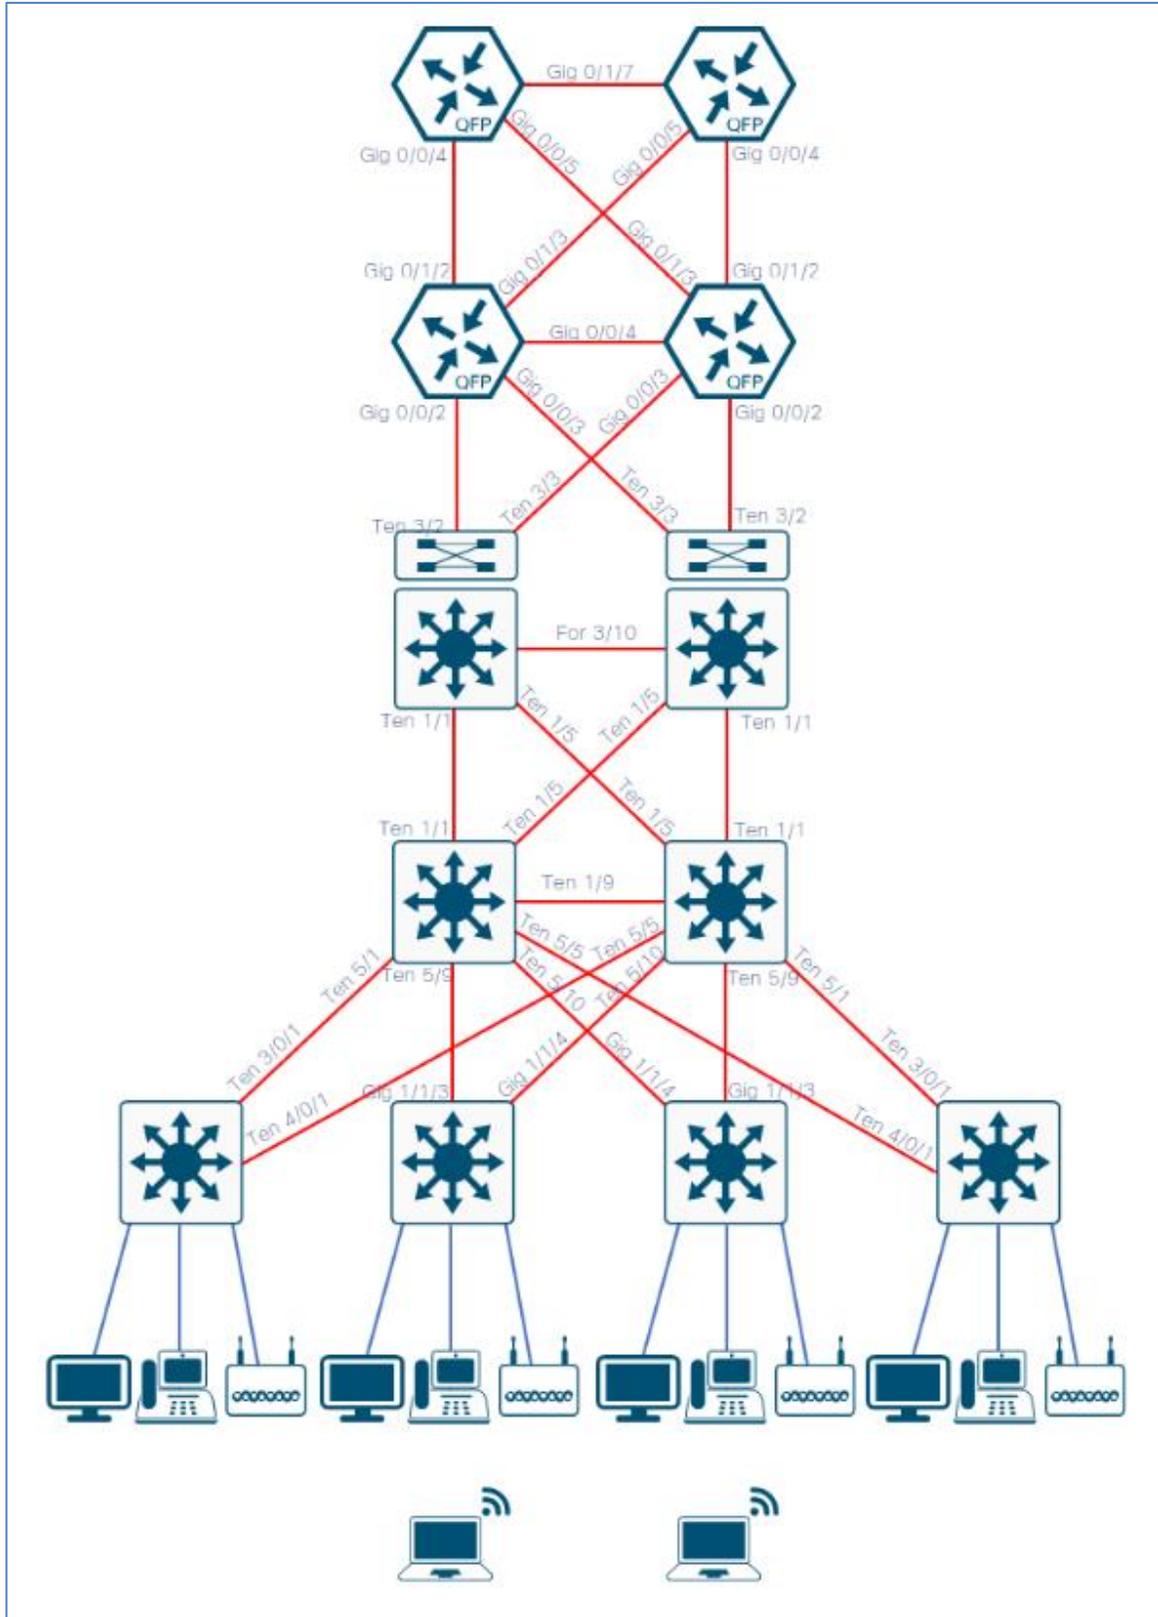
Site-2 Layer-2

Figure 38 Site-2 Layer-2 Topology



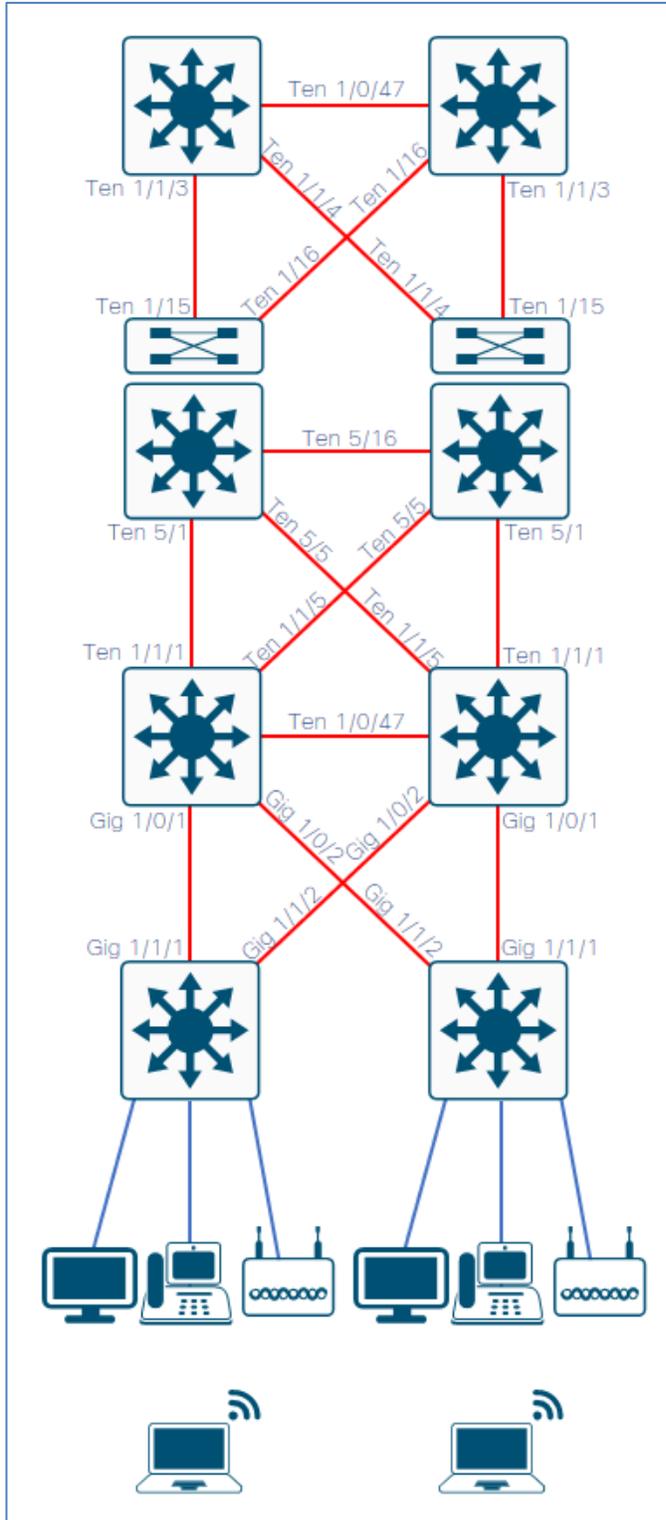
Site-3 Layer-2

Figure 39 Site-3 Layer-2 Topology



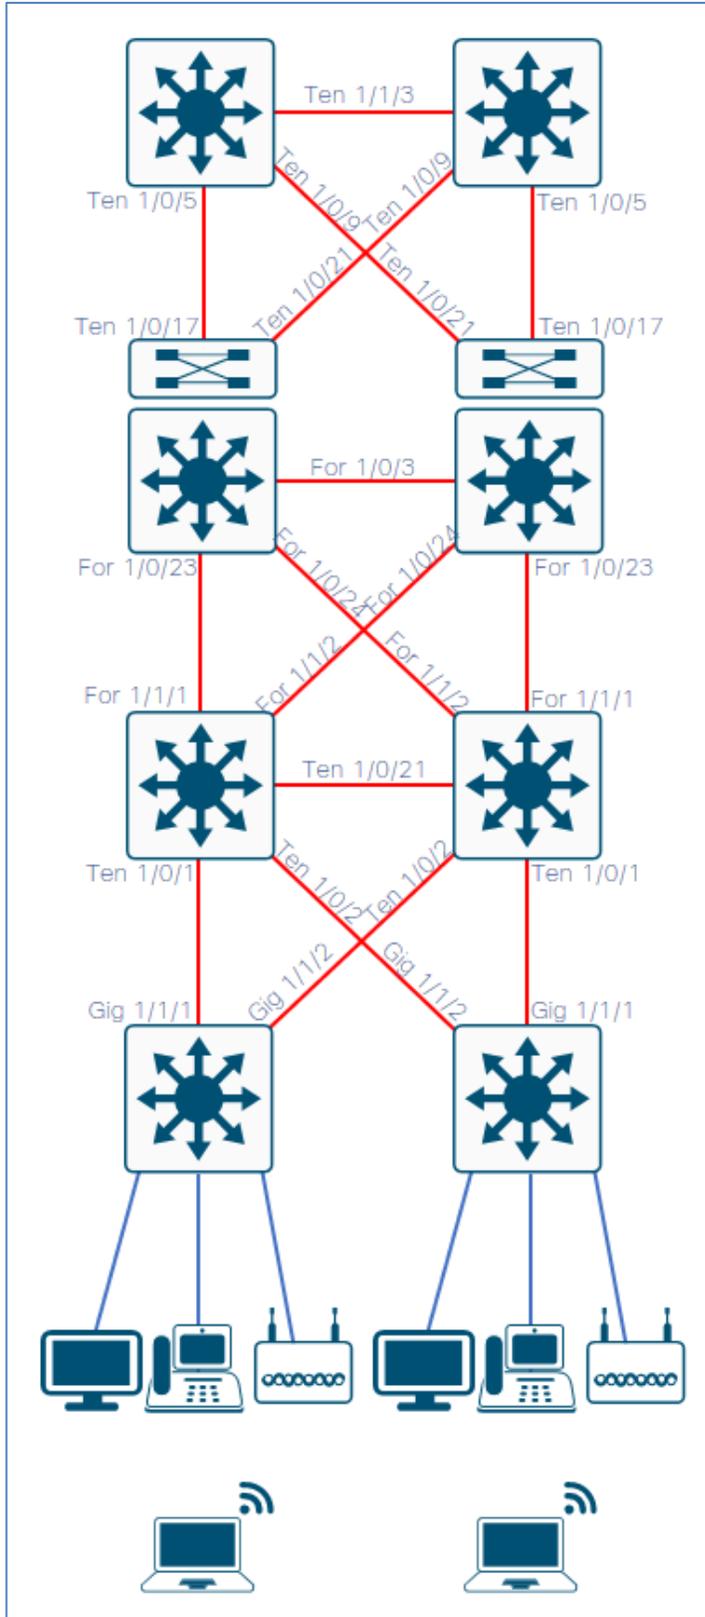
Site-4 Layer-2

Figure 40 Site-4 Layer-2 Topology



Site-5 Layer-2

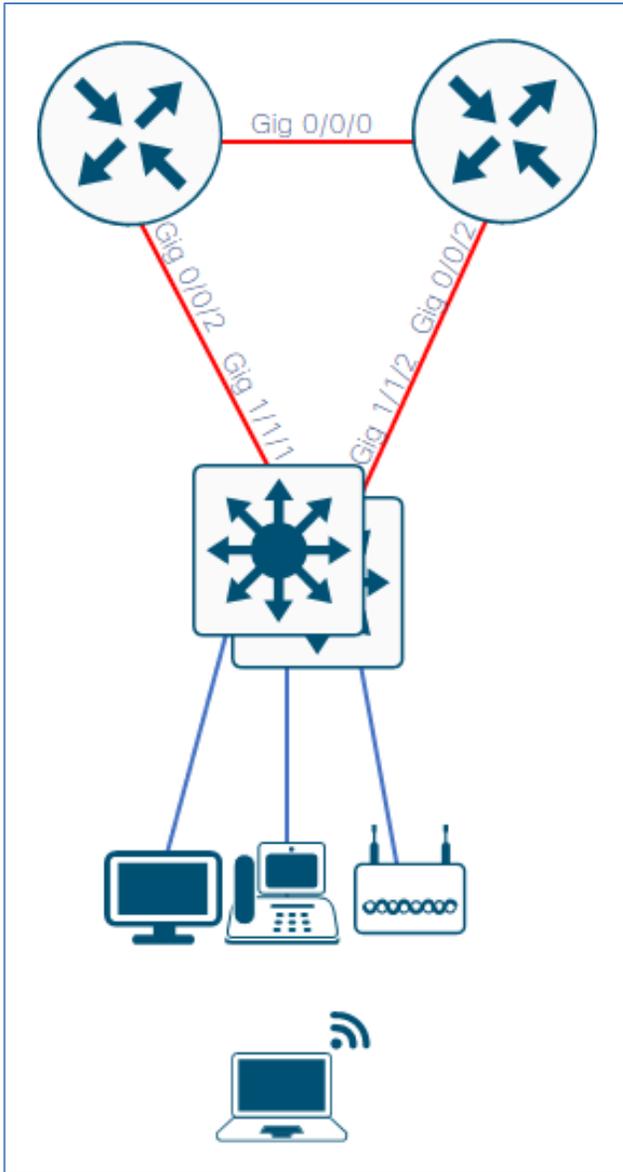
Figure 41 Site-5 Layer-2 Topology



Branch Layer-2

The remaining branch sites have the same general matching topology of a pair of branch edge routers connected to a switch stack downstream.

Figure 42 Branch Sites Layer-2 General Topology



Layer-3 Overview

Individual Sites Loopback IP Schema

Site-1 Layer-3

Figure 43 Site-1 Layer-3 Topology

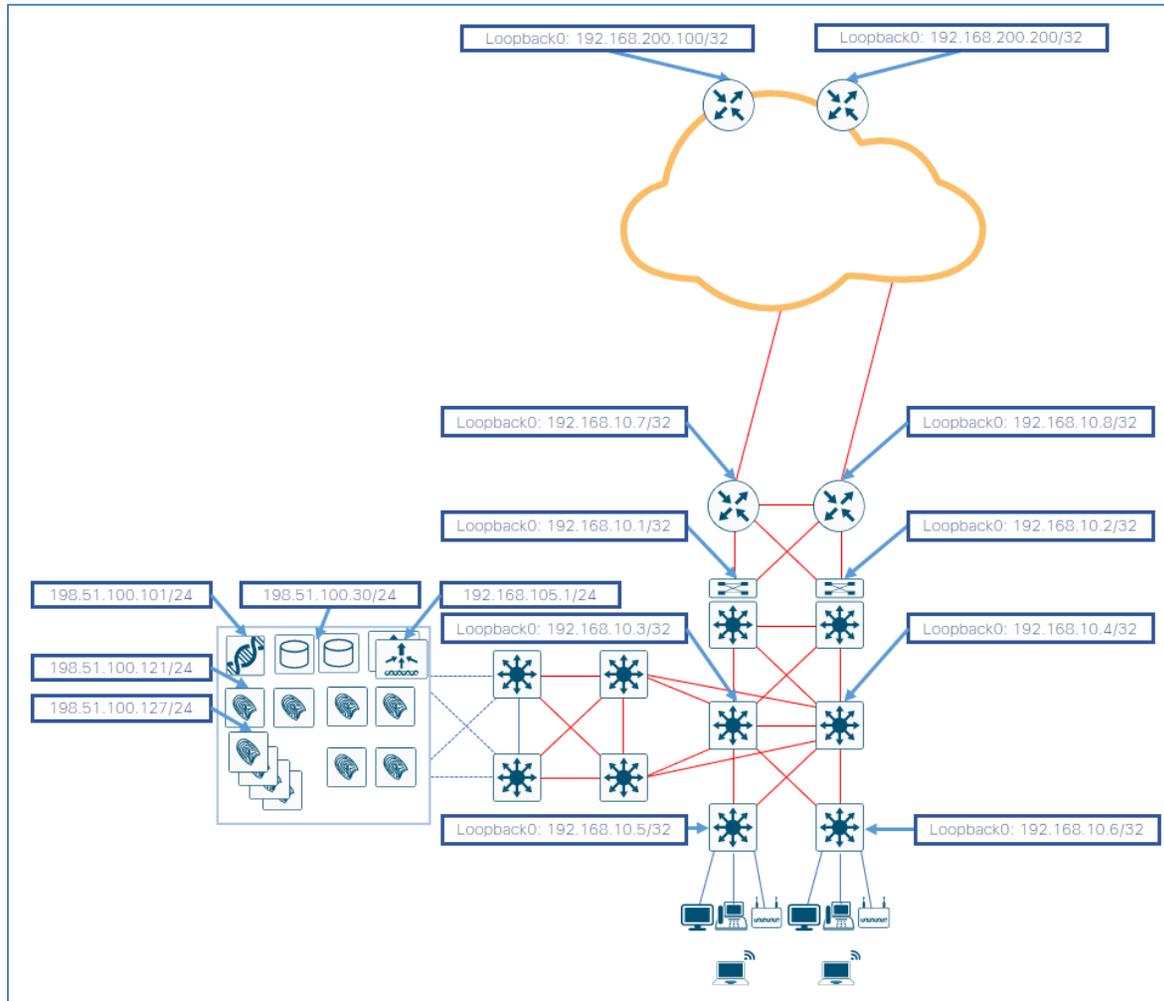


Figure 44 Site-2 Layer-3 Topology

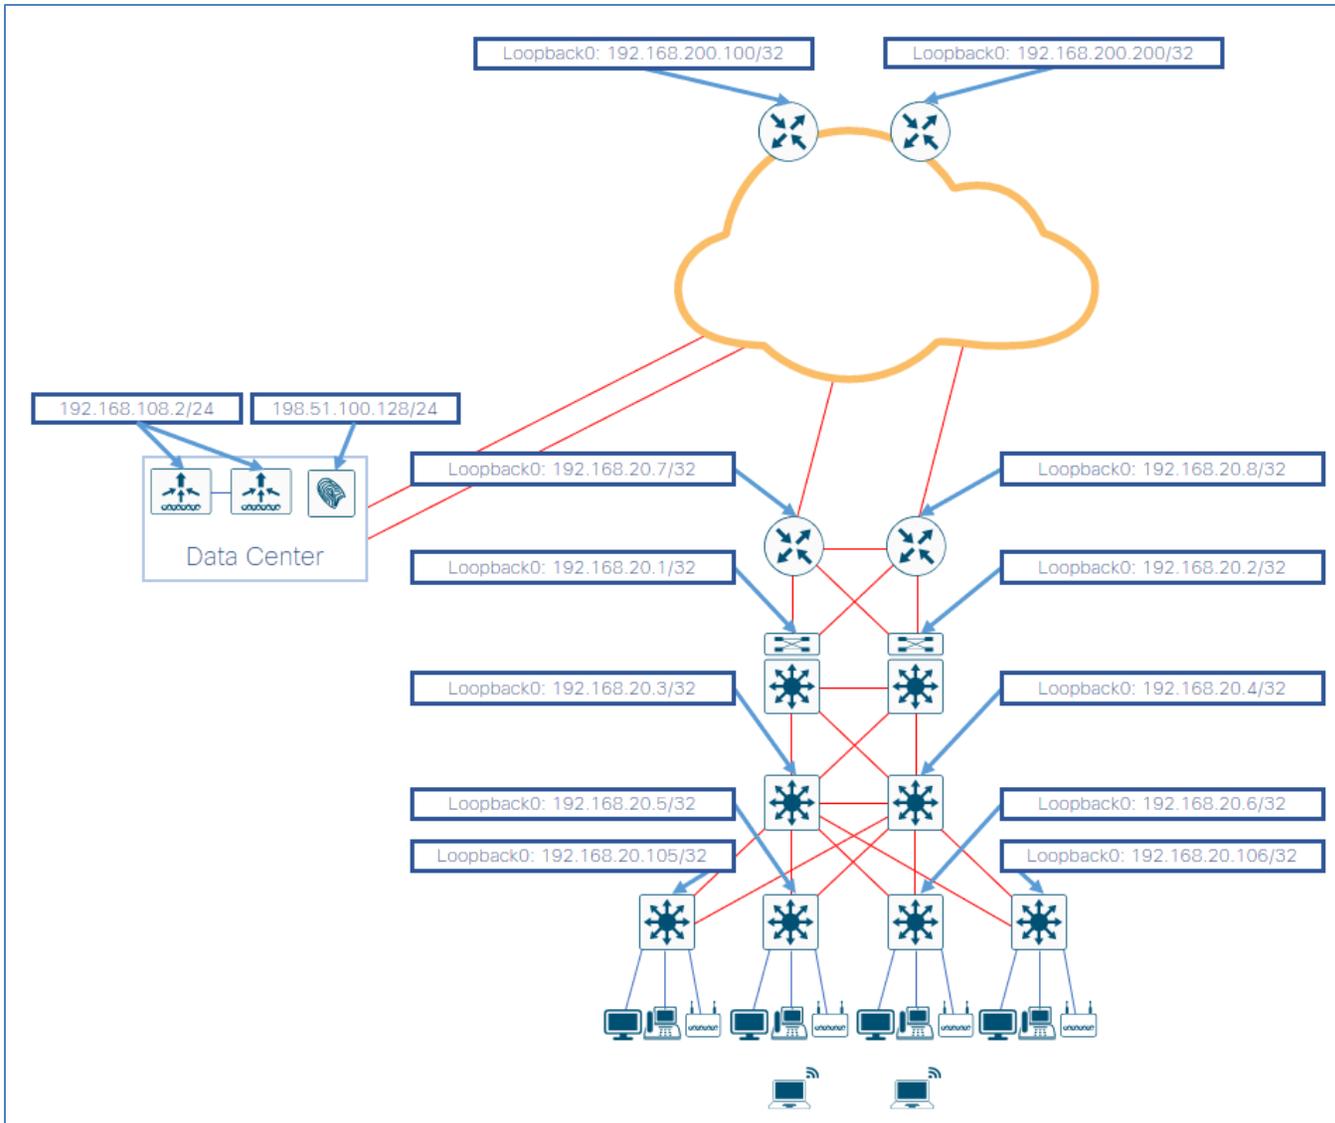


Figure 45 Site-3 Layer-3 Topology

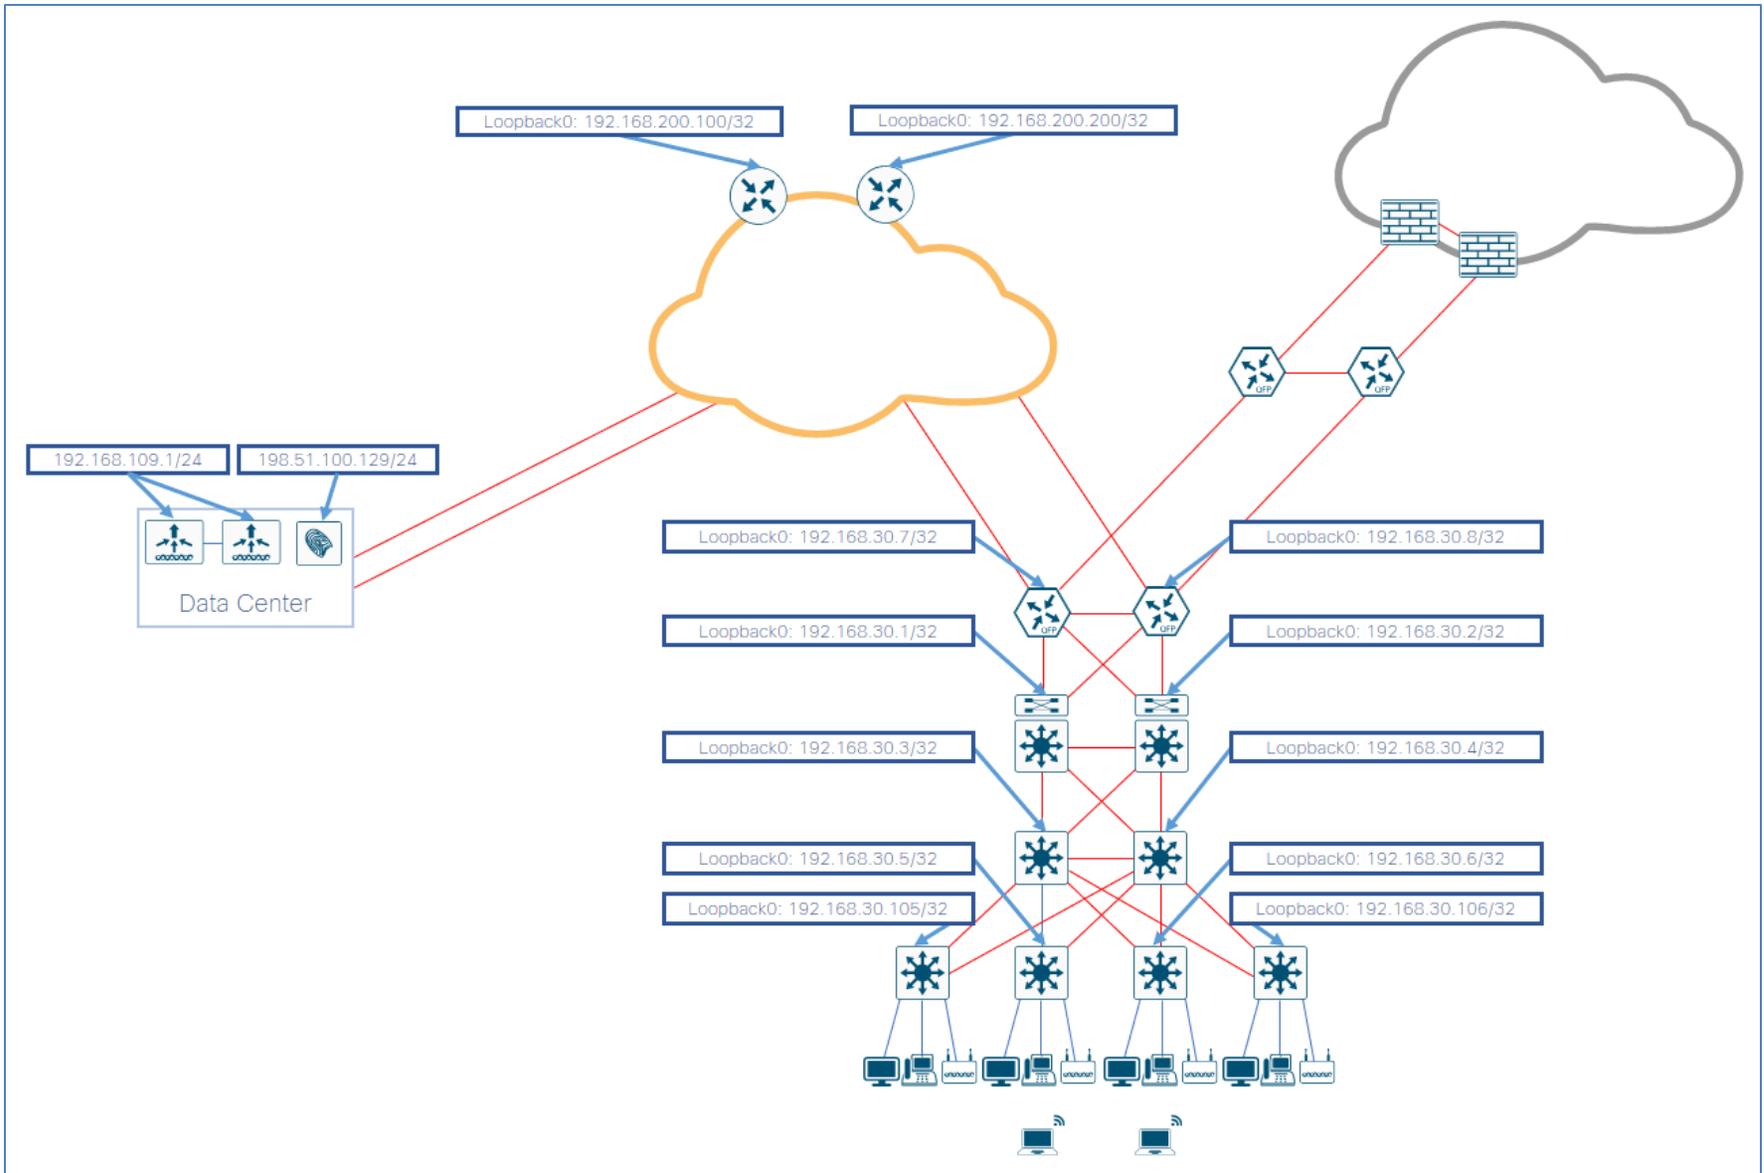


Figure 46 Site-4 Layer-3 Topology

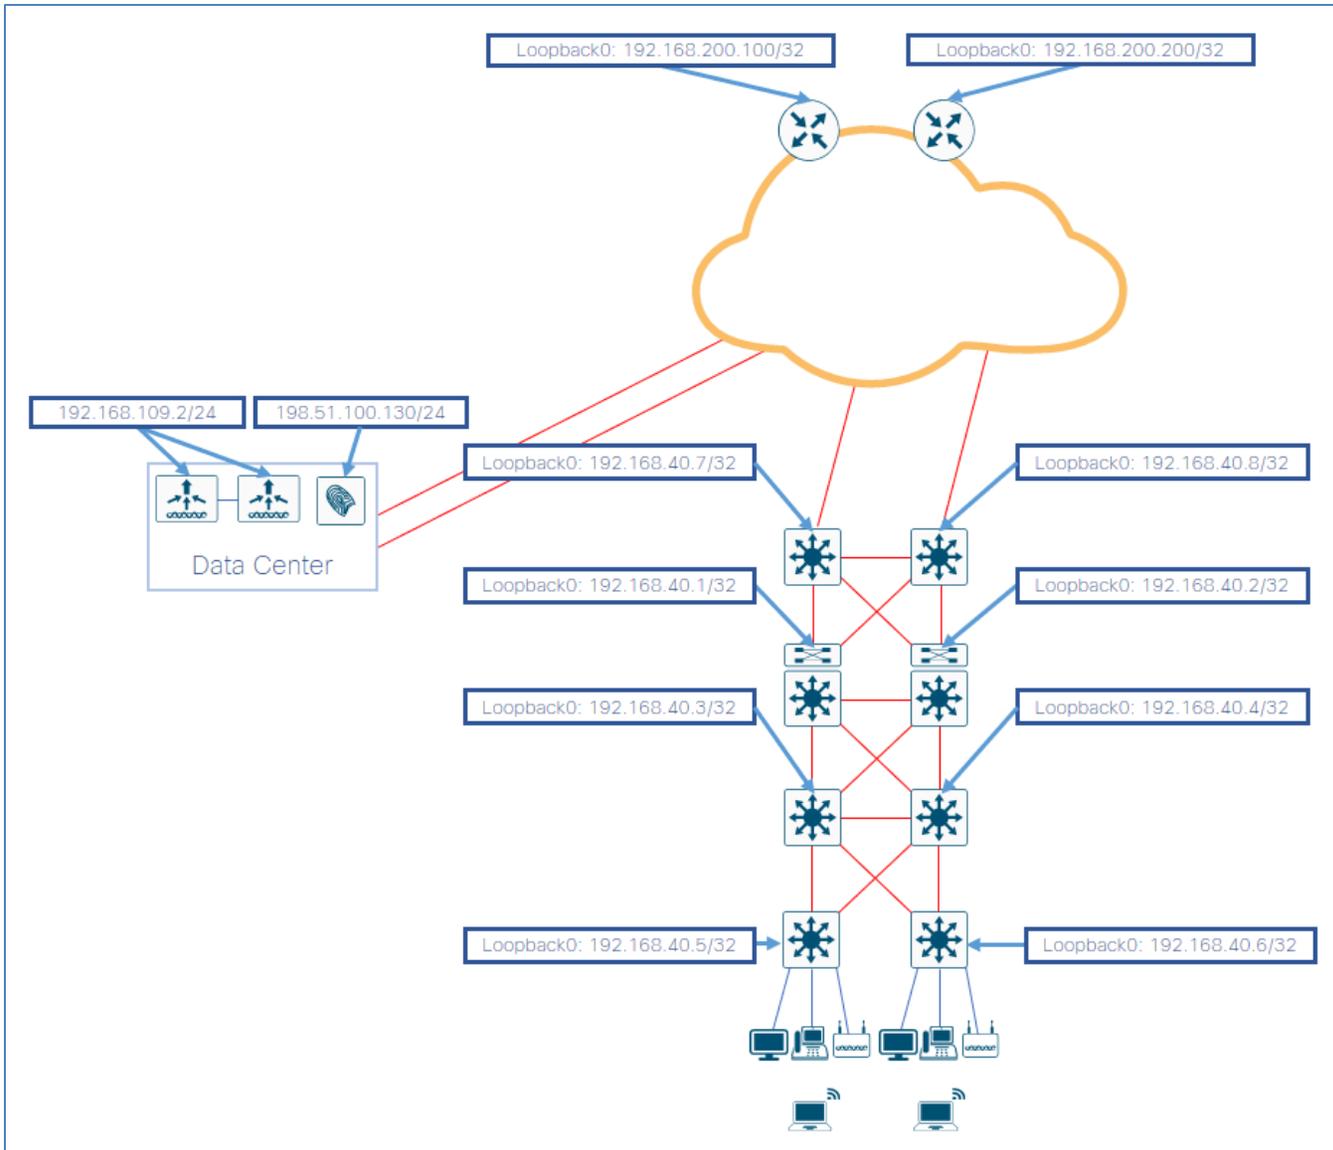
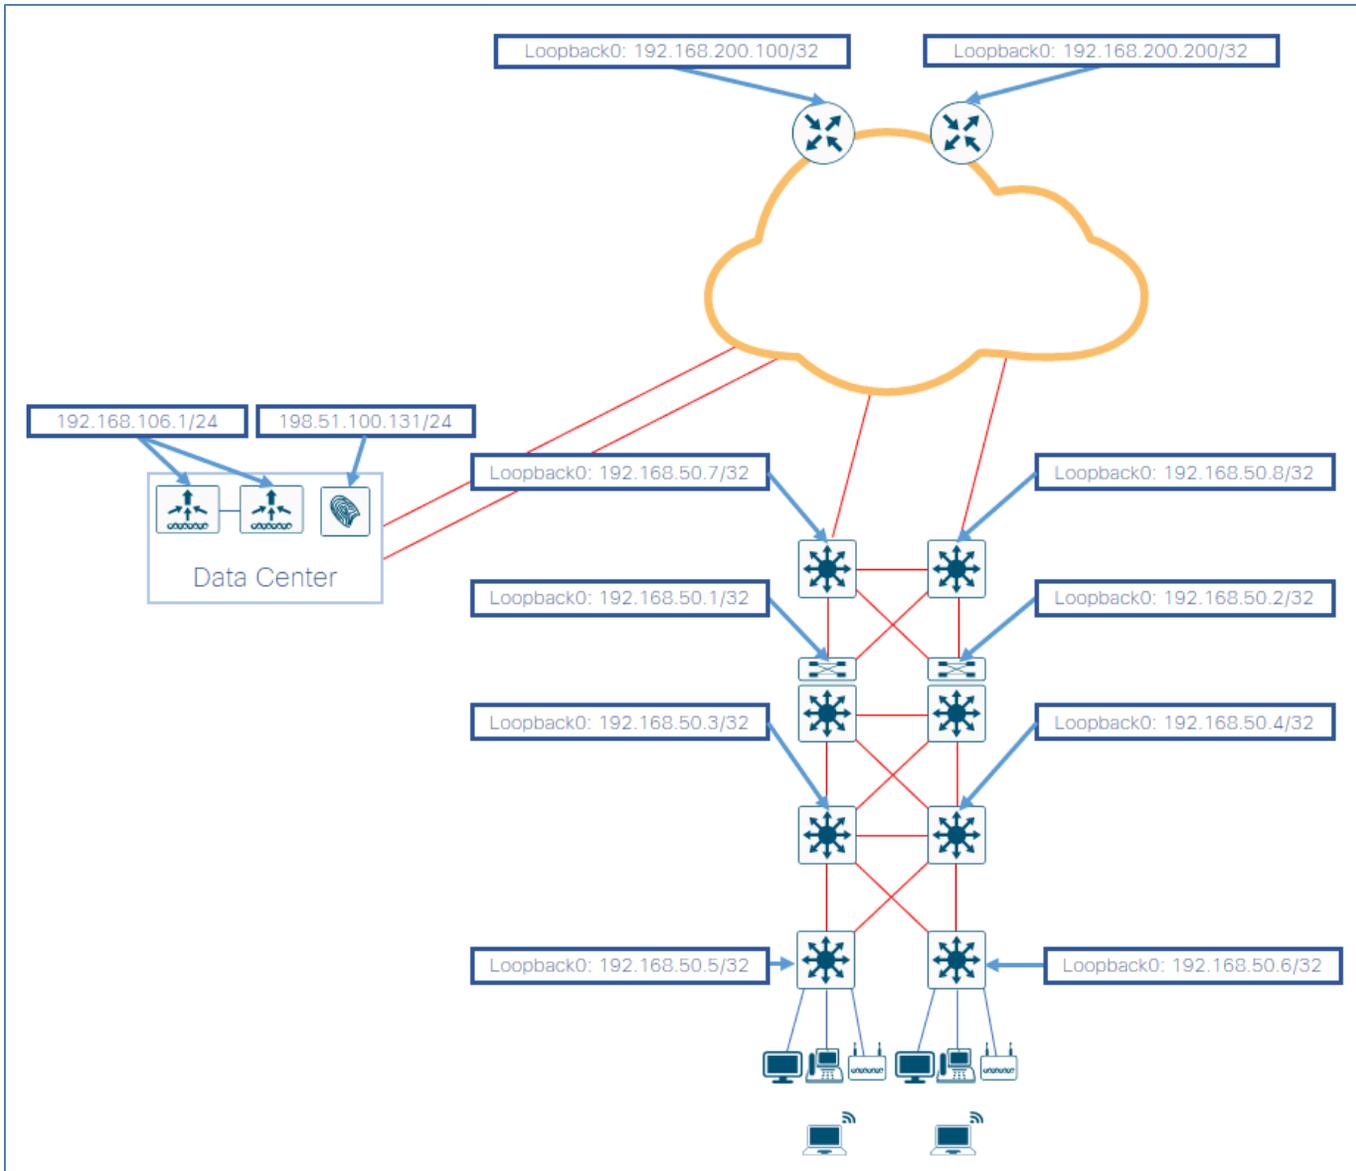
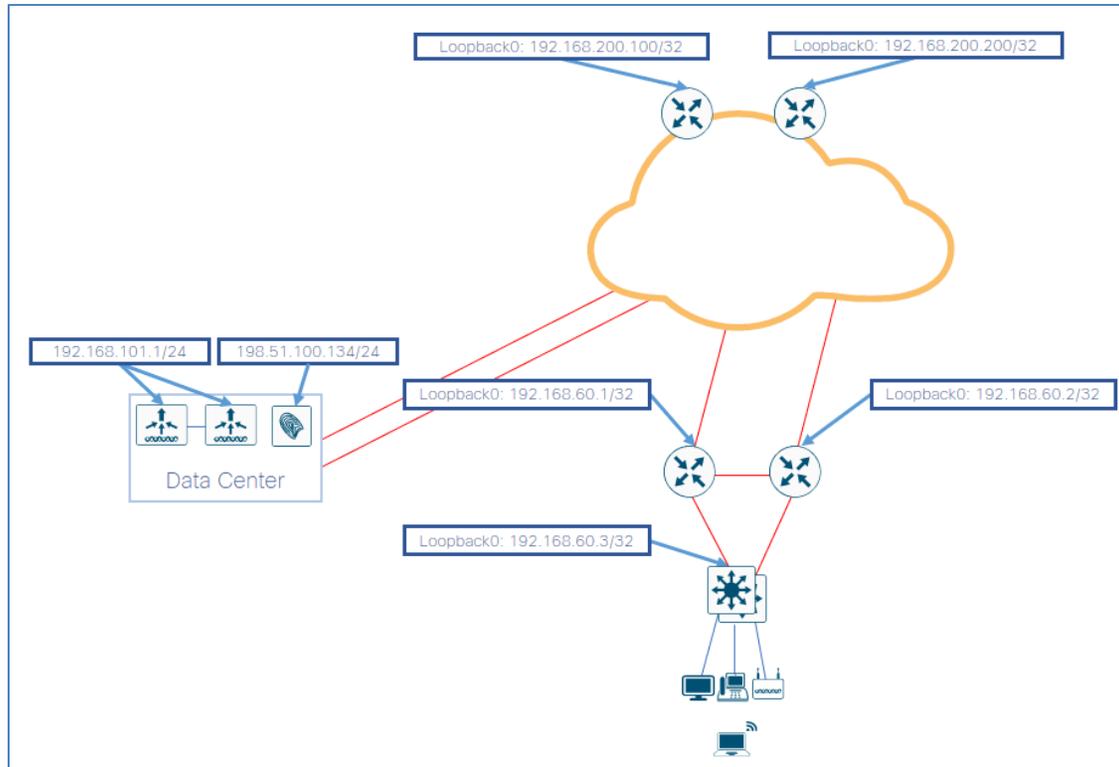


Figure 47 Site-5 Layer-3 Topology



Branch-1 Layer-3

Figure 48 Branch-1 Layer 3 Topology



The remaining branch sites have the same general IP address schema with the branch edge routers consuming the .1 and .2 addresses and the switch stack consuming the .3 address for their respective Loopback 0 interfaces. Each site uses a dedicated PSN which is the next subsequent IP address.

Table 5 Branch Site Loopback, ISE, and WLC IP Schema

Location	Subnet for Loopbacks	ISE PSN IP Address	WLC IP Address
Branch-1	192.168.60.x/32	198.51.100.134	192.168.101.1
Branch-2	192.168.70.x/32	198.51.100.135	192.168.101.2
Branch-3	192.168.80.x/32	198.51.100.136	192.168.102.1
Branch-4	192.168.90.x/32	198.51.100.137	192.168.102.2
Branch-5	192.168.240.x/32	198.51.100.138	192.168.105.1

IP Address Pools

Table 6 IP Address Pools – Site-1 through Site-5

Location	Usage at Site	Pool Name	Network/Mask	IP Gateway	DHCP Server	DNS Server
Global	Global Pool	ONE-SEVEN-TWO	172.16.0.0/12	172.16.0.1	198.51.100.30	198.51.100.30
Site-1	Access Points	AP-SITE-01	172.16.110.0/24	172.16.110.1	198.51.100.30	198.51.100.30
Site-1	BGP Handoff	BGP-SITE-01	172.16.111.0/24	172.16.111.1	-	-
Site-1	CAMPUS VN	CORP-SITE-01	172.16.112.0/24	172.16.112.1	198.51.100.30	198.51.100.30
Site-1	GUEST VN	GUEST-SITE-01	172.16.113.0/24	172.16.113.1	198.51.100.30	198.51.100.30
Site-1	RESEARCH VN	RESEARCH-SITE-01	172.16.114.0/24	172.16.114.1	198.51.100.30	198.51.100.30
Site-2	Access Points	AP-SITE-02	172.16.120.0/24	172.16.120.1	198.51.100.30	198.51.100.30
Site-2	-	BGP-SITE-02	172.16.121.0/24	172.16.121.1	-	-
Site-2	CAMPUS VN	CORP-SITE-02	172.16.122.0/24	172.16.122.1	198.51.100.30	198.51.100.30
Site-2	GUEST VN	GUEST-SITE-02	172.16.123.0/24	172.16.123.1	198.51.100.30	198.51.100.30
Site-2	RESEARCH VN	RESEARCH-SITE-02	172.16.124.0/24	172.16.124.1	198.51.100.30	198.51.100.30
Site-3	Access Points	AP-SITE-03	172.16.130.0/24	172.16.130.1	198.51.100.30	198.51.100.30
Site-3	BGP Handoff	BGP-SITE-03	172.16.131.0/24	172.16.131.1	-	-
Site-3	CAMPUS VN	CORP-SITE-03	172.16.132.0/24	172.16.132.1	198.51.100.30	198.51.100.30
Site-3	GUEST VN	GUEST-SITE-03	172.16.133.0/24	172.16.133.1	198.51.100.30	198.51.100.30
Site-3	RESEARCH VN	RESEARCH-SITE-03	172.16.134.0/24	172.16.134.1	198.51.100.30	198.51.100.30
Site-4	Access Points	AP-SITE-04	172.16.140.0/24	172.16.140.1	198.51.100.30	198.51.100.30
Site-4	-	BGP-SITE-04	172.16.141.0/24	172.16.141.1	-	-
Site-4	CAMPUS VN	CORP-SITE-04	172.16.142.0/24	172.16.142.1	198.51.100.30	198.51.100.30
Site-4	GUEST VN	GUEST-SITE-04	172.16.143.0/24	172.16.143.1	198.51.100.30	198.51.100.30
Site-4	RESEARCH VN	RESEARCH-SITE-04	172.16.144.0/24	172.16.144.1	198.51.100.30	198.51.100.30
Site-5	Access Points	AP-SITE-05	172.16.150.0/24	172.16.150.1	198.51.100.30	198.51.100.30
Site-5	-	BGP-SITE-05	172.16.151.0/24	172.16.151.1	-	-
Site-5	CAMPUS VN	CORP-SITE-05	172.16.152.0/24	172.16.152.1	198.51.100.30	198.51.100.30
Site-5	GUEST VN	GUEST-SITE-05	172.16.153.0/24	172.16.153.1	198.51.100.30	198.51.100.30
Site-5	RESEARCH VN	RESEARCH-SITE-05	172.16.154.0/24	172.16.154.1	198.51.100.30	198.51.100.30

Table 7 IP Address Pools – Branch Site-1 through Branch Site-5

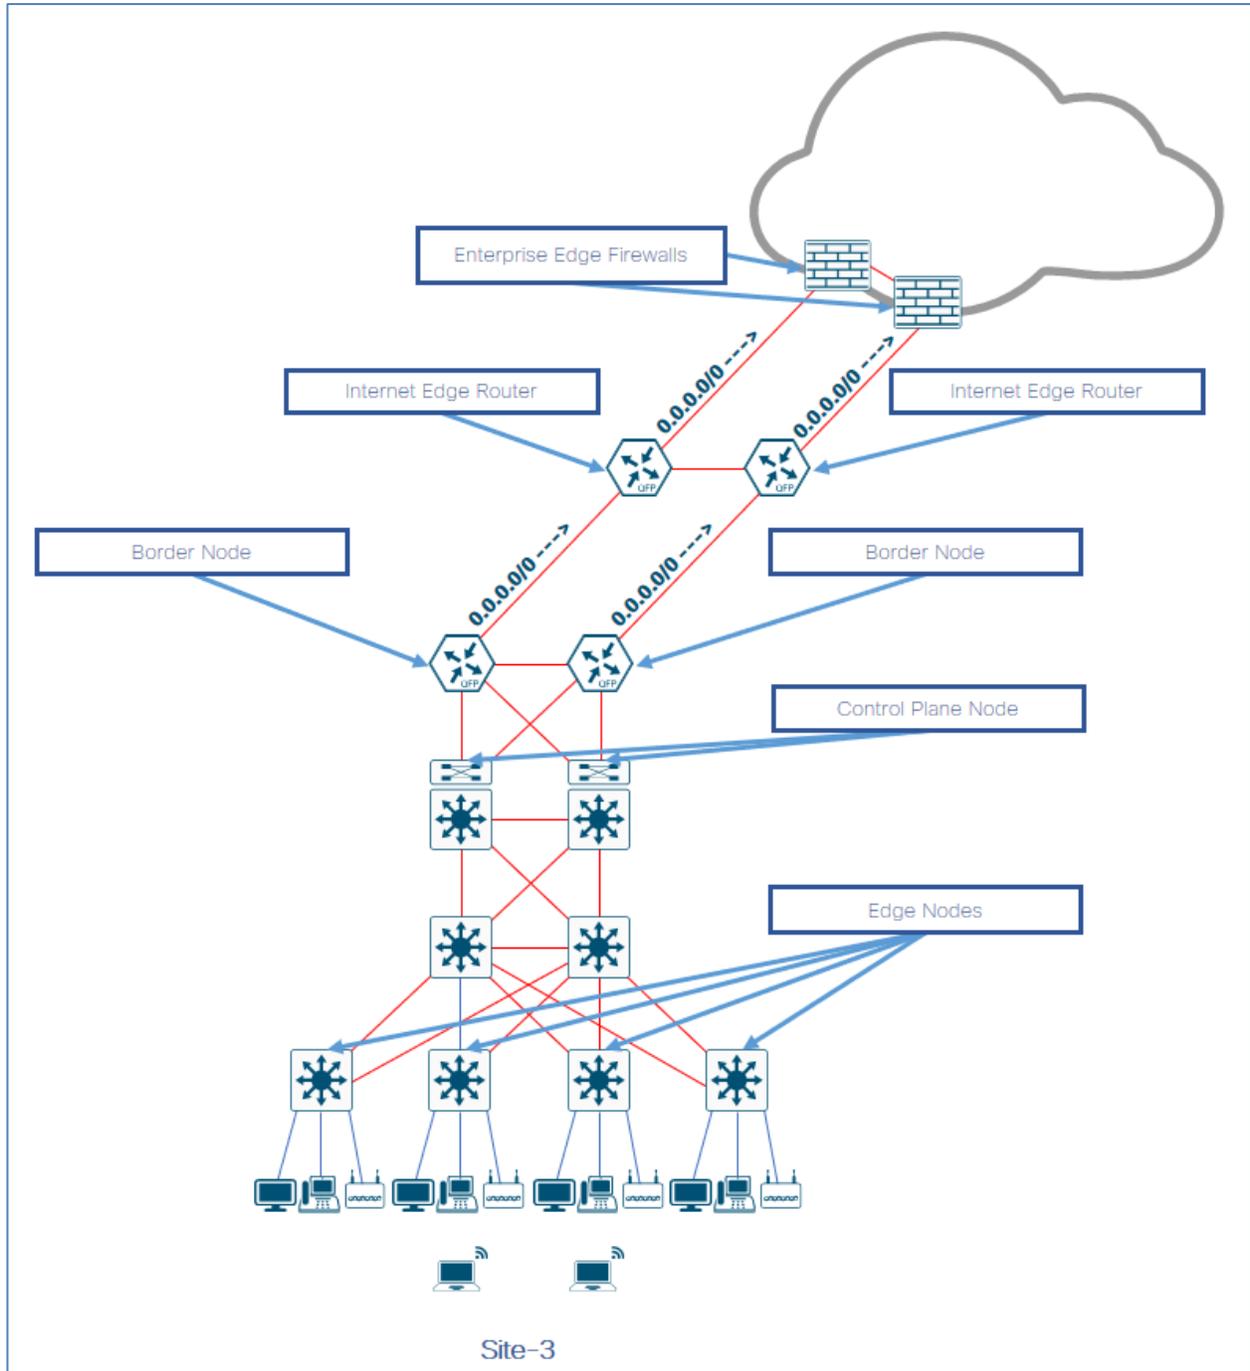
Location	Usage at Site	Pool Name	Network/Mask	IP Gateway	DHCP Server	DNS Server
Branch-1	Access Points	AP-BRANCH-01	172.16.160.0/24	172.16.160.1	198.51.100.30	198.51.100.30
Branch-1	-	BGP-BRANCH-01	172.16.161.0/24	172.16.161.1	-	-
Branch-1	CAMPUS VN	CORP-BRANCH-01	172.16.162.0/24	172.16.162.1	198.51.100.30	198.51.100.30
Branch-1	GUEST VN	GUEST-BRANCH-01	172.16.163.0/24	172.16.163.1	198.51.100.30	198.51.100.30
Branch-1	RESEARCH VN	RESEARCH-BRANCH-01	172.16.164.0/24	172.16.164.1	198.51.100.30	198.51.100.30
Branch-2	Access Points	AP-BRANCH-02	172.16.170.0/24	172.16.170.1	198.51.100.30	198.51.100.30
Branch-2	-	BGP-BRANCH-02	172.16.171.0/24	172.16.171.1	-	-
Branch-2	CAMPUS VN	CORP-BRANCH-02	172.16.172.0/24	172.16.172.1	198.51.100.30	198.51.100.30
Branch-2	GUEST VN	GUEST-BRANCH-02	172.16.173.0/24	172.16.173.1	198.51.100.30	198.51.100.30
Branch-2	RESEARCH VN	RESEARCH-BRANCH-02	172.16.174.0/24	172.16.174.1	198.51.100.30	198.51.100.30
Branch-3	Access Points	AP-BRANCH-03	172.16.180.0/24	172.16.180.1	198.51.100.30	198.51.100.30
Branch-3	-	BGP-BRANCH-03	172.16.181.0/24	172.16.181.1	-	-
Branch-3	CAMPUS VN	CORP-BRANCH-03	172.16.182.0/24	172.16.182.1	198.51.100.30	198.51.100.30
Branch-3	GUEST VN	GUEST-BRANCH-03	172.16.183.0/24	172.16.183.1	198.51.100.30	198.51.100.30
Branch-3	RESEARCH VN	RESEARCH-BRANCH-03	172.16.184.0/24	172.16.184.1	198.51.100.30	198.51.100.30
Branch-4	Access Points	AP-BRANCH-04	172.16.190.0/24	172.16.190.1	198.51.100.30	198.51.100.30
Branch-4	-	BGP-BRANCH-04	172.16.191.0/24	172.16.191.1	-	-
Branch-4	CAMPUS VN	CORP-BRANCH-04	172.16.192.0/24	172.16.192.1	198.51.100.30	198.51.100.30
Branch-4	GUEST VN	GUEST-BRANCH-04	172.16.193.0/24	172.16.193.1	198.51.100.30	198.51.100.30
Branch-4	RESEARCH VN	RESEARCH-BRANCH-04	172.16.194.0/24	172.16.194.1	198.51.100.30	198.51.100.30
Branch-5	Access Points	AP-BRANCH-05	172.16.240.0/24	172.16.240.1	198.51.100.30	198.51.100.30
Branch-5	-	BGP-BRANCH-05	172.16.241.0/24	172.16.241.1	-	-
Branch-5	CAMPUS VN	CORP-BRANCH-05	172.16.242.0/24	172.16.242.1	198.51.100.30	198.51.100.30
Branch-5	GUEST VN	GUEST-BRANCH-05	172.16.243.0/24	172.16.243.1	198.51.100.30	198.51.100.30
Branch-5	RESEARCH VN	RESEARCH-BRANCH-05	172.16.244.0/24	172.16.243.1	198.51.100.30	198.51.100.30

Shared Services and Internet Access

Internet Access

Default Route Overview

Figure 49 Site-3 Default Route



Shared Services

VRF Leaking Overview

Figure 50 VRF Leaking



Appendix D: Route Leaking

In [Operate Section > Process 1](#), the shared services were placed in a dedicated VRF. This is both a convention used for this prescriptive deployment guide and a design decision commonly seen in production networks. However, shared services are often designed and deployed in the Global Routing Table as well.

With shared services in a dedicated VRF, route leaking was very straightforward administratively as it used route-targets, although this was at the expense of creating another VRF to manage. The alternative approach, shared services in the GRT, requires a different approach to leak routes for access to shared services. The process still requires VRF-lite, the Layer-3 (IP-based handoff) automation in Cisco DNA Center, and BGP, although there are more command line elements to configure on the fusion routers. These begin with IP prefix list. A prefix list is created for each VN in the fabric that references each of the subnets for that VN – ostensibly each of the host pools. A route-map is created for each VN in the fabric that matches the prefix list. Finally, the VRF configuration imports and exports routes that are filtered based on these route-maps.

The primary consideration when using this method is the number of touch points. Using IP-based transits when provisioning border nodes, Cisco DNA Center automates the necessary configuration so that the fabric prefixes are advertised via BGP. When share services are in a dedicate VRF, apart from adding a new VN and associated host pool, once route leaking is configured on the fusion router additional touch points should not be required. If a new subnet is added to an existing host pool, it is provisioned and advertised to the fusion router through Cisco DNA Center automation.

When shared services are not in a dedicated VRF, there are touch points on the fusion router every time a new subnet is added to a host pool for any VN. While adding an additional line to a prefix list is relatively straightforward, this process must be completed on all fusion routers for each added subnet as the network grows.

The following example shows an alternative approach to providing access to shared services as shown in [Operate Section > Process 1](#). The first four procedures are the same with the deviation beginning at the fifth. Please review these procedures, as this section begins with the VRFs created on the fusion router. Using the iteration described here, the SHARED_SERVICES VRF does not exist, and therefore the fusion routers' interfaces and BGP configuration will be slightly different than described in the main section of the guide. Unless otherwise noted, this configuration will occur on both fusion routers.

Procedure 1: Create Layer-3 Connectivity Between Border Nodes and Fusion Routers

The process begins by creating the IP connectivity. The VLANs and corresponding IP addresses are pulled from Cisco DNA Center by observing the **View Device Info** tab of the provisioned Internal border nodes.



Pay close attention to interface VLANs 3001 and 3005 on the first fusion router and interface VLANs 3009 and 3013 on the second fusion router. These are the primary difference in the Layer-3 connectivity from the procedures in the main text.

1. In the global configuration on the first fusion router, create the VLANs.

```
vlan 3001-3008
exit
```

2. Configure each SVI to forward for a VRF and to have the applicable IP address based on the information from **View Device Info**.

Remember to use *no shutdown* to bring the SVI to the up/up state.

```
interface vlan 3001
ip address 172.16.111.2 255.255.255.2
no shutdown

interface vlan 3002
vrf forwarding CAMPUS
ip address 172.16.111.6 255.255.255.252
```

```

no shutdown

interface vlan 3003
vrf forwarding RESEARCH
ip address 172.16.111.10 255.255.255.252
no shutdown

interface vlan 3004
vrf forwarding GUEST
ip address 172.16.111.14 255.255.255.252
no shutdown

```

3. Configure a trunk interface that allows the defined VLANs.

```

interface TenGigabitEthernet1/1/2
description Connected to Internal Border 6832-02 | TenGig 1/11
switchport
switchport mode trunk
switchport trunk allowed vlan 3001-3004
no shutdown

```

4. Repeat these steps for the second interface on the first fusion router.

```

interface vlan 3005
ip address 172.16.111.18 255.255.255.252
no shutdown

interface vlan 3006
vrf forwarding CAMPUS
ip address 172.16.111.22 255.255.255.252
no shutdown

interface vlan 3007
vrf forwarding RESEARCH
ip address 172.16.111.26 255.255.255.252
no shutdown

interface vlan 3008
vrf forwarding GUEST
ip address 172.16.111.30 255.255.255.252
no shutdown

interface TenGigabitEthernet1/1/2
description Connected to Internal Border 6832-01 | TenGig 1/11
switchport
switchport mode trunk
switchport trunk allowed vlan 3005-3008
no shutdown

```

5. Verify IP connectivity between the first fusion router and the border nodes using ping commands.

```

ping 172.16.111.1
ping vrf CAMPUS 172.16.111.5
ping vrf RESEARCH 172.16.111.9
ping vrf GUEST 172.16.111.13

ping 172.16.111.17
ping vrf CAMPUS 172.16.111.21
ping vrf RESEARCH 172.16.111.25
ping vrf GUEST 172.16.111.29

```

6. Repeat these steps on the second fusion router.

```

! First Interface and VLAN set
vlan 3009-3016
exit

```

```

interface vlan 3009
  ip address 172.16.111.34 255.255.255.2
  no shutdown

interface vlan 3010
  vrf forwarding CAMPUS
  ip address 172.16.111.38 255.255.255.252
  no shutdown

interface vlan 3011
  vrf forwarding RESEARCH
  ip address 172.16.111.42 255.255.255.252
  no shutdown

interface vlan 3012
  vrf forwarding GUEST
  ip address 172.16.111.46 255.255.255.252
  no shutdown

interface TenGigabitEthernet1/1/2
  description Connected to Internal Border 6832-02 | TenGig 1/12
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 3009-3012
  no shutdown

! Second Interface and VLAN set
interface vlan 3013
  ip address 172.16.111.50 255.255.255.252
  no shutdown

interface vlan 3014
  vrf forwarding CAMPUS
  ip address 172.16.111.54 255.255.255.252
  no shutdown

interface vlan 3015
  vrf forwarding RESEARCH
  ip address 172.16.111.58 255.255.255.252
  no shutdown

interface vlan 3016
  vrf forwarding GUEST
  ip address 172.16.111.62 255.255.255.252
  no shutdown

interface TenGigabitEthernet1/1/2
  description Connected to Internal Border 6832-01 | TenGig 1/12
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 3013-3016
  no shutdown

```

7. Verify IP connectivity between the second fusion router and the border nodes using ping commands.

```

ping 172.16.111.33
ping vrf CAMPUS 172.16.111.37
ping vrf RESEARCH 172.16.111.41
ping vrf GUEST 172.16.111.45

ping 172.16.111.49
ping vrf CAMPUS 172.16.111.53
ping vrf RESEARCH 172.16.111.57
ping vrf GUEST 172.16.111.61

```

Procedure 2: Establish BGP Adjacencies Between Fusion Routers and Border Nodes

The BGP configuration is extremely similar to the one shown in the main text. The primary difference is that shared services are in the global routing table and that neighbor configuration must be done using the default *IPv4 address-family* under *router BGP*. Additionally, the shared services prefixes must be explicitly advertised in BGP. The prefix lists and route-maps require an entry in the BGP table, not the RIB, to match these routes and provide leaking.



Tech Tip

Note that the *address-family ipv4 vrf SHARED_SERVICES* is not used to create the adjacencies. Rather, *address-family ipv4* is used, indicating that the shared services are located in the Global Routing Table.

1. Create the BGP process on the first fusion router.
2. Use the corresponding Autonomous-System defined in the IP-based transit from early steps. As a recommended practice, the Loopback 0 interface is used as the BGP router ID.

```
router bgp 65333
  bgp router-id interface Loopback0
```

3. Under the default *IPv4 address-family* complete the following steps:
 - a. Define the first border node neighbor and its corresponding AS number. Remember that this corresponds with the INFRA_VN on the border nodes.
 - b. Define the update source to use the applicable SVI.
 - c. Activate the exchange of NLRI with the first border node.

```
router bgp 65333
  neighbor 172.16.111.1 remote-as 65001
  neighbor 172.16.111.1 update-source Vlan3001
```

```
address-family ipv4
  neighbor 172.16.111.1 activate
```

4. Repeat these steps for the adjacency with the redundant border node. Remember to use the default *IPv4 address-family*.

```
router bgp 65333
  neighbor 172.16.111.17 remote-as 65001
  neighbor 172.16.111.17 update-source Vlan3001
```

```
address-family ipv4
  neighbor 172.16.111.17 activate
```

5. Advertise the shared services routes into BGP.

```
address-family ipv4
  network 198.51.100.0 mask 255.255.255.0
```

6. Configure the adjacencies for the remaining VRFs using the applicable SVIs as the update source for the peering. These entries will use *address-family ipv4 vrf*.

```
address-family ipv4 vrf CAMPUS
  neighbor 172.16.111.5 remote-as 65001
  neighbor 172.16.111.5 update-source Vlan3002
  neighbor 172.16.111.5 activate
```

```
neighbor 172.16.111.21 remote-as 65001
neighbor 172.16.111.21 update-source Vlan3006
neighbor 172.16.111.21 activate
```

```
address-family ipv4 vrf RESEARCH
  neighbor 172.16.111.9 remote-as 65001
  neighbor 172.16.111.9 update-source Vlan3003
  neighbor 172.16.111.9 activate
```

```
neighbor 172.16.111.25 remote-as 65001
neighbor 172.16.111.25 update-source Vlan3007
neighbor 172.16.111.25 activate
```

```

address-family ipv4 vrf GUEST
  neighbor 172.16.111.13 remote-as 65001
  neighbor 172.16.111.13 update-source Vlan3004
  neighbor 172.16.111.13 activate

  neighbor 172.16.111.29 remote-as 65001
  neighbor 172.16.111.29 update-source Vlan3008
  neighbor 172.16.111.29 activate

```

7. Once this configuration has been done on the first fusion router, repeat it on the second fusion router. Remember to advertise the shared services routes.

```

router bgp 65333
  bgp router-id interface Loopback0

address-family ipv4
  neighbor 172.16.111.33 remote-as 65001
  neighbor 172.16.111.33 update-source Vlan3009
  neighbor 172.16.111.33 activate

  neighbor 172.16.111.49 remote-as 65001
  neighbor 172.16.111.49 update-source Vlan3013
  neighbor 172.16.111.49 activate

  network 198.51.100.0 mask 255.255.255.0

address-family ipv4 vrf CAMPUS
  neighbor 172.16.111.37 remote-as 65001
  neighbor 172.16.111.37 update-source Vlan3010
  neighbor 172.16.111.37 activate

  neighbor 172.16.111.53 remote-as 65001
  neighbor 172.16.111.53 update-source Vlan3014
  neighbor 172.16.111.53 activate

address-family ipv4 vrf RESEARCH
  neighbor 172.16.111.42 remote-as 65001
  neighbor 172.16.111.42 update-source Vlan3011
  neighbor 172.16.111.42 activate

  neighbor 172.16.111.57 remote-as 65001
  neighbor 172.16.111.57 update-source Vlan3015
  neighbor 172.16.111.57 activate

address-family ipv4 vrf GUEST
  neighbor 172.16.111.45 remote-as 65001
  neighbor 172.16.111.45 update-source Vlan3012
  neighbor 172.16.111.45 activate

  neighbor 172.16.111.61 remote-as 65001
  neighbor 172.16.111.61 update-source Vlan3016
  neighbor 172.16.111.61 activate

```

 Tech Tip

In the configuration above, some of the neighbor commands for **172.16.111.1**, **172.16.111.17**, **172.16.111.33**, and **172.16.111.49** were explicitly placed under **address-family ipv4** and others under the **router bgp**. The neighbor commands would all have been accepted by the CLI directly beneath **router bgp**.

Beginning around the 12.2(15)T train, the CLI began to support the MP-BGP address family configuration. This led to a command **bgp upgrade-cli** that [converts](#) the NLRI syntax to the AFI syntax. This had no direct effect on the routing information exchanged; it was a cosmetic update to help simplify syntax and command structures.

The fusion router could be a device running slightly older IOS code in the 12.x family as it would still support the technical requirements. The AFI syntax is shown by default in modern Cisco software, although some commands are still accepted at the default level under the BGP router configuration.

Procedure 3: Create IP Prefix lists to Match Fabric Subnets

After BGP adjacencies are established, configure the routes to be leaked by using prefix lists. Prefix lists can match any prefix length or prefix range. These prefix lists are referenced in the route-maps in the next procedure. For granularity, match the specific prefix-length of the fabric subnets ensuring the correct NLRI is leaked between the Global Routing Table and the VRF tables.

Complete the following configuration on both fusion routers.

1. Create a prefix list that matches the CAMPUS VN subnets for the entire fabric domain.

```
ip prefix-list CAMPUS_VN_NETWORKS seq 5 permit 172.16.111.0/24
ip prefix-list CAMPUS_VN_NETWORKS seq 10 permit 172.16.121.0/24
ip prefix-list CAMPUS_VN_NETWORKS seq 15 permit 172.16.131.0/24
ip prefix-list CAMPUS_VN_NETWORKS seq 20 permit 172.16.141.0/24
ip prefix-list CAMPUS_VN_NETWORKS seq 25 permit 172.16.151.0/24
ip prefix-list CAMPUS_VN_NETWORKS seq 30 permit 172.16.161.0/24
ip prefix-list CAMPUS_VN_NETWORKS seq 35 permit 172.16.171.0/24
ip prefix-list CAMPUS_VN_NETWORKS seq 40 permit 172.16.181.0/24
ip prefix-list CAMPUS_VN_NETWORKS seq 45 permit 172.16.191.0/24
ip prefix-list CAMPUS_VN_NETWORKS seq 50 permit 172.16.241.0/24
```

2. Create a prefix list that matches the GUEST VN subnets for the entire fabric domain.

```
ip prefix-list GUEST_VN_NETWORKS seq 5 permit 172.16.113.0/24
ip prefix-list GUEST_VN_NETWORKS seq 10 permit 172.16.123.0/24
ip prefix-list GUEST_VN_NETWORKS seq 15 permit 172.16.133.0/24
ip prefix-list GUEST_VN_NETWORKS seq 20 permit 172.16.143.0/24
ip prefix-list GUEST_VN_NETWORKS seq 25 permit 172.16.153.0/24
ip prefix-list GUEST_VN_NETWORKS seq 30 permit 172.16.163.0/24
ip prefix-list GUEST_VN_NETWORKS seq 35 permit 172.16.173.0/24
ip prefix-list GUEST_VN_NETWORKS seq 40 permit 172.16.183.0/24
ip prefix-list GUEST_VN_NETWORKS seq 45 permit 172.16.193.0/24
ip prefix-list GUEST_VN_NETWORKS seq 50 permit 172.16.243.0/24
```

3. Create a prefix list that matches the RESEARCH VN subnets for the entire fabric domain.

```
ip prefix-list RESEARCH_VN_NETWORKS seq 5 permit 172.16.114.0/24
ip prefix-list RESEARCH_VN_NETWORKS seq 10 permit 172.16.124.0/24
ip prefix-list RESEARCH_VN_NETWORKS seq 15 permit 172.16.134.0/24
ip prefix-list RESEARCH_VN_NETWORKS seq 20 permit 172.16.144.0/24
ip prefix-list RESEARCH_VN_NETWORKS seq 25 permit 172.16.154.0/24
ip prefix-list RESEARCH_VN_NETWORKS seq 30 permit 172.16.164.0/24
ip prefix-list RESEARCH_VN_NETWORKS seq 35 permit 172.16.174.0/24
ip prefix-list RESEARCH_VN_NETWORKS seq 40 permit 172.16.184.0/24
ip prefix-list RESEARCH_VN_NETWORKS seq 45 permit 172.16.194.0/24
ip prefix-list RESEARCH_VN_NETWORKS seq 50 permit 172.16.244.0/24
```

4. Create a prefix list that matches the shared services subnet behind the fusion router.

```
ip prefix-list SHARED_SERVICES_NETWORKS seq 5 permit 198.51.100.0/24
```

Procedure 4: Create Route-maps that Match IP Prefix Lists

The import and export commands that will be used in future steps selectively leak routes through filtering based on route-maps. The route-maps are created in this procedure and directly reference the prefix lists created in the previous procedure. This ensures only specific routes are matched.

1. Create a route map that matches the CAMPUS VN prefix list.

```
route-map CAMPUS_VN_NETWORKS permit 10
  match ip address prefix-list CAMPUS_VN_NETWORKS
```

2. Create a route map that matches the GUEST VN prefix list.

```
route-map GUEST_VN_NETWORKS permit 10
  match ip address prefix-list GUEST_VN_NETWORKS
```

3. Create a route map that matches the RESEARCH VN prefix list.

```
route-map RESEARCH_VN_NETWORKS permit 10
  match ip address prefix-list RESEARCH_VN_NETWORKS
```

4. Create a route map that matches the shared services prefix list.

```
route-map SHARED_SERVICES_NETWORKS permit 10
  match ip address prefix-list SHARED_SERVICES_NETWORKS
```

Procedure 5: Leak routes

Route leaking is done by importing and exporting route-maps under the VRF configuration. The VRF should export prefixes belonging to itself using a route-map. The VRF should also import desired routes used for access to shared services using a route-map.

The import command will process through the route map, and prefixes that match are imported into the VRF.

1. Configure the CAMPUS VRF for route leaking the shared services prefixes.

```
vrf definition CAMPUS
  address-family ipv4
    import ipv4 unicast map SHARED_SERVICES_NETWORKS
    export ipv4 unicast map CAMPUS_VN_NETWORKS
```

2. Configure the GUEST VRF for route leaking the shared services prefixes.

```
vrf definition GUEST
  address-family ipv4
    import ipv4 unicast map SHARED_SERVICES_NETWORKS
    export ipv4 unicast map GUEST_VN_NETWORKS
```

3. Configure the RESEARCH VRF for route leaking the shared services prefixes.

```
vrf definition RESEARCH
  address-family ipv4
    import ipv4 unicast map SHARED_SERVICES_NETWORKS
    export ipv4 unicast map RESEARCH_VN_NETWORKS
```

4. (Optional) Verify the leaked routes on the CLI of the Internal border nodes.

```
6832-01# show ip route vrf CAMPUS
```

```
Routing Table: CAMPUS
```

```
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP  
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area  
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2  
E1 - OSPF external type 1, E2 - OSPF external type 2  
i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2  
ia - IS-IS inter area, * - candidate default, U - per-user static route  
o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP  
a - application route  
+ - replicated route, % - next hop override, p - overrides from PFR
```

```
Gateway of last resort is not set
```

```
10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks  
C 10.10.10.20/30 is directly connected, Vlan101  
L 10.10.10.21/32 is directly connected, Vlan101  
172.16.0.0/16 is variably subnetted, 7 subnets, 3 masks  
C 172.16.111.4/30 is directly connected, Vlan3019  
L 172.16.111.5/32 is directly connected, Vlan3019  
B 172.16.112.0/24 [200/0] via 192.168.10.1, 5d16h  
C 172.16.112.1/32 is directly connected, Loopback1024  
B 172.16.122.0/24 [200/10] via 192.168.200.100, 18:26:54  
B 172.16.132.0/24 [200/10] via 192.168.200.100, 18:26:54  
B 172.16.152.0/24 [200/10] via 192.168.200.100, 18:26:54  
B 198.51.100.0/24 [20/0] via 172.16.111.6, 00:02:54
```

```
6832-01#
```

```
6832-01# show ip route vrf GUEST
```

```
Routing Table: GUEST
```

```
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP  
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area  
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2  
E1 - OSPF external type 1, E2 - OSPF external type 2  
i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2  
ia - IS-IS inter area, * - candidate default, U - per-user static route  
o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP  
a - application route  
+ - replicated route, % - next hop override, p - overrides from PFR
```

```
Gateway of last resort is not set
```

```
10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks  
C 10.10.10.24/30 is directly connected, Vlan102  
L 10.10.10.25/32 is directly connected, Vlan102  
172.16.0.0/16 is variably subnetted, 7 subnets, 3 masks  
C 172.16.111.12/30 is directly connected, Vlan3021  
L 172.16.111.13/32 is directly connected, Vlan3021  
B 172.16.113.0/24 [200/0] via 192.168.10.1, 5d17h  
C 172.16.113.1/32 is directly connected, Loopback1027  
B 172.16.123.0/24 [200/10] via 192.168.200.100, 18:38:43  
B 172.16.133.0/24 [200/10] via 192.168.200.100, 18:38:43  
B 172.16.153.0/24 [200/10] via 192.168.200.100, 18:38:43  
B 198.51.100.0/24 [20/0] via 172.16.111.14, 00:03:43
```

```
6832-01#
```

```
6832-01# show ip route vrf RESEARCH
```

```
Routing Table: RESEARCH
```

```
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP  
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area  
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2  
E1 - OSPF external type 1, E2 - OSPF external type 2  
i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2  
ia - IS-IS inter area, * - candidate default, U - per-user static route  
o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP  
a - application route  
+ - replicated route, % - next hop override, p - overrides from PFR
```

```
Gateway of last resort is not set

  10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C   10.10.10.28/30 is directly connected, Vlan103
L   10.10.10.29/32 is directly connected, Vlan103
  172.16.0.0/16 is variably subnetted, 7 subnets, 3 masks
C   172.16.111.8/30 is directly connected, Vlan3020
L   172.16.111.9/32 is directly connected, Vlan3020
B   172.16.114.0/24 [200/0] via 192.168.10.1, 5d16h
C   172.16.114.1/32 is directly connected, Loopback1028
B   172.16.124.0/24 [200/10] via 192.168.200.100, 18:32:03
B   172.16.134.0/24 [200/10] via 192.168.200.100, 18:32:03
B   172.16.154.0/24 [200/10] via 192.168.200.100, 18:32:03
B   198.51.100.0/24 [20/0] via 172.16.111.10, 00:04:03
6832-01#
```



Tech Tip

It may take several minutes for the routes to leak and be populated on the Internal border nodes. In Cisco software, [import actions](#) are triggered when a new routing update is received or when routes are withdrawn. During the initial BGP update period, the import action is postponed allowing for quicker BGP convergence. Once BGP converges, incremental BGP updates are evaluated immediately and qualified prefixes are imported as they are received.

Appendix E: IBGP Between Redundant Devices

IBGP for Switching Platforms

Routing platforms such as the Integrated Services Routers (ISRs) and Aggregation Services Routers (ASR) support subinterfaces (dot1q VLAN subinterface). Layer-3 switches may or may not support them. However, Layer-3 switches should universally support Switched Virtual Interfaces (SVIs) and trunk links.

This procedure will first create the VLANs, create the corresponding SVIs, assign them to an existing VRF, and give them an IP address. Once IP connectivity is established and verified, the BGP adjacencies will be created using the SVIs as the update source.



There are multiple reserved VLANs in Cisco DNA Center. VLANs below 1000 are **generally** safe to use for creating the IBGP adjacencies. Because the VLANs are locally significant between the redundant devices, the same VLANs could be used for multiple redundant pairs.

VLAN IDs 1 to 1005 are global in the VTP domain and can be defined on other network devices in the VTP domain. Extended range VLANs can only be used with subinterfaces when the device is in transparent mode. In VTP client or server mode, normal-range VLANs are excluded from subinterfaces.

1. On the CLI of the first External border node router, create the VLANs.

```
vlan 100-104
exit
```

2. Define an SVI that will forward for the Global Routing Table and configure the IP address.

```
interface Vlan100
description Connected to 6832-02 | GRT
ip address 10.10.10.17 255.255.255.252
no shutdown
```

3. Configure the SVIs and IP addresses for the remaining VRFs.

```
interface Vlan101
description Connected to 6832-02 | VRF CAMPUS
vrf forwarding CAMPUS
ip address 10.10.10.21 255.255.255.252
no shutdown
```

```
interface Vlan102
description Connected to 6832-02 | VRF GUEST
vrf forwarding GUEST
ip address 10.10.10.25 255.255.255.252
no shutdown
```

```
interface Vlan103
description Connected to 6832-02 | VRF RESEARCH
vrf forwarding RESEARCH
ip address 10.10.10.29 255.255.255.252
no shutdown
```

4. Define the interface between the redundant border nodes as a trunk and allow the applicable VLANs.

```
interface TenGigabitEthernet1/15
switchport
switchport mode trunk
switchport trunk allowed vlan 100-104
no shutdown
```

5. Complete the corresponding configuration on the redundant border node.

```
vlan 100-104
exit
```

```

interface Vlan100
description Connected to 6832-01 | GRT
ip address 10.10.10.18 255.255.255.252
no shutdown

interface Vlan101
description Connected to 6832-01 | VRF CAMPUS
vrf forwarding CAMPUS
ip address 10.10.10.22 255.255.255.252
no shutdown

interface Vlan102
description Connected to 6832-01 | VRF GUEST
vrf forwarding GUEST
ip address 10.10.10.26 255.255.255.252
no shutdown

interface Vlan103
description Connected to 6832-01 | VRF RESEARCH
vrf forwarding RESEARCH
ip address 10.10.10.30 255.255.255.252
no shutdown

interface TenGigabitEthernet1/15
switchport
switchport mode trunk
switchport trunk allowed vlan 100-104
no shutdown

```

6. Verify connectivity between the SVIs.

```

ping 10.10.10.17
ping vrf CAMPUS 10.10.10.21
ping vrf RESEARCH 10.10.10.25
ping vrf GUEST 10.10.10.29

```

7. On the first border node, enable IBGP between the redundant peer for the Global Routing Table. Use the SVI as the update source.

```

router bgp 65001
neighbor 10.10.10.18 remote-as 65001
neighbor 10.10.10.18 update-source Vlan100

address-family ipv4
neighbor 10.10.10.18 activate

```

8. Enable IBGP between the redundant peer for the remainder of the VRFs. Use the SVI as the update source.

```

address-family ipv4 vrf CAMPUS
neighbor 10.10.10.22 remote-as 65001
neighbor 10.10.10.22 update-source Vlan101
neighbor 10.10.10.22 activate

address-family ipv4 vrf GUEST
neighbor 10.10.10.26 remote-as 65001
neighbor 10.10.10.26 update-source Vlan102
neighbor 10.10.10.26 activate

address-family ipv4 vrf RESEARCH
neighbor 10.10.10.30 remote-as 65001
neighbor 10.10.10.30 update-source Vlan103
neighbor 10.10.10.30 activate

```

9. Complete the corresponding configuration on the redundant border node.

```

router bgp 65001
neighbor 10.10.10.17 remote-as 65001
neighbor 10.10.10.17 update-source Vlan100

```

```

address-family ipv4
  neighbor 10.10.10.17 activate

address-family ipv4 vrf CAMPUS
  neighbor 10.10.10.21 remote-as 65001
  neighbor 10.10.10.21 update-source Vlan101
  neighbor 10.10.10.21 activate

address-family ipv4 vrf GUEST
  neighbor 10.10.10.25 remote-as 65001
  neighbor 10.10.10.25 update-source Vlan102
  neighbor 10.10.10.25 activate

address-family ipv4 vrf RESEARCH
  neighbor 10.10.10.29 remote-as 65001
  neighbor 10.10.10.29 update-source Vlan103
  neighbor 10.10.10.29 activate

```

10. Verify the BGP adjacency for the Global Routing Table SVI.

```
6832-01# show ip bgp summary
```

```

BGP router identifier 192.168.10.3, local AS number 65001
BGP table version is 120, main routing table version 120
3 network entries using 432 bytes of memory
7 path entries using 588 bytes of memory
6/3 BGP path/bestpath attribute entries using 960 bytes of memory
2 BGP rrinfo entries using 48 bytes of memory
2 BGP AS-PATH entries using 48 bytes of memory
1 BGP community entries using 24 bytes of memory
3 BGP extended community entries using 72 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 2172 total bytes of memory
1 received paths for inbound soft reconfiguration
BGP activity 410/389 prefixes, 880/833 paths, scan interval 60 secs

```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
10.10.10.18	4	65001	40709	40717	120	0	0	1w4d	2
172.16.111.2	4	65333	31708	31727	120	0	0	1w6d	1
192.168.10.1	4	65001	41517	40730	120	0	0	1w4d	1
192.168.10.2	4	65001	41580	40723	120	0	0	1w4d	1

```
6831-01#
```

11. Verify the BGP adjacencies for the VRFs.

```
6832-01# show ip bgp vpnv4 all summary
```

```
BGP router identifier 192.168.10.3, local AS number 65001
BGP table version is 2123, main routing table version 2123
18 network entries using 2808 bytes of memory
40 path entries using 3360 bytes of memory
26/12 BGP path/bestpath attribute entries using 4368 bytes of memory
2 BGP rrinfol entries using 48 bytes of memory
2 BGP AS-PATH entries using 48 bytes of memory
1 BGP community entries using 24 bytes of memory
3 BGP extended community entries using 72 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 10728 total bytes of memory
BGP activity 410/389 prefixes, 880/833 paths, scan interval 60 secs
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
10.10.10.22	4	65001	40709	40714	2123	0	0	1w4d	2
10.10.10.26	4	65001	40711	40717	2123	0	0	1w4d	2
10.10.10.30	4	65001	40714	40707	2123	0	0	1w4d	2
172.16.111.6	4	65333	32241	32286	2123	0	0	1w6d	1
172.16.111.10	4	65333	31715	31733	2123	0	0	1w6d	1
172.16.111.14	4	65333	31727	31685	2123	0	0	1w6d	1
192.168.10.1	4	65001	41519	40732	2123	0	0	1w4d	14
192.168.10.2	4	65001	41582	40725	2123	0	0	1w4d	14

```
6832-01#
```

About this guide

ALL DESIGNS, SPECIFICATIONS, STATEMENTS, INFORMATION, AND RECOMMENDATIONS (COLLECTIVELY, "DESIGNS") IN THIS MANUAL ARE PRESENTED "AS IS," WITH ALL FAULTS. CISCO AND ITS SUPPLIERS DISCLAIM ALL WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE. IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THE DESIGNS, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

THE DESIGNS ARE SUBJECT TO CHANGE WITHOUT NOTICE. USERS ARE SOLELY RESPONSIBLE FOR THEIR APPLICATION OF THE DESIGNS. THE DESIGNS DO NOT CONSTITUTE THE TECHNICAL OR OTHER PROFESSIONAL ADVICE OF CISCO, ITS SUPPLIERS OR PARTNERS. USERS SHOULD CONSULT THEIR OWN TECHNICAL ADVISORS BEFORE IMPLEMENTING THE DESIGNS. RESULTS MAY VARY DEPENDING ON FACTORS NOT TESTED BY CISCO.

CCDE, CCENT, Cisco Eos, Cisco Lumin, Cisco Nexus, Cisco StadiumVision, Cisco TelePresence, Cisco WebEx, the Cisco logo, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unified Computing System (Cisco UCS), Cisco UCS B-Series Blade Servers, Cisco UCS C-Series Rack Servers, Cisco UCS S-Series Storage Servers, Cisco UCS Manager, Cisco UCS Management Software, Cisco Unified Fabric, Cisco Application Centric Infrastructure, Cisco Nexus 9000 Series, Cisco Nexus 7000 Series, Cisco Prime Data Center Network Manager, Cisco NX-OS Software, Cisco MDS Series, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website 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. (0809R)

© 2019 Cisco Systems, Inc. All rights reserved.

Feedback & Discussion

For comments and suggestions about our guides, please join the discussion on [Cisco Community](#).