Cisco Cloud Fabric

A Cisco Validated Design case study

February 8, 2026

Cisco's new SaaS-based cloud-managed campus fabric solution is now generally available. Customers can deploy and manage a BGP EVPN VXLAN campus fabric via the Meraki Dashboard UI. The campus fabric configurations are orchestrated from the cloud, eliminating the need for complex manually created CLI configuration. Combining new Meraki capabilities such as Routed Ports, VRFs, and cloud CLI, with the existing robust onboarding, orchestration and management capabilities brings "Meraki Easy" to the world of campus fabric.

This design guide provides an overview of the case study used to validate the Campus Cloud Fabric solutions. It discusses the architecture and components of the solution, including the routed underlay, VXLAN-based fabric overlays, BGP EVPN control plane, and VRF-based segmentation.

The intended audience is for network architects, design engineers, and implementation engineers responsible for planning, deploying, and operating campus networks. It is also relevant for technical leaders and operations teams evaluating cloud-managed campus fabric solutions. Readers are expected to have a working knowledge of enterprise campus networking concepts.

This design guide focuses on a validated deployment that follows Cisco best-practice recommendations, including the use of a Layer 3 routed underlay, dynamic routing with OSPF and BGP, and redundant fabric borders connected to VRF-aware upstream handoff devices. The guide highlights key design considerations and workflows for access, wireless, DHCP, Adaptive Policy/Trustsec, and multi-VRF deployments to support scalable and resilient Campus Cloud Fabric operations through the cloud UI.

Use Case: Cloud-managed campus fabric with Cisco Cloud Fabric

In this validated use case, a three-tier spine-and-leaf campus fabric deploys with redundant spines, leaves, and dedicated border devices. A Layer 3 routed underlay using OSPF provides fast convergence and predictable connectivity, while VXLAN overlays deliver scalable segmentation across multiple VRFs. Border devices use automated eBGP peering to integrate the fabric with upstream networks, maintaining VRF separation beyond the fabric edge.

This approach allows Cisco customers to:

- Simplify campus fabric deployment through UI-driven workflows
- Reduce operational risk by eliminating manual overlay CLI configuration
- Maintain Cisco-recommended Layer 3 design best practices
- Scale segmentation using VRFs without redesigning the physical network
- Integrate access, wireless, DHCP, and Adaptive Policy/Trustsec consistently across the fabric

The validated design demonstrates how Cisco Cloud Fabric enables network architects and operators to deploy and operate a modern campus fabric that combines the architectural rigor of traditional Cisco campus designs with the simplicity and automation of cloud-managed operations, all while remaining aligned with Cisco enterprise networking standards and lifecycle practices.

Validated deployment scope and case study context

Cisco Cloud Fabric supports designs with dedicated border devices and border-on-spine deployments (border-on-leaf is not currently supported). The upstream devices that connect to the borders may vary in capabilities resulting in design differences. Given these permutations, two Cisco Validated case study networks were created for validation testing.

This document briefly references a border-on-spine case study; however, the remainder of the document focuses on the full-featured IOS-XE-based device case study that follows best practice recommendations using redundant borders and redundant upstream handoff devices.

Case study: Border-on-spine topology

This case study uses a border-on-spine topology with a pair of MX105 security appliances in an active/warm standby configuration as the handoff devices. In this configuration, the MX functions as a single logical device. The MX does not support VRFs and therefore acts as a fusion device, combining the underlay and overlay routing domains into a single routing domain outside the fabric. Static routes are used for underlay routing between the MX and the borders, and eBGP is used for overlay routing between the MX and the borders. This configuration is supported but not ideal.

Case study: Full-featured IOS-XE-based devices

This case study uses a pair of full-featured IOS-XE–based devices in a VRF-Lite active/active configuration. OSPF is used between the handoff devices and the borders for dynamic underlay routing, and eBGP is used between the borders and the handoff devices for dynamic overlay routing. VRF separation is maintained to and beyond the handoff devices, with a fusion device routed further upstream. This is the recommended deployment option and, as such, is detailed in this document.

In both case study setups, the links between the borders and the handoff devices are IEEE 802.1Q trunks. These trunks carry both underlay and overlay traffic. For overlay routing, eBGP is used, and the configurations are automated on the border side as part of the fabric workflow. In typical fabric networks, the underlay IGP is redistributed into BGP on the borders, and eBGP is also used for underlay routing between the borders and the handoff devices, typically in a VRF-Lite handoff when the handoff devices are VRF-aware. However, redistribution from OSPF into BGP is not currently available on the switches. As a result, the available options are static routing or extending OSPF from the underlay to the handoff devices. The validated setup uses OSPF due to its dynamic nature.

Architectural and Foundational Components

Fabrics – Fabrics are built using an underlay and overlay model.

Fabric Device Roles – Fabrics are typically deployed following a Clos architecture using a spine-and-leaf approach. The Cisco Cloud Fabric architecture uses a three-tier model similar to traditional Core–Distribution–Access layer designs, with spines at the core, leaves at the distribution layer, and traditional switches and Access Points at the access layer. The border function is deployed on a limited number of devices within the fabric to connect the fabric to external networks. Cisco Cloud Fabric supports both dedicated border deployments and border-on-spine deployments.

Fabric Underlay – In campus fabric environments, the underlay uses an IGP with Layer 3 point-to-point links to eliminate spanning tree, enable rapid convergence, and support equal-cost multi-pathing (ECMP). Cisco Cloud Fabric uses Routed Ports and OSPF to establish underlay reachability. The fabric control plane runs MP-BGP EVPN (AFI 25 / SAFI 70) over the underlay to distribute overlay endpoint reachability, with BGP peering established using OSPF reachability. The data plane uses VXLAN encapsulation to transport overlay traffic. The underlay is manually configured in the current Cisco Cloud Fabric solution. The fabric underlay connects to external handoff devices through border nodes using OSPF routing.

Fabric Overlay(s) – An overlay is a logical network in which traffic is encapsulated and transported between underlay devices. Cisco Cloud Fabric uses IP-based underlay packets sourced and destined to devices within the underlay routing domain, with overlay traffic carried using VXLAN encapsulation. Virtual Network Identifier (VNI) and Security Group Tag (SGT) information is conveyed in the VXLAN header. VRFs define logical overlays on fabric devices, including leaves and borders, where VXLAN tunnels terminate. Fabric borders use eBGP for dynamic routing into and out of overlay networks. Overlay eBGP peering on fabric borders toward external handoff devices is fully automated, while external handoff devices are manually configured to interoperate.

Solution considerations

Handoff considerations

The border role must be present on at least one fabric device to connect the fabric to the outside world. Fabrics are typically configured with two borders for high availability. As stated in the Case Study Context section, Cisco Cloud Fabric supports designs with dedicated border devices and border-on-spine deployments (border-on-leaf is not currently supported). In either case, one or more upstream devices are required at the other end of the border handoff.

DHCP considerations

When the Cisco Cloud Fabric creates fabric overlay subnets, DHCP is automatically configured to relay requests to external DHCP servers. These relayed packets are sourced from an underlay loopback address on the fabric leaf where the requesting client's IP gateway resides. As a result, IP reachability is required from the DHCP infrastructure outside the fabric environment to the underlay loopback range inside the fabric, and the DHCP infrastructure must be VRF-aware.

Meraki device default behavior

In Meraki cloud-controlled mode, the default device configuration sets all ports as trunks with VLAN 1 configured as the native VLAN. The devices will attempt to obtain IP and DNS settings via DHCP and connect to the Meraki Cloud for configuration and management.

Best of breed hardware

Many modern Cisco switches and wireless devices can be configured locally or from the Meraki cloud. The supported hardware models can be switched between modes as required. The Cloud Fabric solution requires that device configurations are controlled by the Cloud. If existing catalyst switches are running 17.15.n and are Meraki-monitored with locally controlled CLI, they must be removed from their current Meraki network. Once the cloud-driven cleanup scripts complete, upgrade the switches to 17.18.2, and re-add them to the target Meraki network where the fabric will be created.

Note: When adding them back into that network, the option for them to be cloud-managed must be selected. The devices do not need to be unclaimed and reclaimed; they only need to be removed and re-added to the Meraki network.

Cloud reachability considerations

In environments where additional cabling and IP reachability to the Meraki Cloud infrastructure is available, it is convenient to use that for device to cloud management communications. Currently, this requires the use of additional "front panel" ports to connect to the dedicated cabling, and that IP connectivity and related routing is in the global routing table of the managed devices. In the Cisco Cloud Fabric solution, the UAC (Uplink Autoconfigure) and the resulting IP connectivity to the cloud-delivered control plane are in-band in the fabric underlay network. This validated deployment does not use additional dedicated cables for UAC traffic. This is relevant during the conversion from the default Layer 2 trunks to the recommended Layer 3 routed connections between the devices that will become fabric devices.

Note: It is important not to break UAC management's connectivity to the cloud during that conversion. Dedicated management cabling makes that simpler but is not always an option.

Fabric design best practice considerations

The Cisco Cloud Fabric solution is based on a 3-tier architecture with access layer routing occurring at SVIs configured on the fabric leaves at tier-2. The leaf routing can be deployed in one of three ways:

In order of preference

1. Routed SVI on leaf with unique subnet(s) per leaf
2. Routed SVI on leaf as Distributed Anycast Gateway (DAG) - same subnet on multiple leaves
3. Routed SVI on leaf as DAG with bridging - same subnet and bridging on multiple leaves

The ideal situation will have unique subnets deployed off each leaf which does not require a DAG. This approach provides the greatest scalability and is preferred whenever possible. When the same subnet is required on multiple leaves, a DAG routed configuration is used, and when bridging is required, a DAG bridged configuration is used. These options can be combined within a fabric; however, Cisco best practice is to route unique subnets and use DAG routed or DAG bridged designs only when necessary, minimizing the use of less preferred options. These best-practice recommendations are based on years of large-scale campus fabric customer deployment experience.

One additional Cisco best practice recommendation is to use a Layer 3 routed underlay, which currently requires some manual configuration. Layer 2 trunks, STP and SVIs can be used and may be advantageous in some brownfield migration scenarios, particularly when no spare cabling exists between fabric devices. However, a Layer 3 underlay is the preferred and proven approach with years of customer proven scalability and reliability and should always be the target end state.

Dot1x/Trustsec/Adaptive policy considerations

Cloud-provisioned 802.1X authentication for wired and wireless access is supported, along with dynamic VLAN assignment (by name or number) and filter-list assignment. Micro-segmentation using SGTs, including dynamic classification during 802.1X authentication, SGT propagation via Cisco Metadata Header and/or AutoVPN, and egress enforcement, is provided through Meraki Adaptive Policy on supported platforms. Because Cisco Cloud Fabric does not alter the access-layer architecture, these capabilities remain fully supported and unchanged.

Trunks between the access layer and fabric leaves can be configured for inline tagging using the Peer SGT Capable option. SGTs are preserved for overlay traffic, as they are carried in the VXLAN header across the fabric, and can optionally be propagated through the border using inline tagging.

Note: Border handoffs must be trunks with SVIs. The option to include the CMD header on routed port traffic is not currently supported, and the automation assumes the handoff links are trunks with SVIs.

Access layer to leaf EtherChannel

In Cisco Cloud Fabric, the access layer devices are not fabric-aware and operate unchanged and typically have trunks connecting to their upstream device. In a Cloud Fabric, the upstream device is a leaf that acts as the fabric edge and functions as a VXLAN Tunnel Endpoint (VTEP); border nodes also operate as VTEPs. EtherChannels are supported between access-layer devices and their upstream leaf to provide increased bandwidth and redundancy.

Access layer to leaf high availability considerations

High availability for access devices connecting to fabric leaves is provided through Multi-Chassis EtherChannel (MEC). Back-panel stacking forms a leaf stack, with multiple links from different stack members to the access device bundled into an EtherChannel using the aggregate function in the UI. This can be to a single access switch or an access switch stack providing even greater redundancy with a stack on both ends of the connection. Currently, an access device can only be connected to a single leaf or leaf stack.

MTU considerations

The UI defaults the system MTU to 9198, which maps to the system MTU command; however, the recommended best practice is an MTU of 9100, configured at the network level and applied to all switches in the network.

Note: All the devices at a given site are grouped together and referred to as a single network in the UI.

Fabric automation configures EBGP peer SVIs with an IP MTU of 9100 by default, which can be overridden if required. External devices connected to border nodes must be configured with matching MTUs for both underlay and overlay traffic. With an MTU of 9100, underlay traffic remains unfragmented up to 9100 bytes, while IPv4 VXLAN encapsulation adds 50 bytes of overhead, allowing unfragmented overlay payloads up to 9050 bytes.

Spanning tree considerations

Meraki deploys Multiple Spanning Tree (MST) by default. The best practice recommendation for Cisco Cloud Fabric is to use Rapid Per VLAN Spanning Tree (RPVST+). STP is configured at the network level in the UI, ensuring all switches at a site use the same STP version; this is part of the Layer 3 underlay preparation performed before running the fabric workflow.

Note: Care must be taken when changing this setting so as not to disrupt UAC connectivity.

The Routed Underlay deployment steps in this document provide guidance for this configuration. During the Layer 3 underlay conversion process, RPVST+ is enabled; spine bridge priorities are set to 4096 (spine1) and 8192 (spine2), and most interfaces are shut down with remaining trunks tightly restricted. After migration, STP no longer runs between fabric nodes, and leaf bridge priorities are set to 0, resulting in a leaf being the root bridge for accessing VLANs on the trunks to the downstream access devices.

Wireless considerations

The access layer operates as it traditionally does and is not fabric aware. Access Points will trunk to an access switch or leaf, while wireless capabilities and outcomes vary according to the physical environment and the fabric options deployed.

- Unique routed subnets per leaf - Using unique routed subnets in each leaf, as recommended best practice, clients are required to obtain a new IP address when roaming between APs on different leaves. An SSID may bridge to the same VLAN name or number; however, the VLAN represents a unique broadcast domain on each leaf. For clarity, fabric VLAN 100 (leaf1) is a separate broadcast domain and subnet from VLAN 100 (leaf2). This design is best suited for deployments where a leaf or leaf stack serves a single building or space, and seamless wireless coverage between buildings or spaces is not required.
- Common subnet on two or more leaves - When seamless roaming is required between buildings or spaces with contiguous wireless coverage, a routed DAG is used. This design routes the same IP subnet on multiple leaves, with each participating leaf using the same VLAN number and the SSID bridging to that VLAN. As a client roams between APs on different leaves, the client's IP address remains

unchanged. The upstream fabric detects client movement and updates routing to forward traffic to the new servicing leaf, effectively enabling a fabric leaf roam.

- Bridging is required – When bridging is also required, the routed DAG could include the bridge option and thus become a bridged DAG. This is the least preferred deployment option and should only be used when necessary and judiciously.

Deployment planning prerequisites and considerations

The following information should be gathered in advance of any configuration work.

Underlay network IP and VLAN information

Allocate sufficient IP address space accounting for the following:

- Underlay point to point links – Each leaf connects to each spine; the spines connect to each other, and if dedicated borders are used, each border connects to each spine. Additionally, the border handoff links must be considered in the underlay design. Links between fabric devices are converted to Layer 3 links, and /31 subnets are recommended. The links from the borders to the handoff devices remain trunks, only carry specific underlay and overlay VLANs, and function as routed interconnects using SVIs.
- Underlay IP address range entered during the fabric workflow – The automation uses IP addresses from the selected range for the underlay loopback 100 interface on each fabric device. Additional IP addresses from this range include the loopback 600 address on each spine for MSDP peering and the loopback 300 address on each spine, which is the PIM anycast RP address. Assume an additional 32 host addresses must be reserved for other infrastructure SVI IPs allocated from this pool. Ensure sufficient IP address space is provided for current and future needs, including adding new leaves.

  Note: Note that changing this IP address range requires rebuilding the fabric.

- Temporary DHCP pools – Pools are used when converting fabric devices with default configurations to a Layer 3 underlay configuration. These devices initially obtain addresses via DHCP on VLAN 1 and are then manually converted to routed interconnects. In the event of an RMA or new fabric device addition, this pool or a similar temp pool must be activated to facilitate the onboarding and conversion to Layer 3 underlay. This is also necessary if a fabric device is reset to factory default settings.
- Permanent DHCP pools for access layer devices – Access layer devices are managed in VLAN 1 by default. Cisco's best practice recommendation is to use a different VLAN. VLAN 2 was used in the Cisco Validated setup. Management traffic routes on an SVI that is manually created as part of normal setup. SVIs for access-layer devices terminate on their upstream leaf, requiring a unique subnet per leaf, which must be sized to support current and future downstream devices, including switches, access points, and cameras that require UAC control-plane connectivity to the cloud.
- Optional traditional subnets and VLANs – Subnets can be routed on a leaf that is not part of the fabric, which are no different than the management network required on each leaf. Subnets must be unique per leaf and manually configured. If traditional subnets are created, they will be part of the underlay routing domain, which may be relevant in a migration scenario. And existing deployment can be migrated in stages:
  - First to a traditional three-tier design with unique subnets per distribution switch,
  - Then to Layer 3 links between core and distribution, and
  - Finally, to fabric.
- Subnets can then be gradually transitioned from underlay SVI routing to fabric overlay SVI routing. If a traditional subnet is configured on a fabric leaf, a corresponding VLAN is required. The same VLAN number can be used across leaves if required, as those VLANs are discrete broadcast containers and are not connected over the routed connections between the spines and leaves. This assumes the use of a best practice Layer 3 underlay.

  Note: Note it is ideal if all the underlay networks can be summarized into a single prefix for summarization at the handoff devices out to the rest of the Intranet.

- Underlay OSPF information – It is important to plan a unique OSPF Area number for this handoff for the underlay between the borders and handoff devices. OSPF Area 0 is used between the spines and leaves, and spines and borders. Be prepared to match the interface MTU and media type on the handoff devices. The validation setup used OSPF Area 1 between the borders and handoff devices.
- Overlay networking IP and VLAN information – The subnets and associated VLAN numbers are entered in the fabric workflow. The planning depends on the options being configured:
    - Routed – 1 Subnet and VLAN number for each leaf selected (the VLAN number can reused if desired and can be beneficial in certain wireless scenarios)
    - Routed DAG – 1 subnet and VLAN number per set of selected leaves
    - Bridged DAG – 1 subnet and VLAN number per set of selected leaves

    Note: Note that multiple instances of each option can be deployed; options may be combined in any manner, and the target leaf or leaves are selected independently for each deployment. For example, two routed subnets can be defined, with one existing only on a subset of leaves. Less-preferred DAG options should be deployed only on the leaves where they are required.

- BGP information – A new BGP AS number is required for the fabric and is specified during the fabric workflow; a private AS may be used if needed. The workflow also collects handoff device BGP details, including the remote AS number and any MD5 authentication strings to automate border eBGP configuration.

NNJ 204 planning sheets

The following section documents the pre-planning details for the CV 204 NNJ network.

IP address, VLAN, and DHCP information

Subnet	Description	Fabric Leaf/Leaves	VLAN	DHCP Server Location
10.204.0.0/16	**Assigned range for this location / network**			
10.204.192.0/18	**Underlay and any traditional access layer subnets**			
10.204.253.0/24	Initial DHCP pool for default underlay (DHCP Server on Handoff-01)			Handoff-01
10.204.250.0/24	Manually configured /31s between spines and leaves/borders and /30s overlay Borders Handoffs			
10.204.255.0/24	Underlay Subnet for Fabric workflow			
10.203.251.0/25	leaf-01 access layer UAC Management		VLAN 2	leaf-01
10.203.251.128/25	leaf-02 access layer UAC Management		VLAN 2	leaf-02
10.204.252.0/25	leaf-03 access layer UAC Management		VLAN 2	leaf-03
10.204.252.128/25	leaf-04 access layer UAC Managemen		VLAN 2	leaf-04
10.204.211.0/24	leaf-01 traditional user subnet		VLAN 11	Corp Server 10.100.0.5 via Underlay
10.204.212.0/24	leaf-02 traditional user subnet		VLAN 12	Corp Server 10.100.0.5 via Underlay
10.204.213.0/24	leaf-03 traditional user subnet		VLAN 13	Corp Server 10.100.0.5 via Underlay
10.204.214.0/24	leaf-04 traditional user subnet		VLAN 14	Corp Server 10.100.0.5 via Underlay
10.204.0.0/18	**Overlay Range**			
10.204.11.0/24	leaf-01 routed fabric subnet - no DAG	leaf-01	VLAN 200	Corp Server 10.100.0.5 with VRF via Underlay
10.204.12.0/24	leaf-02 routed fabric subnet - no DAG	leaf-02	VLAN 200	Corp Server 10.100.0.5 with VRF via Underlay
10.204.13.0/24	leaf-03 routed fabric subnet - no DAG	leaf-03	VLAN 200	Corp Server 10.100.0.5 with VRF via Underlay
10.204.14.0/24	leaf-04 routed fabric subnet - no DAG	leaf-04	VLAN 200	Corp Server 10.100.0.5 with VRF via Underlay
10.204.21.0/24	1st routed DAG fabric subnet	leaf-01, leaf-02	VLAN 221	Corp Server 10.100.0.5 with VRF via Underlay
10.204.21.0/24	2nd routed DAG fabric subnet	leaf-02, leaf-03,leaf-04	VLAN 222	Corp Server 10.100.0.5 with VRF via Underlay
10.204.31.0/24	1st Bridged DAG fabric subnet	leaf-02, leaf-03,leaf-04	VLAN 231	Corp Server 10.100.0.5 with VRF via Underlay

BGP and OSPF routing information

OSPF (Underlay)

Handoff Description	Border Interface	Handoff Interface	VLAN SVI	Subnet	Trunk Interface	OSPF Area	MTU
Border-01 to Handoff-01	10.204.250.10	10.204.250.11	VLAN 3	10.204.250.10/31	Border g1/0/14 to Handoff g1/0/14	1	9100
Border-01 to Handoff-02	10.204.250.14	10.204.250.15	VLAN 5	10.204.250.14/31	Border g1/0/13 to Handoff g1/0/14	1	9100
Border-02 to Handoff-01	10.204.250.12	10.204.250.13	VLAN 2	10.204.250.12/31	Border g1/0/14 to Handoff g1/0/13	1	9100
Border-02 to Handoff-02	10.204.250.16	10.204.250.17	VLAN 4	10.204.250.16/31	Border g1/0/13 to Handoff g1/0/13	1	9100

OSPF Area Number	Name	Type
0	backbone	normal
1	handoffs	normal
11	leaf1	stub
12	leaf2	stub
13	leaf3	stub
14	leaf4	stub

BGP (Overlay)

Routing Domain	BGP Autonomous System Number / ASN
External Handoff Fabric VRF	65207
NNJ204 Fabric overlay	65208

Handoff Description	Border Interface	Handoff Interface	VLAN SVI	Subnet	Trunk Interface	MD5 String	MTU
Border-01 to Handoff-01	10.204.0.14	10.204.0.13	VLAN 7	100.204.0.12/30	Border g1/0/14 to Handoff g1/0/14	C1sco12345	9100
Border-01 to Handoff-02	10.204.0.18	10.204.0.17	VLAN 9	100.204.0.16/30	Border g1/0/13 to Handoff g1/0/14	C1sco12345	9100
Border-02 to Handoff-01	10.204.0.6	10.204.0.5	VLAN 6	100.204.0.4/30	Border g1/0/14 to Handoff g1/0/13	C1sco12345	9100
Border-02 to Handoff-02	10.204.0.10	10.204.0.9	VLAN 8	100.204.0.8/30	Border g1/0/13 to Handoff g1/0/13	C1sco12345	9100

RPVST+ Spanning Tree priorities

Fabric Device	Bridge Priority	VLAN List	Notes
spine1	4096	1-4094	Just for underlay conversion and RMA or new leaf add
spine2	8192	1-4094	Just for underlay conversion and RMA or new leaf add
leaf1	0	2,11,200,221	should match the truk allow list betwene leaf and access devices
leaf2	0	2,12,200,221-222,231	should match the truk allow list betwene leaf and access devices
leaf3	0	2,13,200,222,231	should match the truk allow list betwene leaf and access devices
leaf4	0	2,14,200,222,231	should match the truk allow list betwene leaf and access devices

Routed underlay

The validated setup is a greenfield deployment built on top of a manually configured Layer 3 underlay. It is important to have the underlay network complete with stable routing before attempting the fabric workflow. This section outlines building the CV underlay by following recommended best practices.

Initial topology

Initial setup includes all devices be in the default configuration:

- All interfaces are trunks with a native VLAN of 1 and an allowlist of 1-1000.
- All switches run MST.
- The upstream handoff devices are fully configured and out of scope in this document (their final configs are available in the appendices).
- A temporary initial DHCP pool is configured in VLAN 1 on Handoff-01.
- All devices have UAC control plan management connectivity to the cloud and have been added to a network called CV-204-NNJ.
- All switches are running IOS-XE 17.8.12.

Figure 1. **Initial Layer 2 topology**

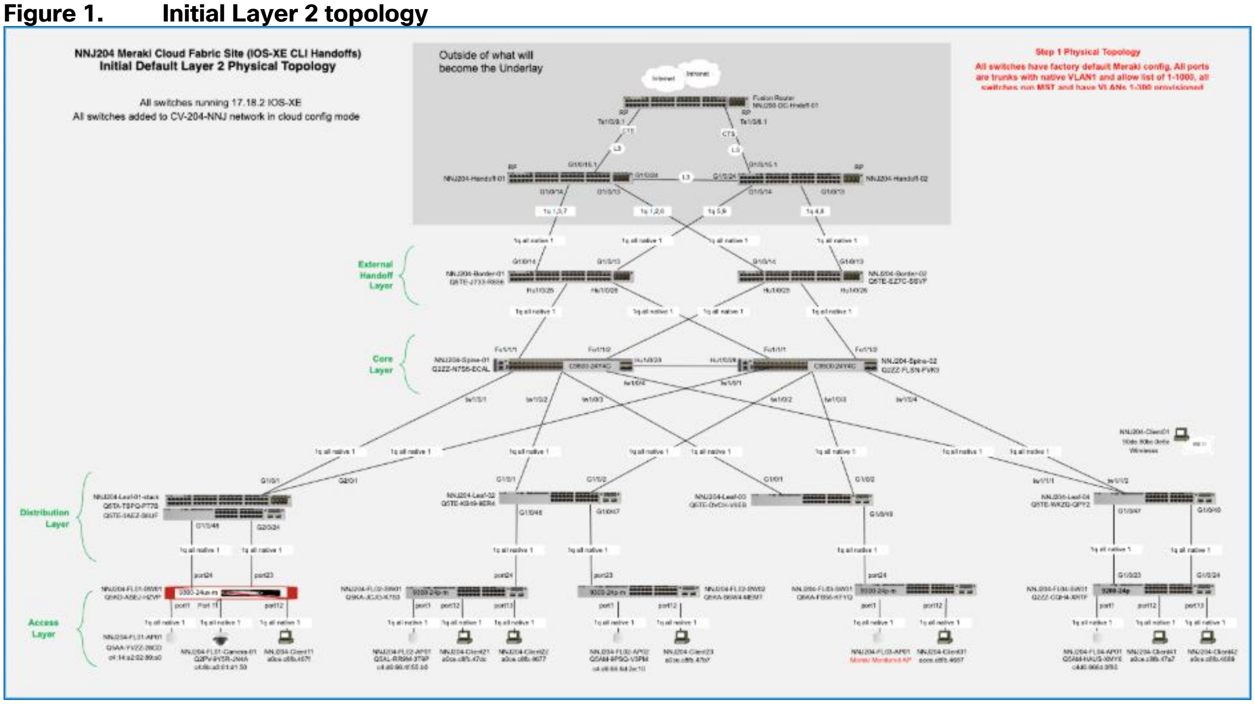

Figure 2. **Initial Layer 3 topology**

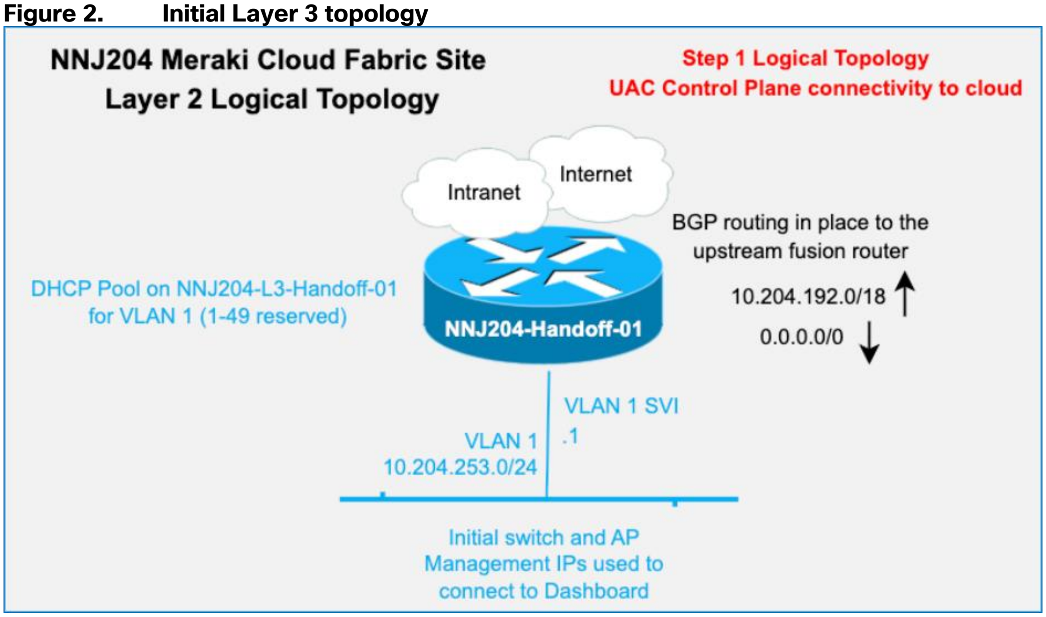

Layer 3 underlay conversion process

Care must be taken when converting the fabric links from Layer 2 trunks to Layer 3 routed interfaces. The process assumes best practice topology is in place such that each spine is connected to each leaf and to each border, and that each border is connected to each handoff, and the spines are connected to each other. This provides fully redundant physical paths. Using an outside-in approach, begin with the handoff-to-border links and convert one side of each redundant link to Layer 3 with OSPF routing.

Note: The first link being converted to Layer 3 on a given switch will automatically become the new UAC link. As links are converted from the borders to the spines and then spines to leaves, OSPF routing must provide reachability to the Internet. Once all fabric switches are using Layer 3 and OSPF routing for their UAC connectivity, the other half of the redundant links can be converted to complete the Layer 3 underlay.

The specific steps used in the validated network are as follows:

1. Set the network-wide MTU value for all switches to 9100.

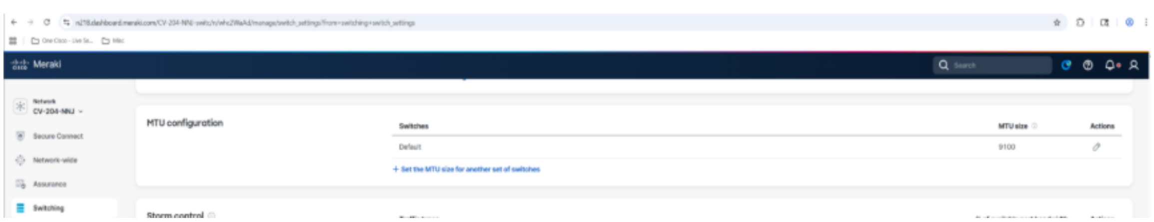

2. Adjust the Spanning Tree settings according to current planning.
   a. Before changing the settings, it is critical to shut down all ports on the leaves except for their uplinks to the spines.
   b. It is also important to then adjust the trunk allowlist on the remaining active links on all fabric devices to only include the required VLANs.

   VLAN 1 is needed for temporary management. In the validation setup underlay VLANs 2-5 were required for the border to handoff trunks. The goal is to reduce the required instances of STP to the minimum required before switching to RPVST+ to avoid traffic disruptions.

3. Set the network-wide OSPF settings according to current planning.

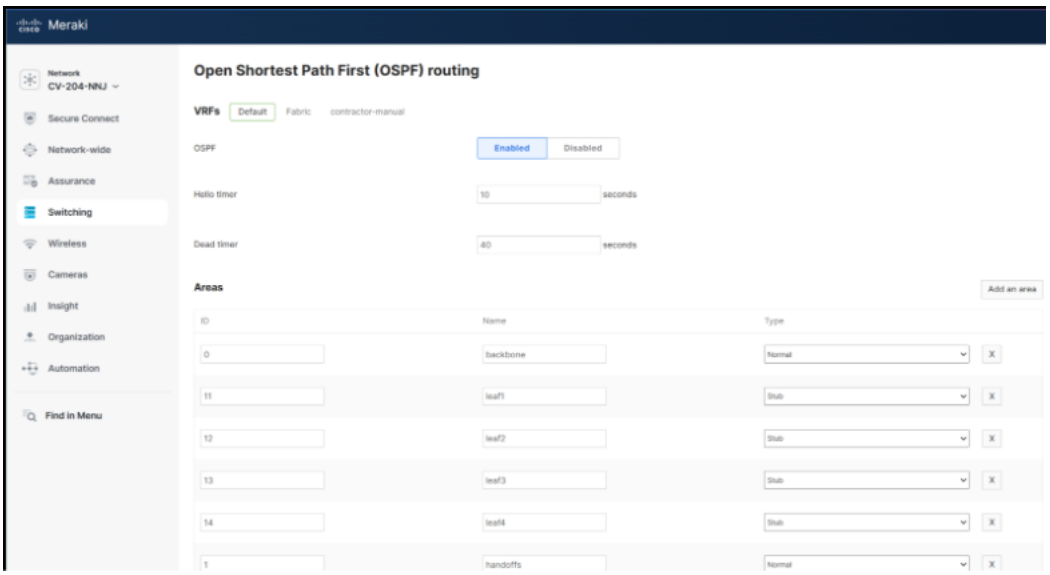

4. Identify the current uplink port on a switch to allow converting the other port. The initial temporary UAC should be in VLAN 1. There should be two uplink trunks on each switch carrying VLAN 1. Ensure that the link not currently used for UAC is chosen for modification. Use the `show uac uplink` command using the Cloud CLI capability to quickly determine which link is currently used for UAC.

Note:	Since the bridge priority of spine1 was set to 4096, the root port for VLAN 1 should be very predictable and that should align with the Port Used: field in the UAC display.

```
NNJ203-Leaf-01>show uac uplink
Uplink Autoconfig: Enable
Uplink Allow-list enforce: IPv4:No  IPv6:No
Configured IPv4 Uplink interface: Vlan 1 (Default)
Uplink IPv4 interface: Vlan 1
        IP Address: 10.203.253.18/255.255.255.0
        Type:       DHCP
        SVI:        Configured
        Port Used:  GigabitEthernet1/0/13
        GW IP:      10.203.253.1
        GW MAC:     cc03.d9ff.5d81
        Score:      7
Configured IPv6 Uplink interface: Vlan 1 (Default)
Uplink IPv6 interface: None
Uplink Reachable: IPv4
```

```
NNJ203-Leaf-01>show int trunk

Port        Mode             Encapsulation  Status        Native vlan
Gi1/0/13    on               802.1q         trunking      1
Gi1/0/14    on               802.1q         trunking      1

Port        Vlans allowed on trunk
Gi1/0/13    1
Gi1/0/14    1

Port        Vlans allowed and active in management domain
Gi1/0/13    1
Gi1/0/14    1

Port        Vlans in spanning tree forwarding state and not pruned
Gi1/0/13    1
Gi1/0/14    none
```

```
NNJ203-Leaf-01>show span

VLAN0001
  Spanning tree enabled protocol rstp
  Root ID    Priority    4097
             Address     24d5.e41d.6300
             Cost        20000
             Port        13 (GigabitEthernet1/0/13)
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    32769  (priority 32768 sys-id-ext 1)
             Address     246c.847a.fc80
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time  300 sec

Interface           Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- ------------------------
Gi1/0/13            Root FWD 20000     128.13   P2p
Gi1/0/14            Altn BLK 20000     128.14   P2p
```

5. Convert the "alternate" port that is not currently the UAC uplink to Layer 3. Note that when the change is saved, it will automatically navigate the user to the Layer 3 settings for this port.

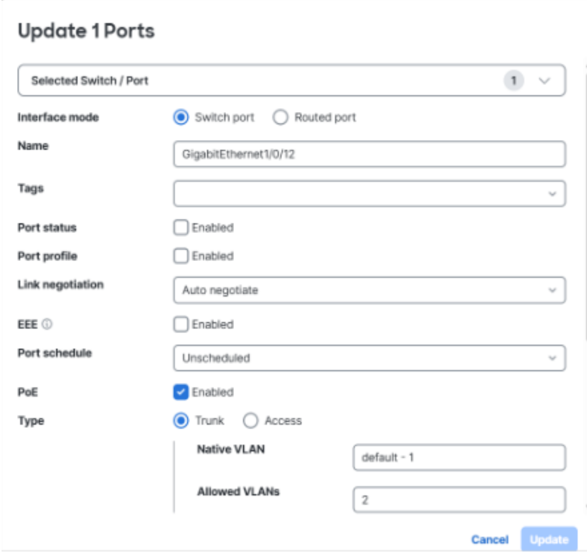

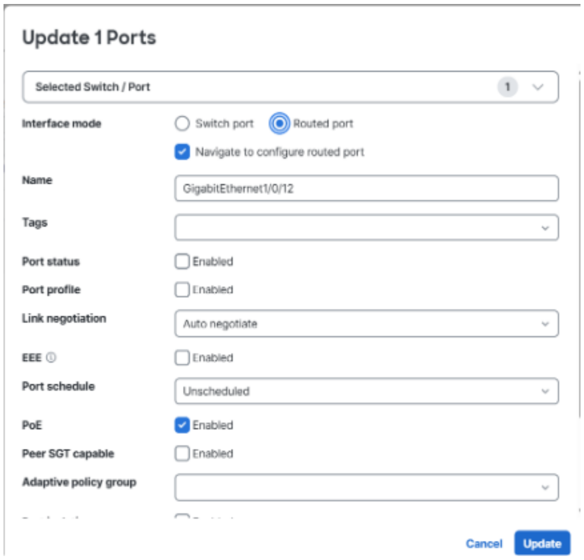

Since this is the first Layer 3 interface on this switch, the **Preferred Uplink > IPv4 Preferred Management Connectivity** setting is automatically selected. This configuration cannot be saved unless the option is selected. If another interface is later converted to Layer 3, the management function can be moved. A static route is automatically created according to the specified next-hop default gateway.

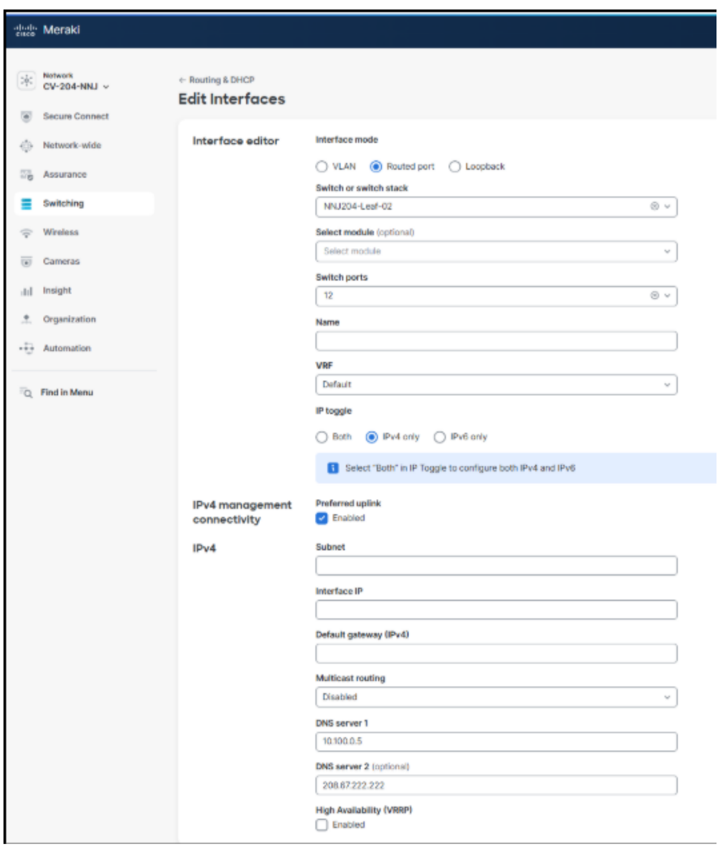

There are also options for DHCP and OSPF. Because OSPF is required for the Layer 3 management routing, OSPF must be enabled and the interface placed in the desired area defined previously. P2P is recommended to avoid unnecessary DR/BFR election. DHCP is typically off for fabric links. It may be temporarily enabled for day-two RMA or new leaf add activities.

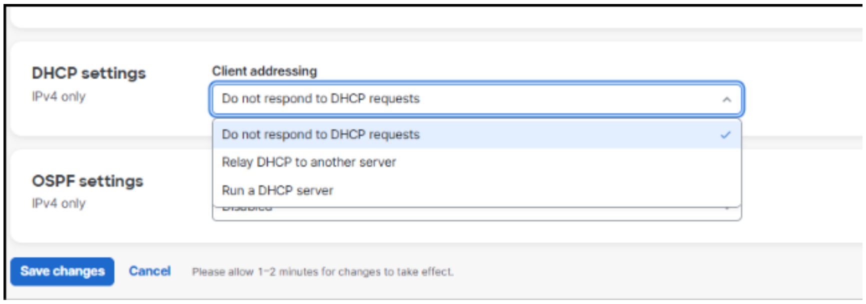

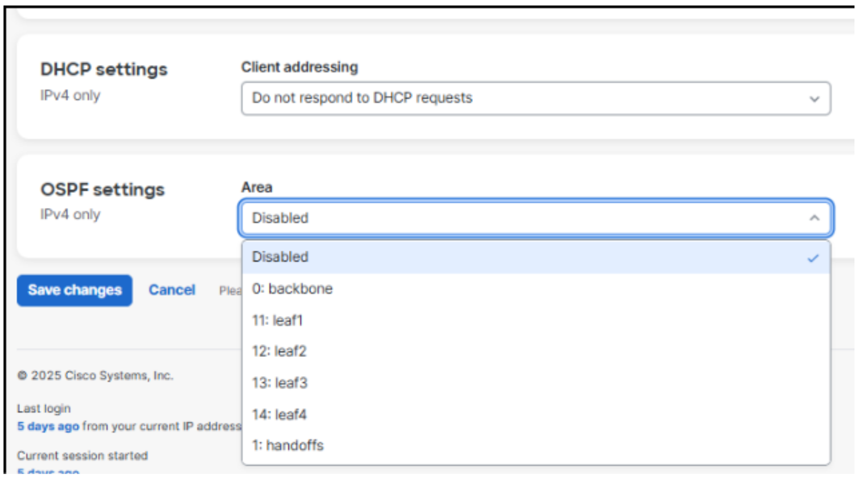

The resulting Layer 3 interface and static route definitions can be found by navigating to **Switching > Configure > Routing and DHCP** in the UI.

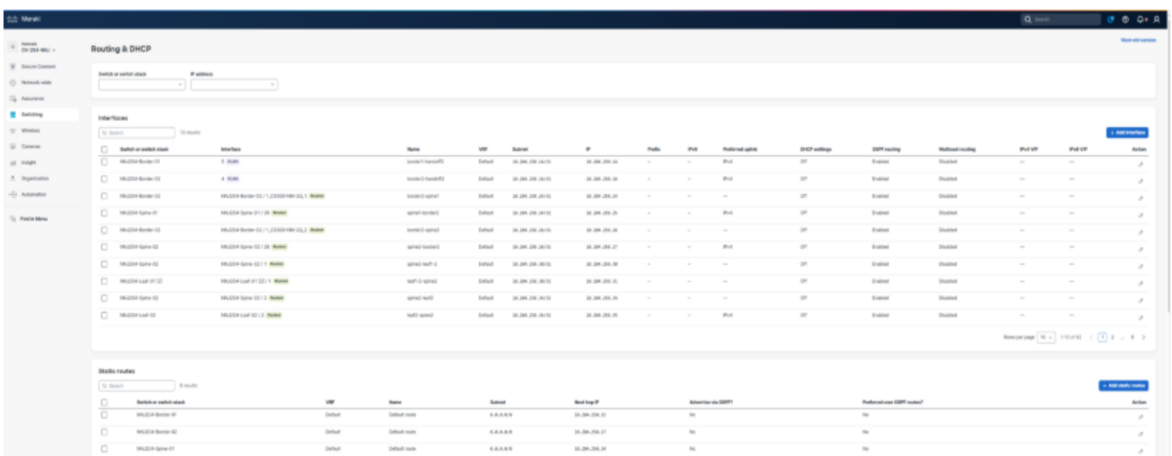

> **Note:** The default behavior for the static route that is automatically created, is to have the available additional OSPF settings both set to **No**. This is important, aligned to best practice, and will result in a static route being added to the configuration with an administrative distance (AD) of 120 and not redistributed into OSPF. Since OSPF uses an AD of 110, the OSPF learned default route advertised from the handoff devices into the borders will be preferred in the active routing table. The static route is essentially a floating static backup route.

When the Layer 3 underlay conversion is complete, the underlay routing table will have discrete OSPF learned or connected routes for all the fabric links and the default for traffic outside the underlay via the borders. Access layer management subnets, and any traditional non-fabric subnets will also reside in the underlay routing table, which is the global routing table.

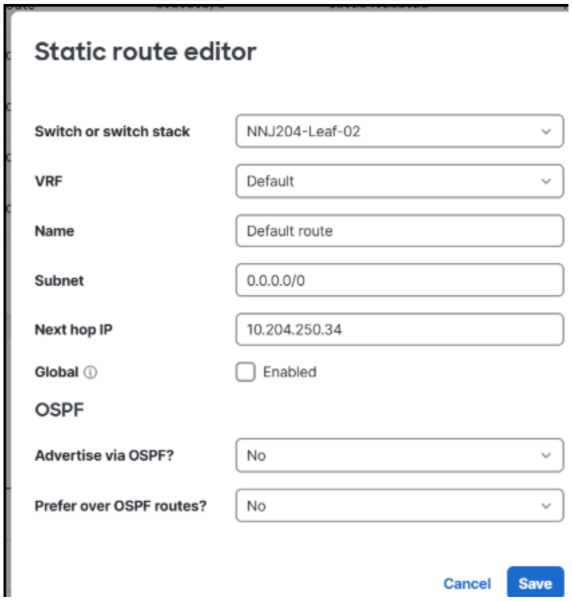


```
Welcome to the interactive CLI IOS XE terminal
You are in Read-only Mode

Establishing connection to your device. Please wait...
Connection established successfully

NNJ204-Leaf-02>show run | inc ip route
ip route 0.0.0.0 0.0.0.0 10.204.250.34 120
```

```
NNJ204-Leaf-02>show ip route
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, m - OMP
       n - NAT, Ni - NAT inside, No - NAT outside, Nd - NAT DIA
       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
       H - NHRP, G - NHRP registered, g - NHRP registration summary
       o - ODR, P - periodic downloaded static route, l - LISP
       a - application route
       + - replicated route, % - next hop override, p - overrides from PfR
       & - replicated local route overrides by connected

Gateway of last resort is 10.204.250.34 to network 0.0.0.0

O*E1  0.0.0.0/0 [110/13] via 10.204.250.34, 23:56:04, GigabitEthernet1/0/2
                [110/13] via 10.204.250.32, 23:56:04, GigabitEthernet1/0/1
```

Note: When an interface is converted from a routed port back to a switched port, any associated Layer 3 configuration is automatically removed.

6. Complete the conversion of all underlay links to Layer 3 and confirm. On each fabric device, verify the correct number of OSPF neighbors, IP routing including redundant default routes, and minimal spanning trees. There should be no STPs on the spines.

   At this stage, the DHCP scope for VLAN 1 on Handoff-01 can be removed but is often retained to support a day-two RMAs or factory resets. If a spine is replaced or reset, a temporary DHCP pool must be created on a border, and one border interface temporarily converted back to Layer 2 to support the onboarding of the RMAed spine. A similar process is required for a leaf with the DHCP and the temporary Layer 2 link living on the spine and is the most common scenario when adding new leaves. In all cases, the objective is to use a temporary VLAN 1 for initial cloud connectivity and then convert to Layer 3 connectivity.

7. Once the routed underlay is stable, the leaf-to-access layer trunks can be enabled. These trunk links carry the UAC management VLAN for all downstream cloud-managed devices. VLAN 2 was used for this in the validation setup. Before activating the ports, clear the VLAN allowlist to include just the management of VLAN and any other traditional subnets if present. The access layer devices will run MST initially until they connect to the cloud and are configured to run RPVST+. These devices require DHCP, and these DHCP scopes will remain in place.
    a. Add Layer 3 SVI for the access layer VLANs on each leaf and set the DHCP settings as required.
    b. Remember these VLANS are unique broadcast domains and not connected between leaves allowing for the same VLAN numbers to be used if desired.
    c. Unique IP subnets are required.
    d. Remember to enable OSPF on the SVIs in the desired non-zero OSPF Area. Set them to OSPF Passive as there will be no downstream OSPF neighbors in the access layer.
8. It is best practice to enable automatic fallback to preferred uplink in the global switch settings and set the preferred uplink VLAN on the access switches. This ensures the UAC traffic remains in the underlay as intended.
9. Set up EtherChannels for increased throughput and redundancy.
10. Set up access ports and enable dot1x for desired access devices.
11. Enable CTS in-line tagging using the **Peer SGT capable** setting on both sides of the leaf ports facing the access layer devices and border ports facing the handoff devices. Any APs attached to access switch ports must also have Peer SGT enabled. Always select **Adaptive policy group 2: Infrastructure** when enabling the CMD header for infrastructure-to-infrastructure links.

Ending underlay topology

When complete underlay connectivity is in place, all devices are registered in the UI and stable; the setup is ready for the fabric workflow to create the desired fabrics and related configurations.

Figure 3. Underlay Layer 2 topology

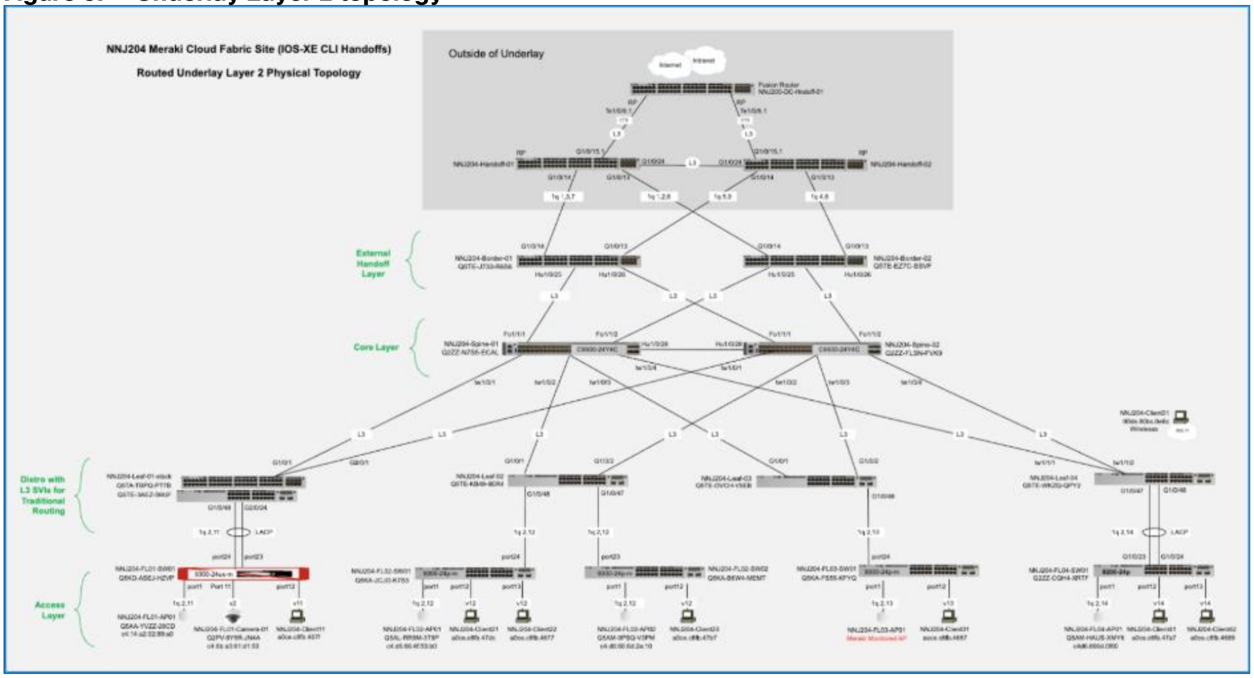

Figure 4. Underlay Layer 3 topology

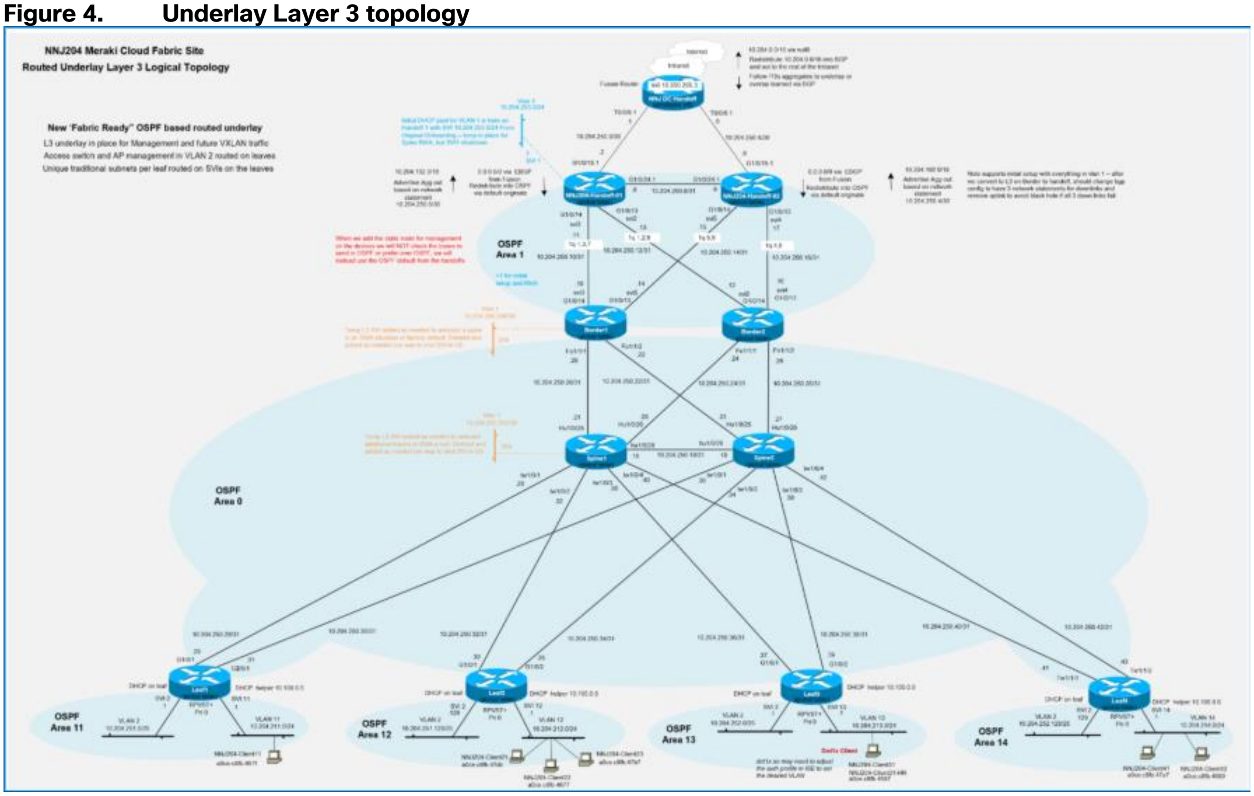

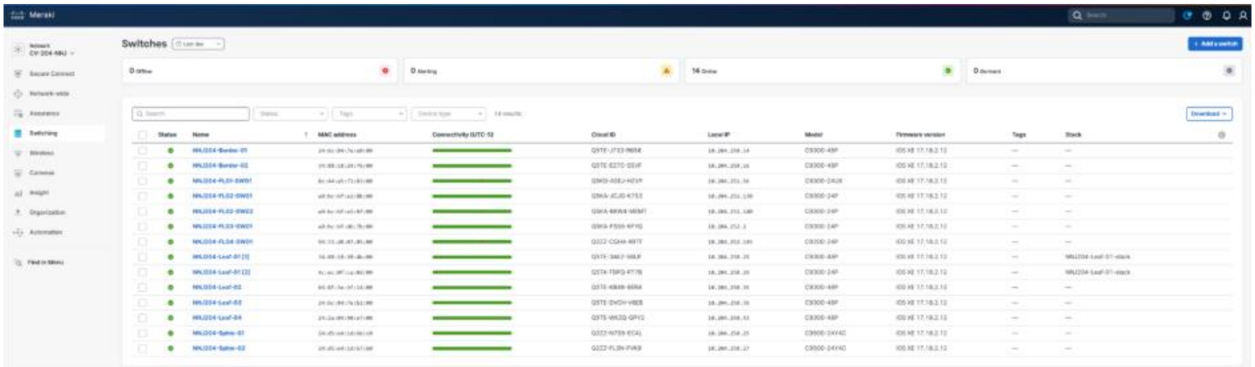

At this point, the underlay is ready for overlay provisioning via the fabric workflow.

Fabric overlay setup

Fabric setup

Once the underlay is prepared and stable, the fabric workflow is used to create and maintain the fully automated campus fabric. The first portion of the workflow collects the information required for the infrastructure including the name of the fabric, the fabric BGP AS, and the underlay type and subnet.

The best practice recommendation is to use a Layer 3 underlay, which is known as a **Custom** underlay in the UI. As the workflow progresses, VRF and eBGP handoff information are added along with one or more fabric subnets, which can be Routed, Routed DAG, and/or Bridged DAG subnets in any combination and leaf distribution.

Figure 5. **Fabric planning sheet for NNJ204**

Fabric Settings

Field	Value
Fabric Name	nnj204-fabric
Fabric BGP AS	65208
Selected Networks	CV-204-NNJ
Underlay Loopback IP Poo	10.204.255.0/24
Underlay core IP Pool	N/A as setup is using L3 Underlay
Custom Underlay	"Enabled"

Device Roles

Device	Fabric Role
NNJ204-Spine-01	Spine
NNJ204-Spine-02	Spine
NNJ204-Border-01	Border
NNJ204-Border-02	Border
NNJ204-Leaf-01	Leaf
NNJ204-Leaf-02	Leaf
NNJ204-Leaf-03	Leaf
NNJ204-Leaf-04	Leaf

Fabric Subnets

Subnet Name	VLAN name	Type	VLAN ID	SVI IP and Mask	DHCP Server(s)	VRF	Leaves	Anycast Gateway	Broadcast Replicatio
leaf1-routed	leaf1-routed	routed (no DAG)	200	10.204.11.1/24	10.100.0.5	Fabric	leaf-01	unchecked	uncheckec
leaf2-routed	leaf2-routed	routed (no DAG)	200	10.204.12.1/24	10.100.0.5	Fabric	leaf-02	unchecked	uncheckec
leaf3-routed	leaf3-routed	routed (no DAG)	200	10.204.13.1/24	10.100.0.5	Fabric	leaf-03	unchecked	uncheckec
leaf4-routed	leaf4-routed	routed (no DAG)	200	10.204.14.1/24	10.100.0.5	Fabric	leaf-04	unchecked	uncheckec
Routed-DAG-1	Routed-DAG-01	routed (no DAG)	221	10.204.21.1/24	10.100.0.5	Fabric	leaf-01, leaf-02	checked	uncheckec
Routed-DAG-2	Routed-DAG-02	routed (no DAG)	222	10.204.22.1/24	10.100.0.5	Fabric	leaf-02, leaf-03,leaf-04	chedked	uncheckec
Bridged-DAG-1	Routed-DAG-02	routed (no DAG)	231	10.204.23.1/24	10.100.0.5	Fabric	leaf-02, leaf-03,leaf-04	checked	checked

BGP Layer3 Connection Information

Switch	Name	VRF	VLAN	MTU	IP/Mask
Border-01	border1-handoff1	Fabric	7	9100	10.204.0.14/30
Border-01	border1-handoff2	Fabric	9	9100	10.204.0.18/30
Border-02	border2-handoff1	Fabric	6	9100	10.204.0.6/30
Border-02	border2-handoff2	Fabric	8	9100	10.204.0.10/30

BGP Peer Information

Neighbor IP	Remote AS	VRF	Source Int	MD5 String
Peers on Border-01				
10.204.0.5	65207	Fabric	10.204.0.6/30	C1sco12345
10.204.0.9	65207	Fabric	10.204.0.10/30	C1sco12345
Peers on Border-02				
10.204.0.13	65207	Fabric	10.204.0.14/30	C1sco12345
10.204.0.17	65207	Fabric	10.204.0.18/30	C1sco12345

Fabric workflow

The screen captures and related information below outlines the current fabric workflow.

1. Navigate to **Organization > Configure** in the cloud UI to view the fabric workflow. The initial fabric values are set along with the fabric device roles. One or more VRFs are added.

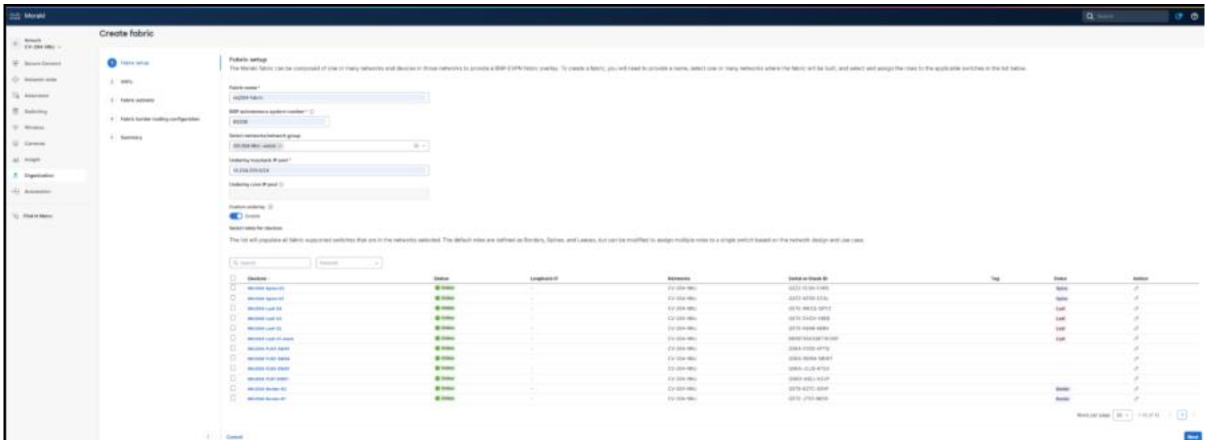

Note: There is one VRF called **Fabric** that will be created by default as part of the workflow.

Note: Other VRFs created outside the fabric workflow will display, but they are not eligible for use in the workflow.

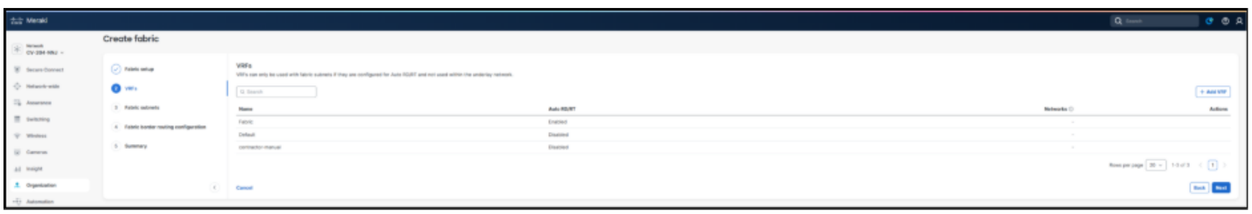

One or more fabric subnets are added, and the leaves where the subnet should be deployed are selected.

2. Select **Anycast Gateway** and/or **Broadcast Replication** to control whether the fabric subnet is Routed, Routed DAG, or Bridged DAG.

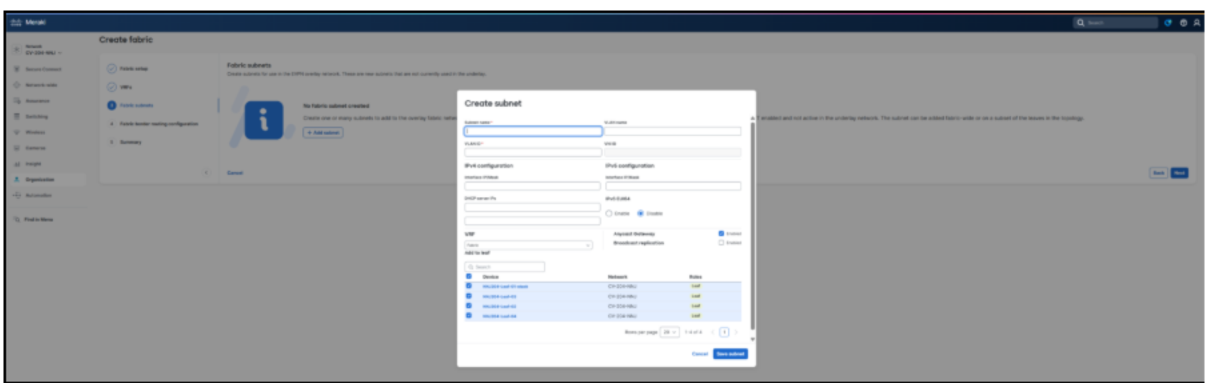

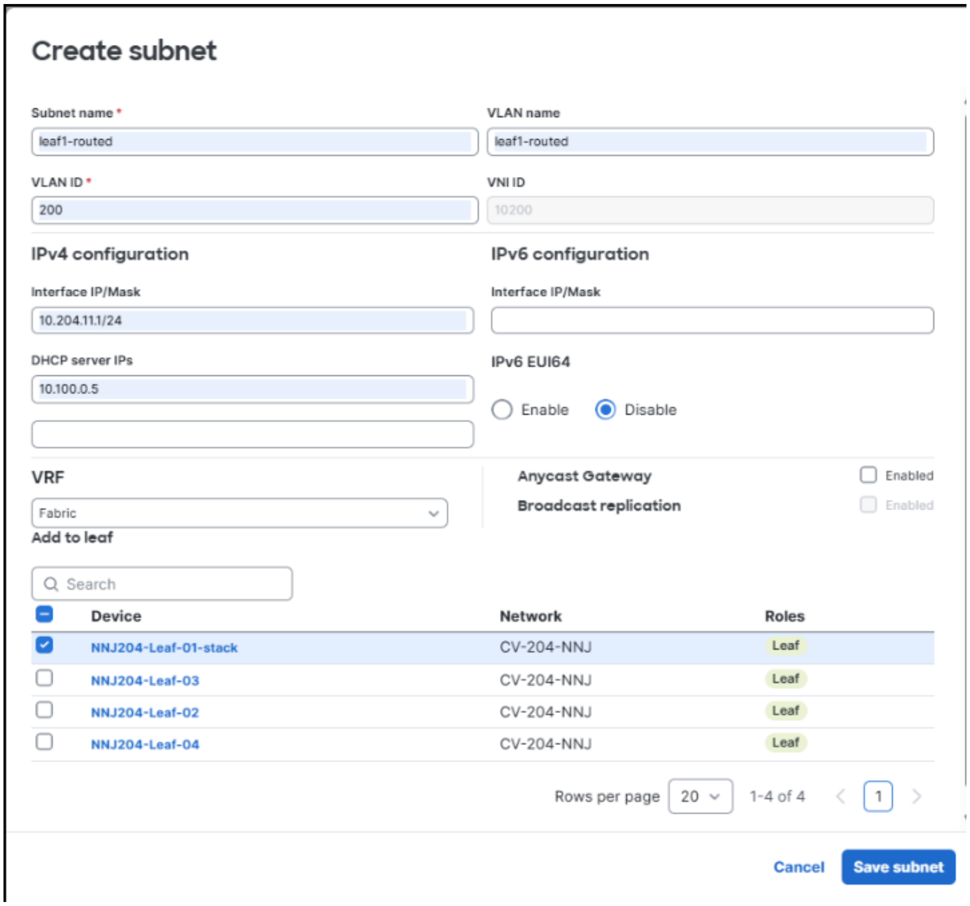

3. Select the **Configure L3 Interface** option to add the settings for each Layer 3 handoff link. The validation setup process is repeated four times; one for each handoff link in the fabric VRF.

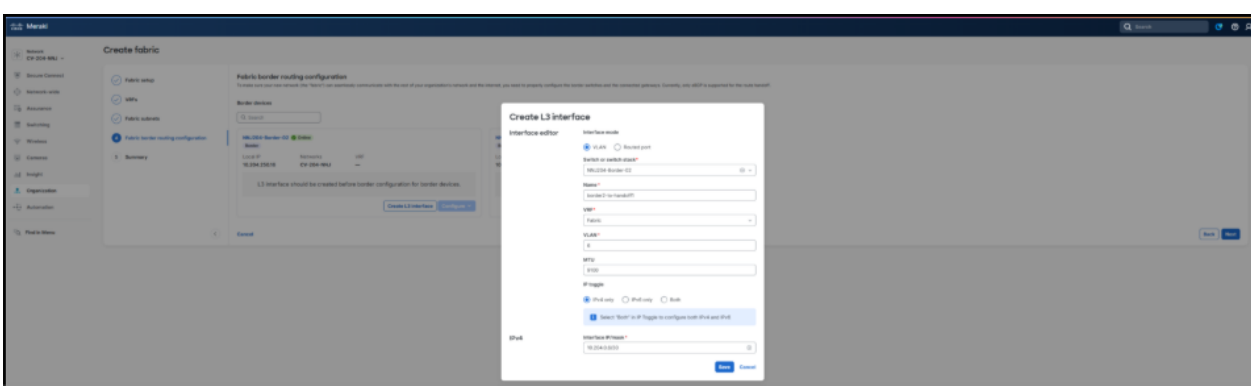

4. After adding the Layer 3 link information, click **Configure > Create eBGP Peer** to add the corresponding eBGP peering information for each link.

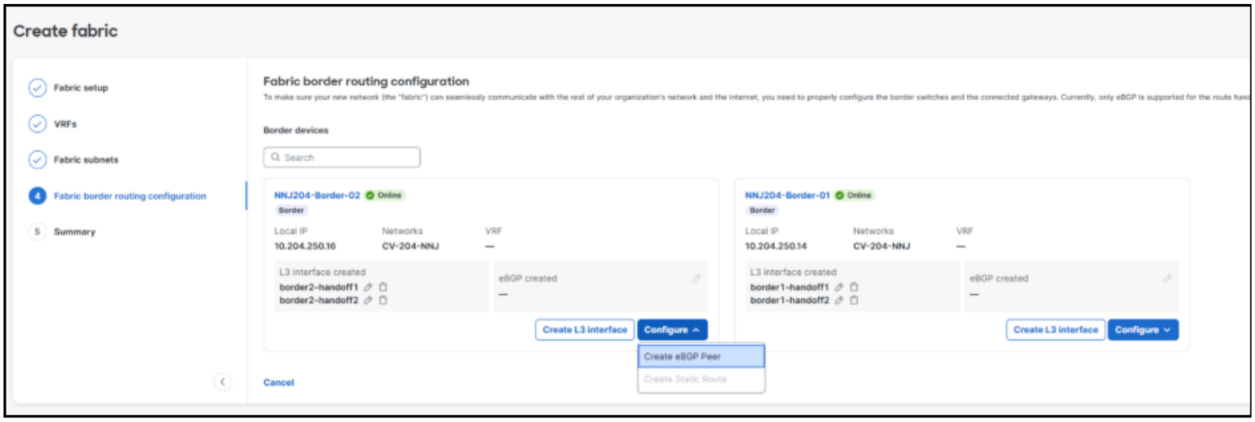

This is done uniquely on each border resulting in two peers on each border in the validated setup.

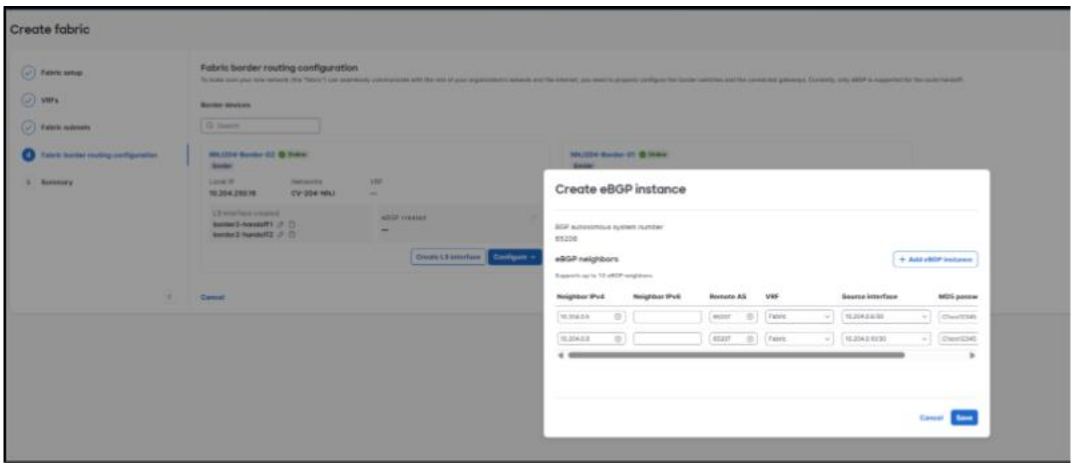

The screenshot below displays the completed border configuration screen. This information can be updated at any time if the border configurations require modification.

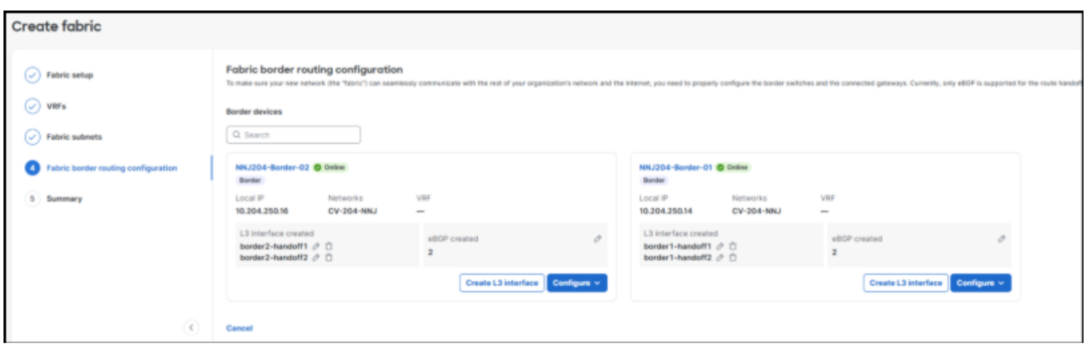

5. Click **Save to Staging** once the values are set as intended.

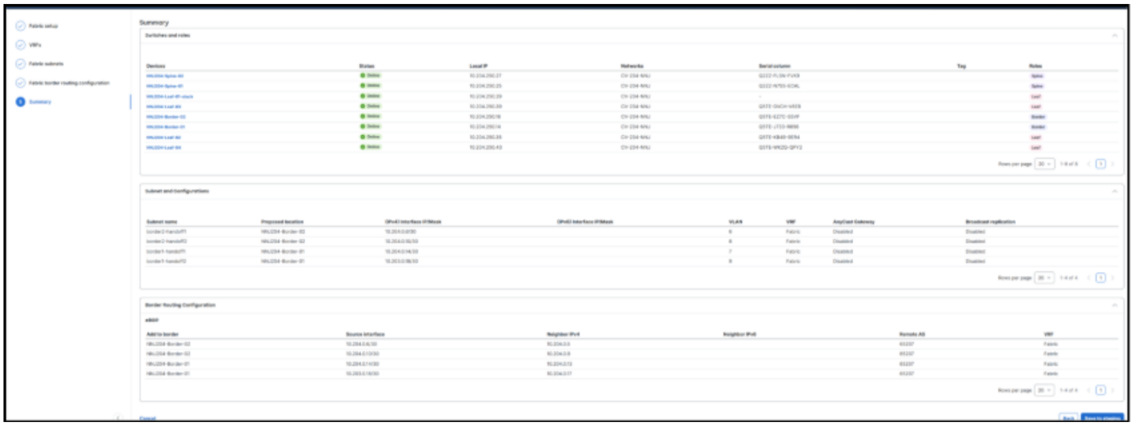

6. Once the information is saved to staging, select **Preview Changes** to display the high-level sections of configuration that will be deployed.

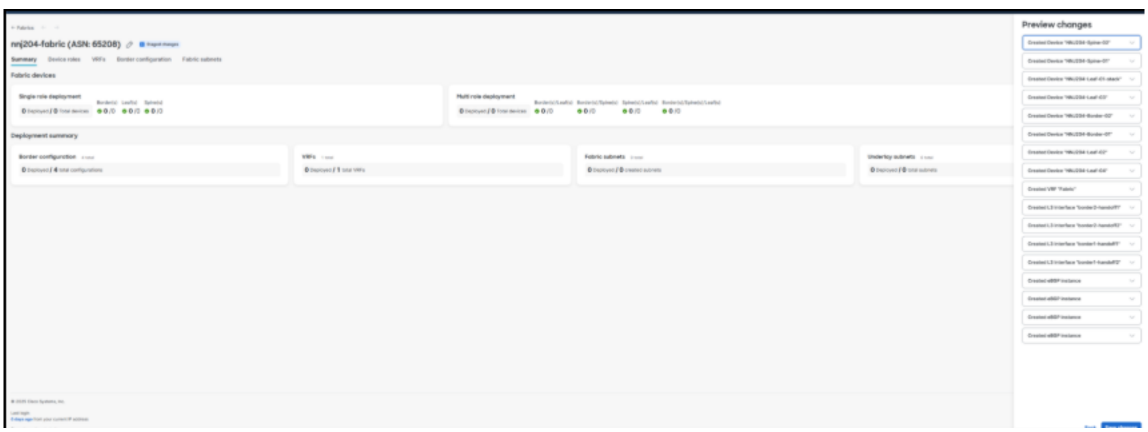

7. Click **Deploy** to start the deployment.

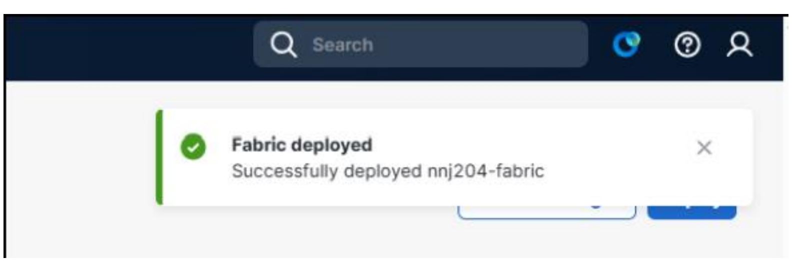

Within a few minutes the configurations will be delivered to the devices from the cloud, and the fabric subnets and related capabilities are ready for use.

The entire fabric can be deleted with a fully automated cleanup using the **Delete** option.

Final topology

Figure 6. Fabric Layer 2 physical topology

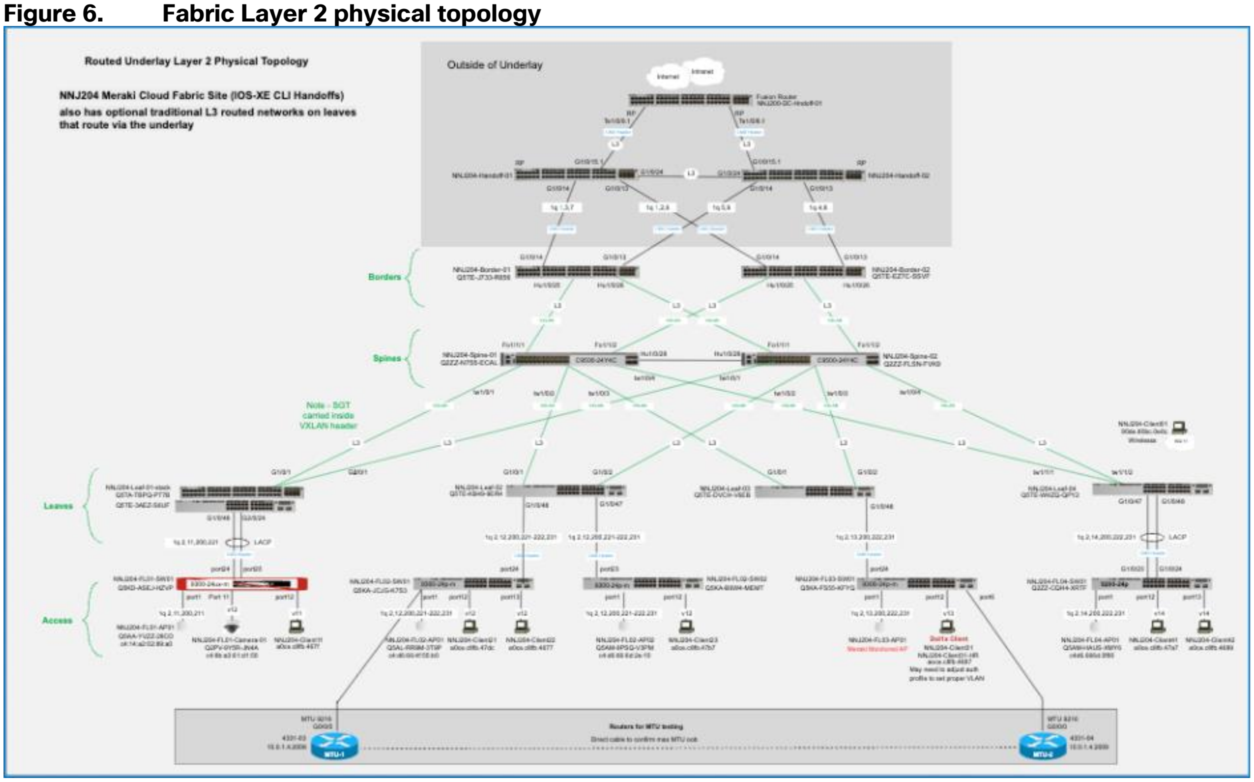

Figure 7. Fabric Layer 3 logical topology

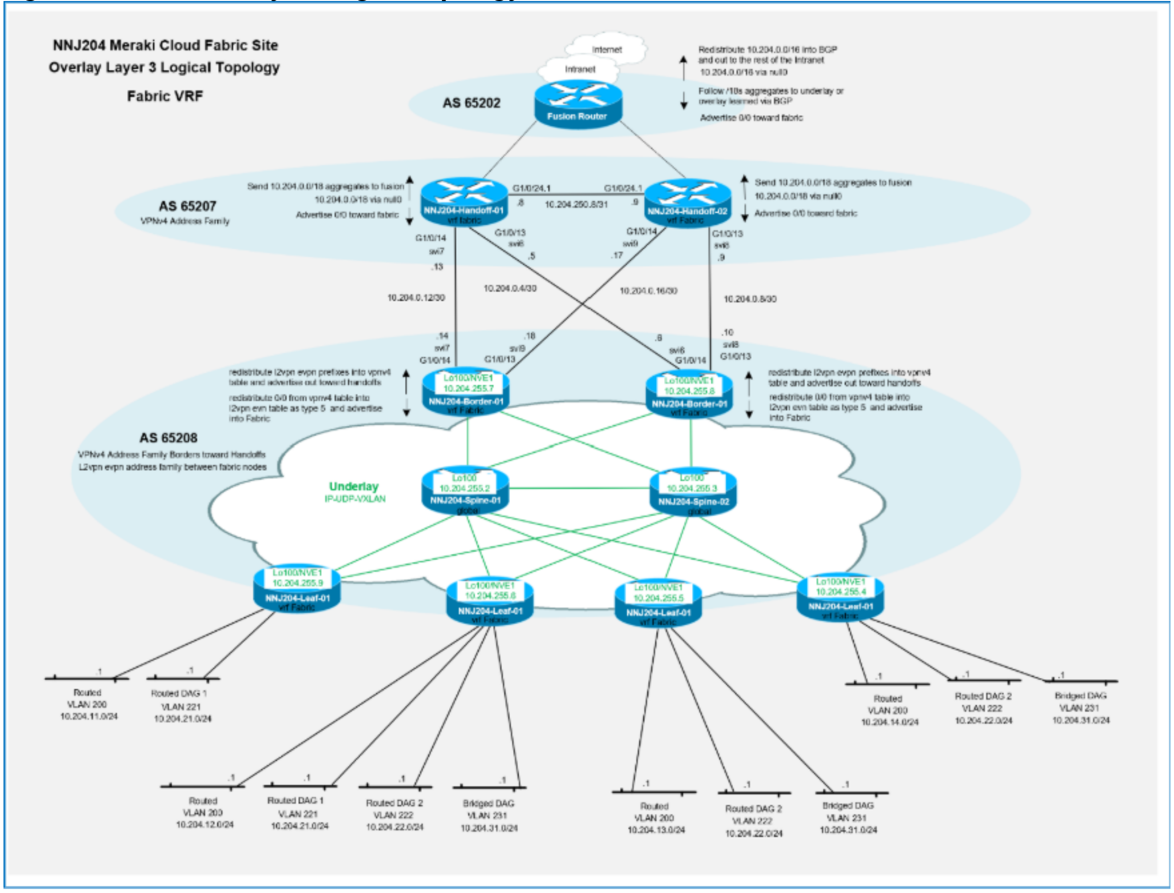

Multiple VRFs

It is possible to create multiple VRFs in a fabric. An additional VFR was created on two access switches in the validated setup to isolate guest traffic on wired guest ports.

VRF workflow

Figure 8. Planning sheet for adding NNJ204 guest access

Fabric Settings - additional guest VRF

Field	Value
Fabric Name	nnj204-guest

Fabric Subnets

Subnet Name	VLAN name	Type	VLAN ID	SVI IP and Masl	DHCP Server(s)	VRF	Leaves	Anycast Gateway	Broadcast Replication
leaf1-guest	leaf1-guest	routed (no DAG)	64	10.204.64.1/24	10.100.0.5	nnj204-guest	leaf-01	unchecked	unchecked
leaf2-guest	leaf2-guest	routed (no DAG)	65	10.204.65.1/24	10.100.0.5	nnj204-guest	leaf-02	unchecked	unchecked

BGP Layer3 Connection Information (border side has higher host IP address)

Switch	Name	VRF	VLAN	MTU	IP/Mask
Border-01	border1-handoff1-guest	nnj204-guest	64	9100	10.204.127.0/3
Border-02	border2-handoff2-guest	nnj204-guest	65	9100	10.204.127.4/3

BGP Peer Information

Neighbor IP	Remote AS	VRF	Source Int	MD5 String
Peers on Border-01				
10.204.127.1	65207	nnj204-guest	10.204.127.2/30	C1sco12345
Peers on Border-02				
10.204.127.5	65207	nnj204-guest	10.204.127.6/30	C1sco12345

1. A new VRF is added.

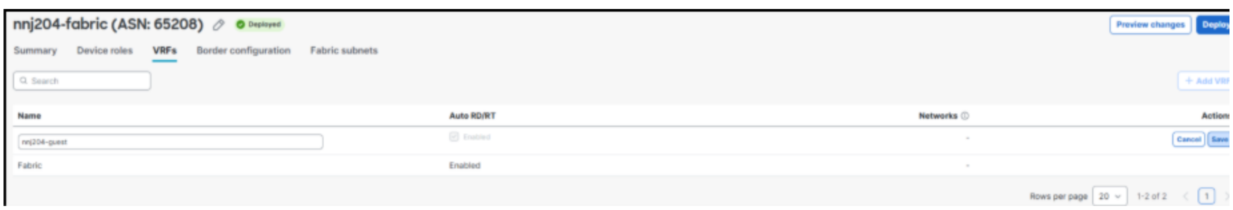

2. A new routed fabric subnet is created for Leaf 1.

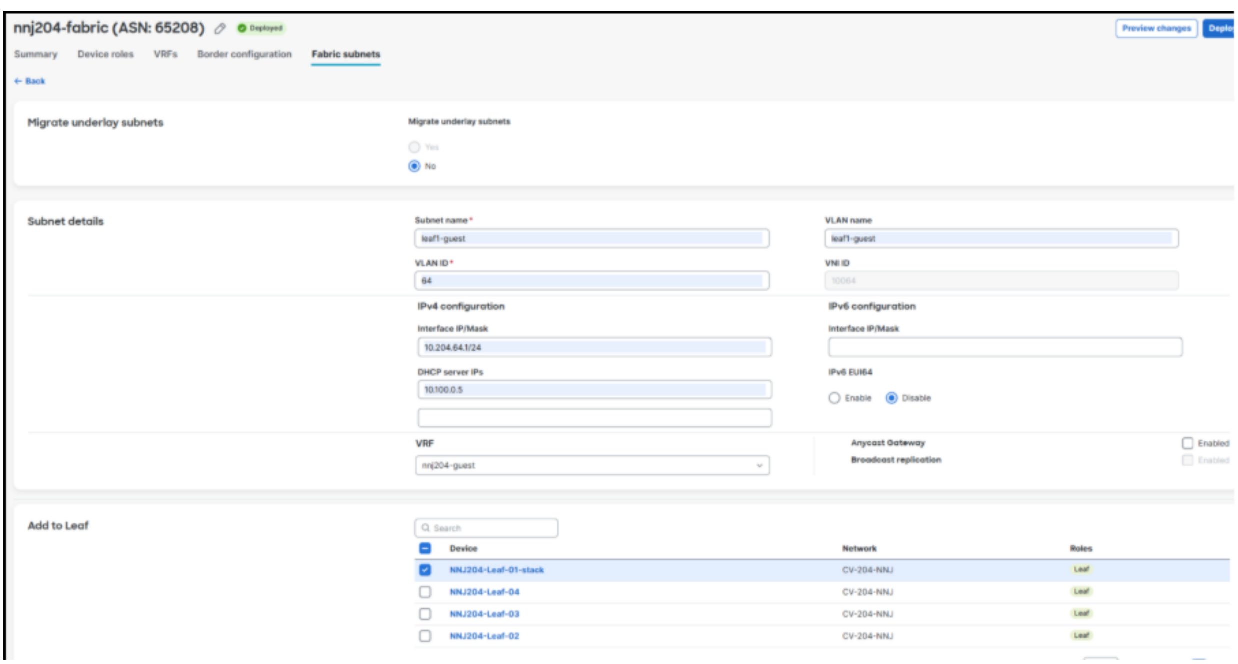

3. A new routed fabric subnet is created for Leaf 2.

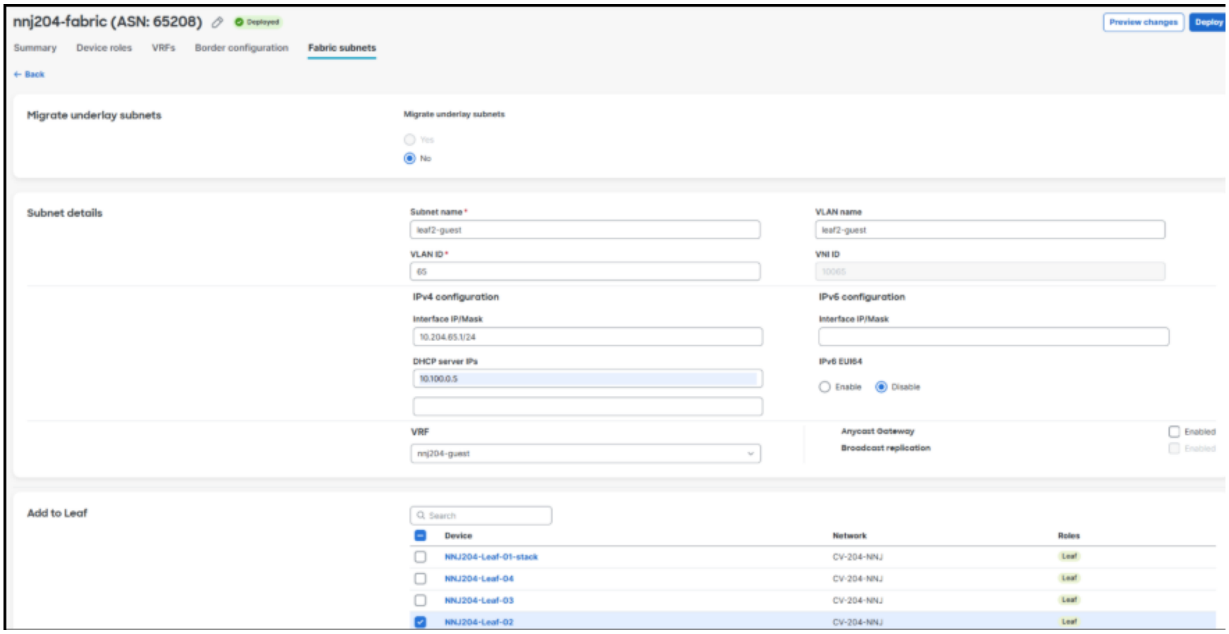

4. Connectivity from the new guest VRF fabric subnets to the network outside of the fabric is provisioned by adding two border interfaces and their related BGP peering information

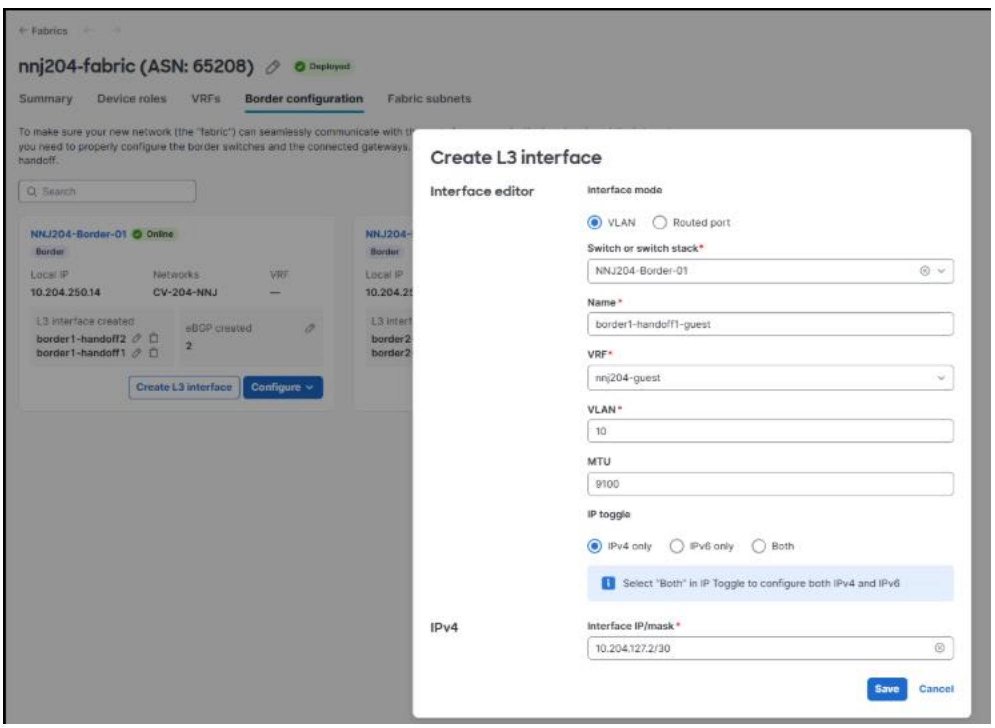

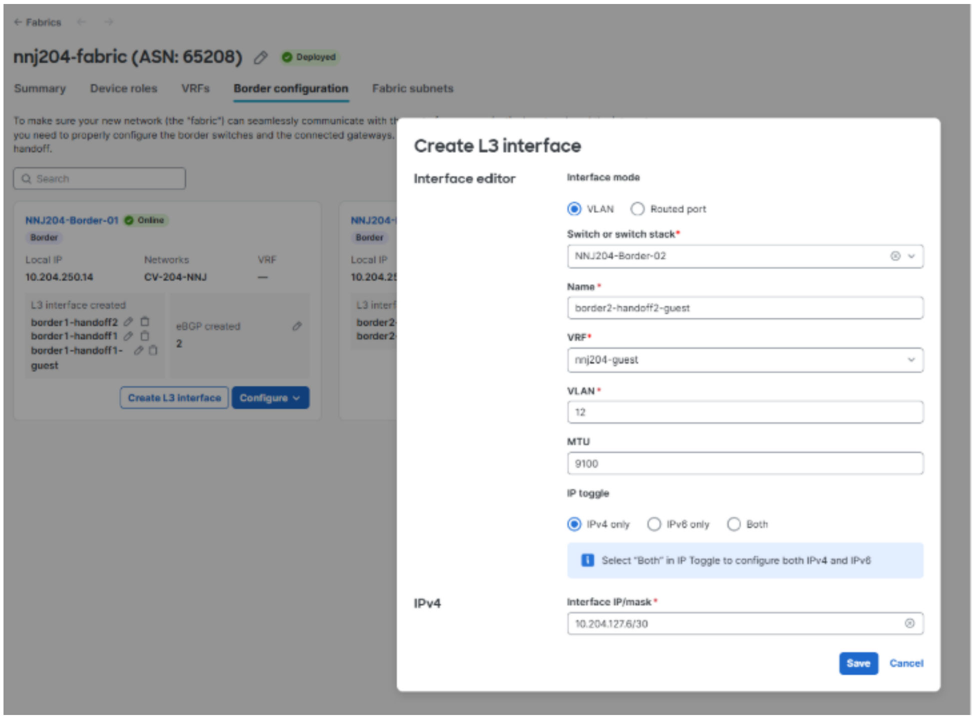

Fabrics
nnj204-fabric (ASN: 65208)
Deployed
Summary
Device roles
VRFs
Border configuration
Fabric subnets
To make sure your new network (the "fabric") can seamlessly communicate with th
you need to properly configure the border switches and the connected gateways.
handoff.
Search
NNJ204-Border-01
Online
Border
Local IP
Networks
VRF
10.204.250.14
CV-204-NNJ
L3 interface created
border1-handoff2
border1-handoff1
border1-handoff1-guest
eBGP created
2
Create L3 interface
Configure
Create L3 interface
Interface editor
Interface mode
VLAN
Routed port
Switch or switch stack*
NNJ204-Border-02
Name *
border2-handoff2-guest
VRF*
nnj204-guest
VLAN *
12
MTU
9100
IP toggle
IPv4 only
IPv6 only
Both
Select "Both" in IP Toggle to configure both IPv4 and IPv6
IPv4
Interface IP/mask *
10.204.127.6/30
Save
Cancel

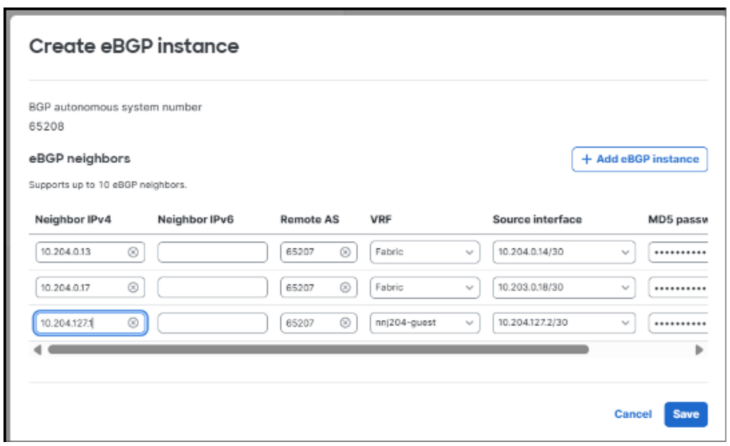

Create eBGP instance
BGP autonomous system number
65208
eBGP neighbors
+ Add eBGP instance
Supports up to 10 eBGP neighbors.
Neighbor IPv4
Neighbor IPv6
Remote AS
VRF
Source interface
MD5 passw
10.204.0.13
65207
Fabric
10.204.0.14/30
10.204.0.17
65207
Fabric
10.203.0.18/30
10.204.127.1
65207
nnj204-guest
10.204.127.2/30
Cancel
Save

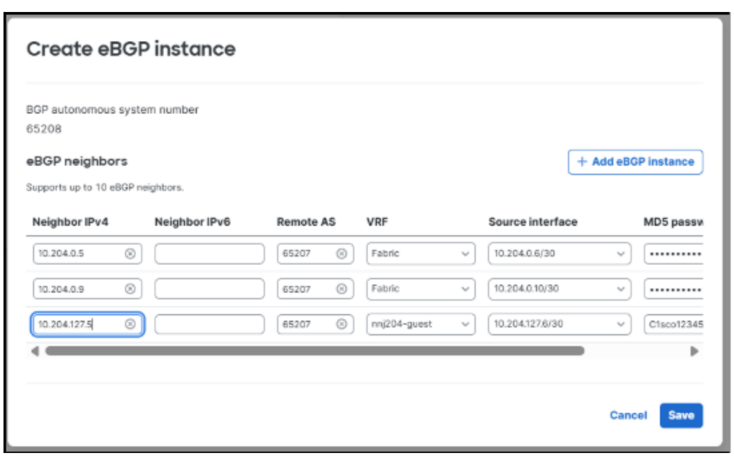

Create eBGP instance
BGP autonomous system number
65208
eBGP neighbors
+ Add eBGP instance
Supports up to 10 eBGP neighbors.
Neighbor IPv4
Neighbor IPv6
Remote AS
VRF
Source interface
MD5 passw
10.204.0.5
65207
Fabric
10.204.0.6/30
10.204.0.9
65207
Fabric
10.204.0.10/30
10.204.127.5
65207
nnj204-guest
10.204.127.6/30
C1sco12345
Cancel
Save

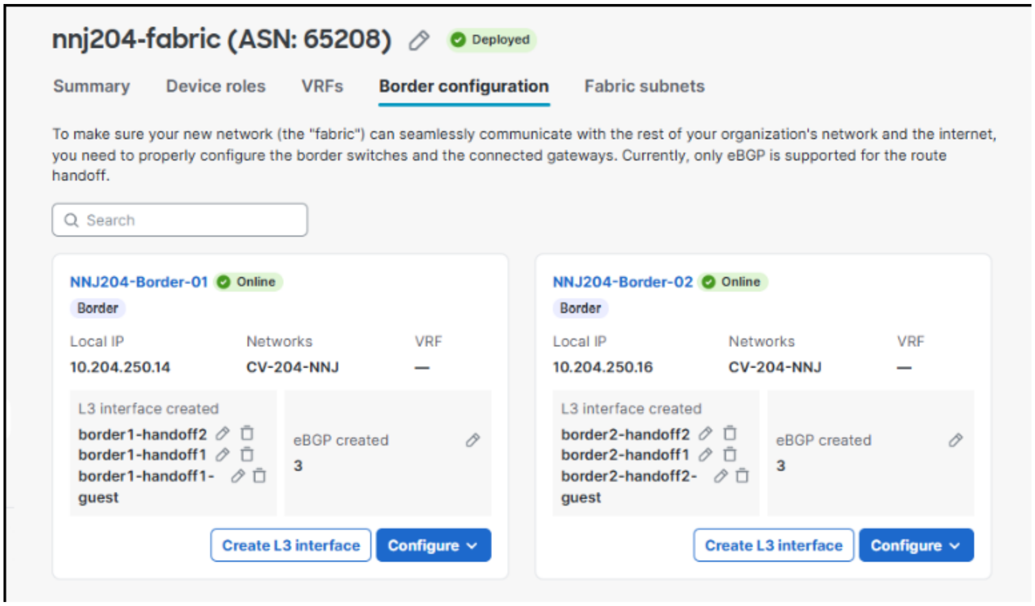

5. To complete the work, changes are reviewed, saved, and deployed. In several minutes, the new VRF connectivity will be ready for use.

   Wired guests can come into play in various deployment scenarios. The validation setup uses both statically set guest ports with the guest VLAN and SGT hard coded, and dynamic 802.1x authentication with the VLAN and SGT set from the ISE policy.

6. Static guest port set to guest VLAN 65 with guest SGT (6) set.

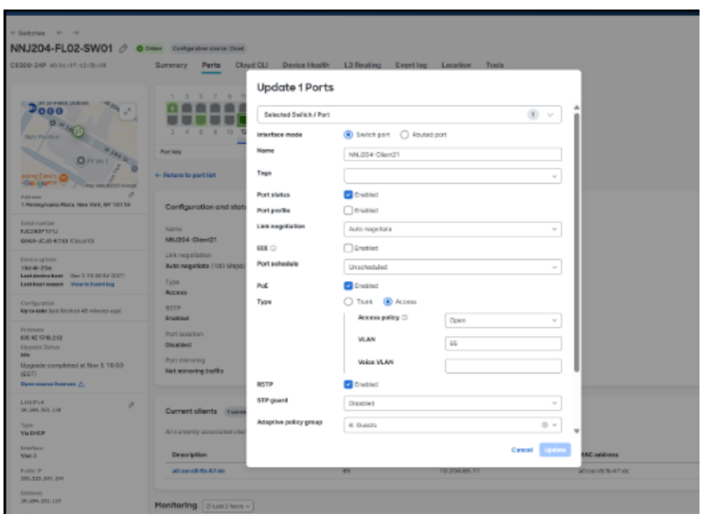

7. Dynamic VLAN and SGT assignment via 802.1x that also requires the fallback VLAN in the access policy.

Note: When an Access policy is set, the Adaptive policy group cannot be set on the port.

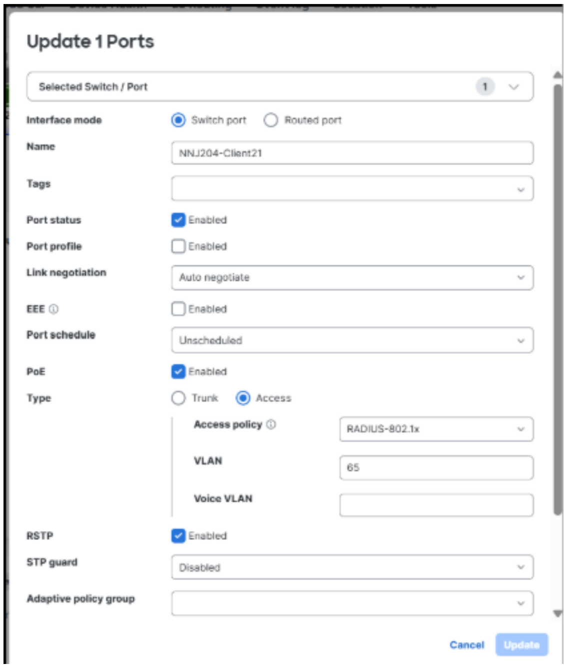

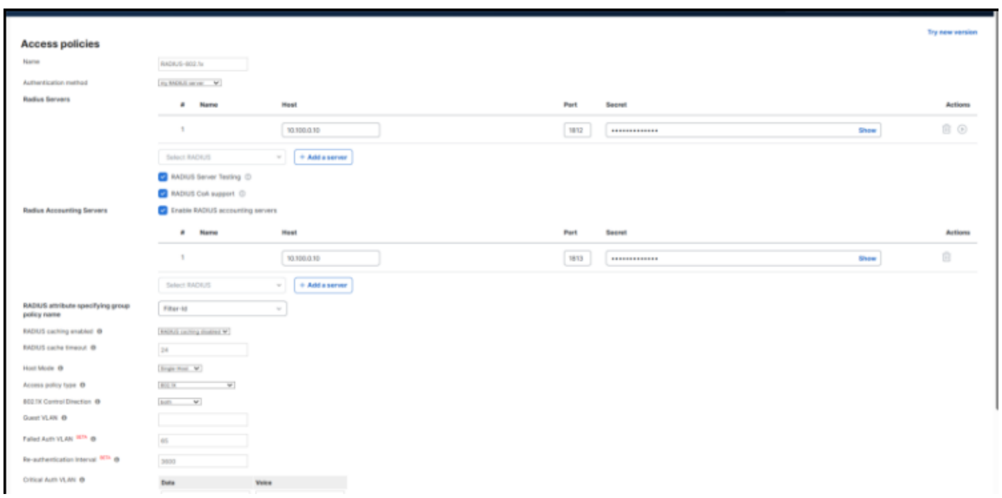

Final topology with guest VRF

Figure 9. Fabric with guest VFR – Layer 2 physical topology

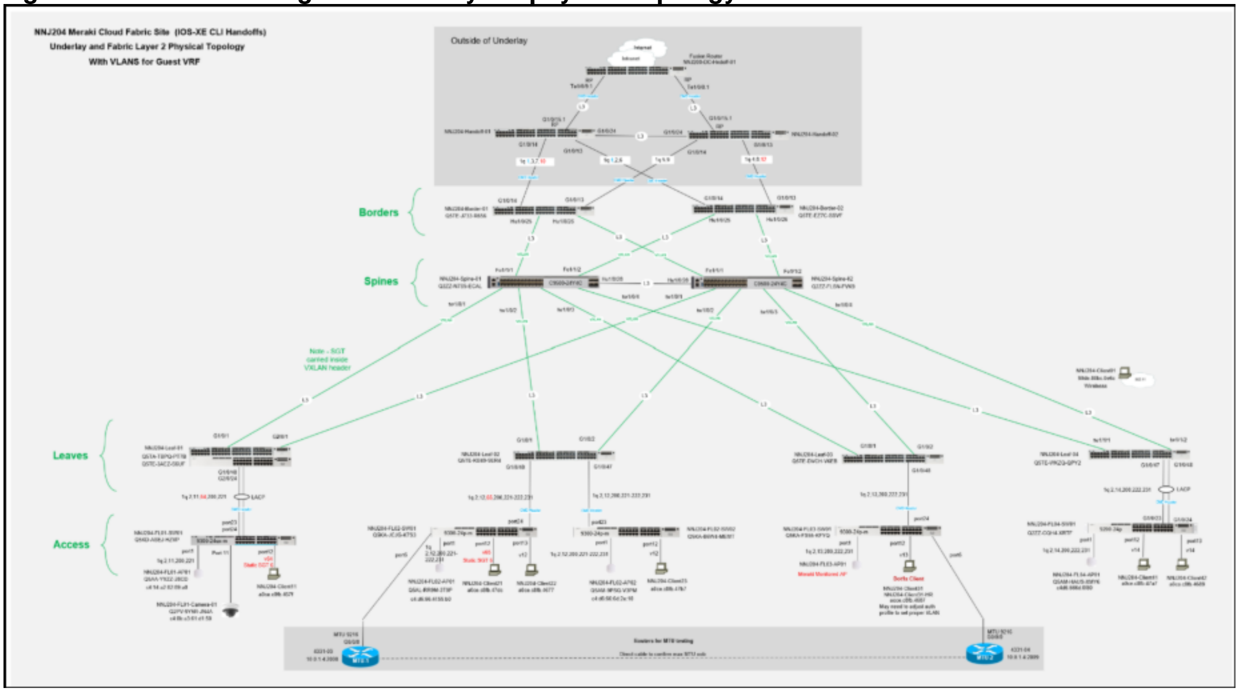

Figure 10. Fabric with Guest VRF – Layer 3 logical topology

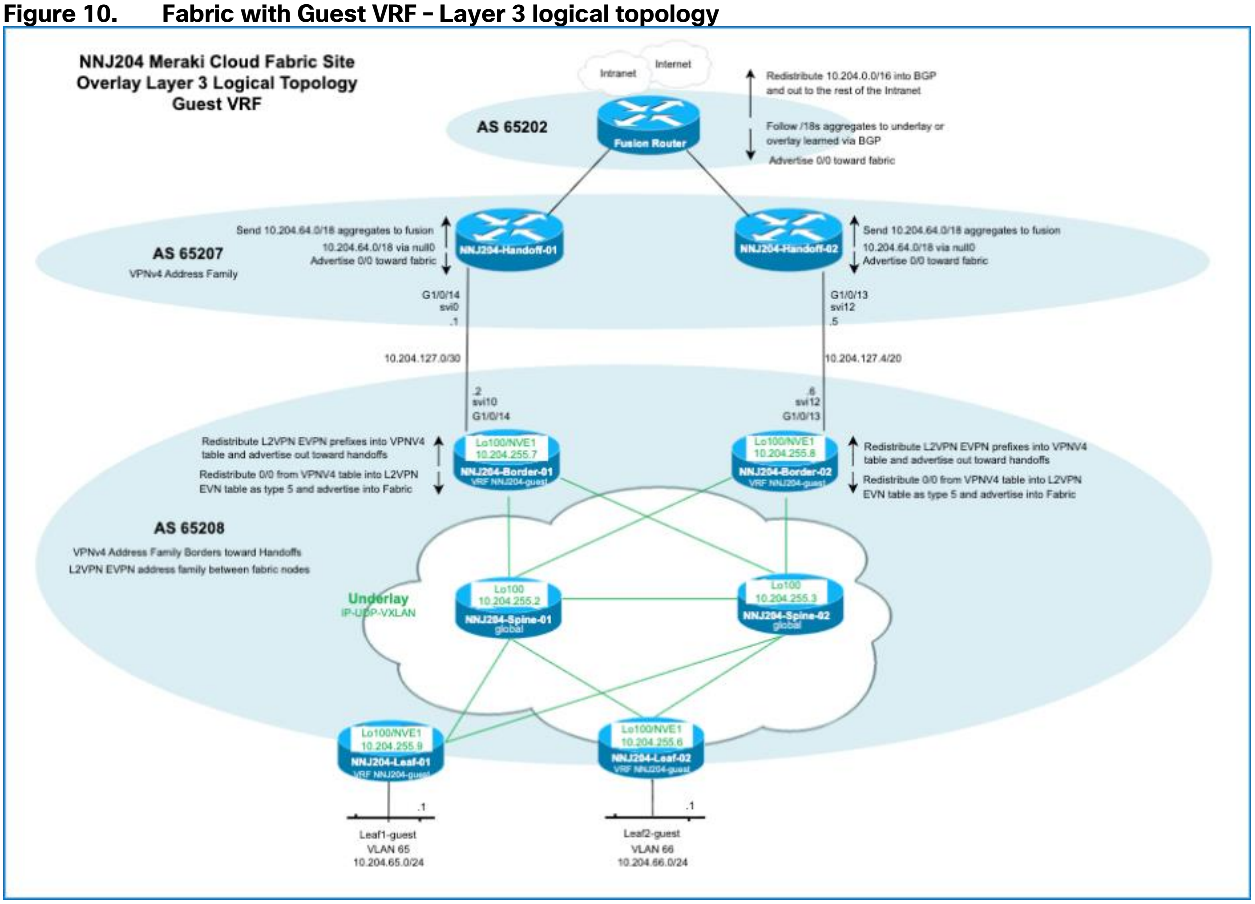

Conclusion

Through this Cisco Cloud Fabric validate case study, Cisco provides a future-ready foundation for mission-critical networks, combining deterministic performance, post-quantum security, Zero Trust enforcement, and intelligent automation in a unified architecture. By decoupling physical transport from policy-driven overlays, enforcing air-gapped trust boundaries, and standardizing on validated configurations, organizations can reduce risk, simplify operations, and scale securely across diverse environments. Powered by Cisco Secure Routers and enhanced by integrated observability and orchestration, this architecture ensures continuous availability, rapid adaptability, and long-term resilience for the most demanding missions.