Architecture Guide

Secure Network Reference Architecture (SNRA)

June 1, 2026

Executive summary

Modern campus and enterprise networks are evolving to support hybrid work, AI-driven applications, cloud adoption, distributed users, and an expanding threat landscape. Traditional siloed architectures often rely on fragmented management platforms, inconsistent policy models, and manually intensive operations that limit scalability, visibility, resiliency, and agility. Organizations require a simpler modernization approach that improves security while preserving operational flexibility and existing investments.

This architecture guide presents the Secure Network Reference Architecture (SNRA), a validated end-to-end framework for secure campus, branch, and distributed enterprise environments. SNRA integrates secure infrastructure, scalable segmentation, identity-driven policy, intelligent operations, and ecosystem integration into a unified architecture. The framework aligns networking, security, policy, and operations into a modular architecture that supports different operational models and deployment preferences while maintaining consistency across sites.

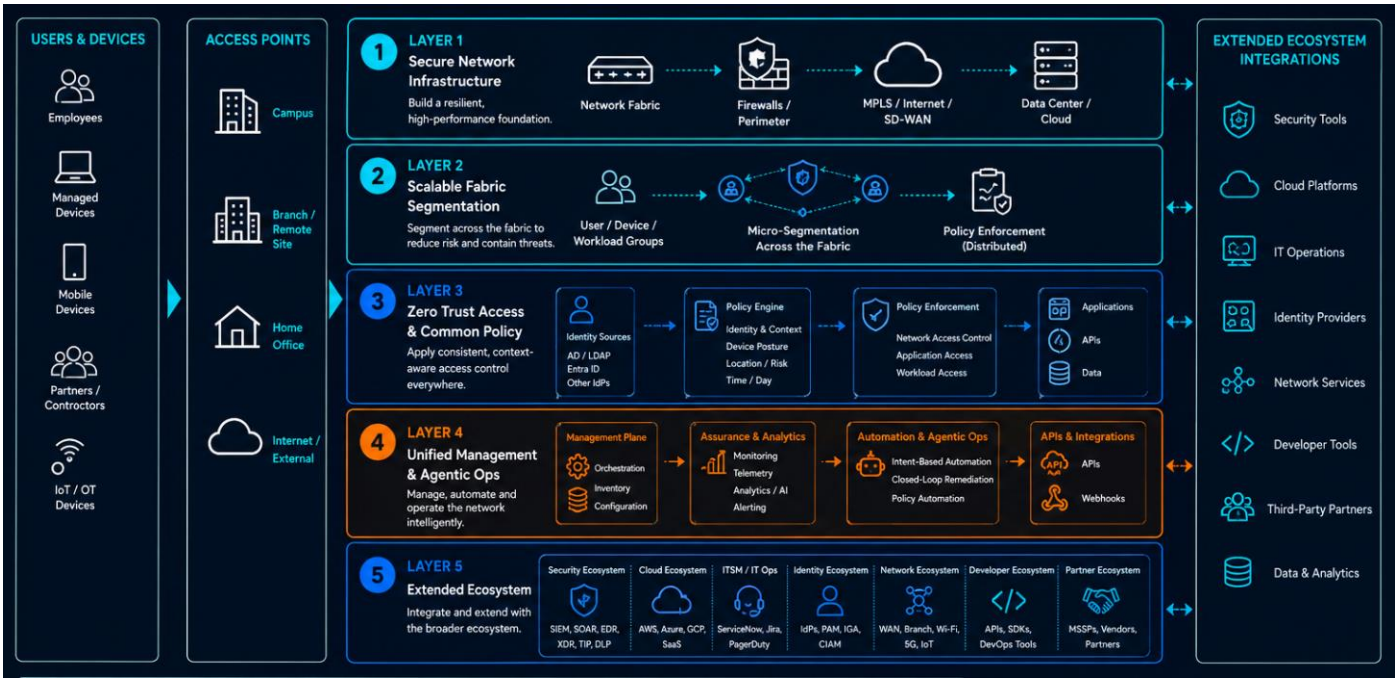

Figure 1. SNRA End-to-End Architecture

SNRA supports three deployment types:

- Cloud-managed
- On-premises
- Programmable

The cloud-managed model emphasizes simplified operations, centralized visibility, rapid onboarding, and scalable lifecycle management through SaaS-based platforms such as Cisco Meraki Dashboard and Security Cloud Control. The on-premises model is designed for organizations requiring localized control, governance alignment, or deeper operational customization. The programmable model supports environments where infrastructure automation, APIs, and DevOps or NetOps integration are strategic priorities.

SNRA is organized across five architecture domains:

- Secure Network Infrastructure
- Scalable Fabric Segmentation
- Zero Trust Access and Hybrid Mesh Firewall
- Unified Management and Agentic Ops
- Extended Ecosystem

These domains provide a scalable framework for secure connectivity, segmentation, policy enforcement, operations, and ecosystem integration across campus and branch environments.

Key benefits include:

- Improved security posture through embedded Zero Trust principles
- Identity-driven access control with consistent policy enforcement
- Reduced attack surface through segmentation and least-privilege access
- Simplified operations through centralized management and automation
- Improved visibility through integrated telemetry and analytics
- Flexible deployment choice across cloud-managed, on-premises, and programmable models
- Scalable architecture suitable for small, medium, and large enterprise environments

Security is integrated throughout the architecture rather than implemented as a standalone function. Identity, segmentation, policy enforcement, telemetry, and threat response operate together to provide consistent protection for users, devices, and workloads.

This guide focuses on architectural guidance and design rationale rather than implementation procedures. It explains what to build, why architectural decisions matter, and how to align technology investments to business outcomes such as stronger security, operational efficiency, lifecycle simplification, and improved user experience.

Introduction

Enterprise networks must securely connect users, devices, applications, sites, and cloud resources while delivering consistent performance, visibility, and control.

The Secure Network Reference Architecture (SNRA) aligns to Cisco's Future-Proofed Workplaces and Cisco Secure Networking strategy by enabling secure, adaptable, and operationally efficient workplace environments that support evolving business and workforce requirements.

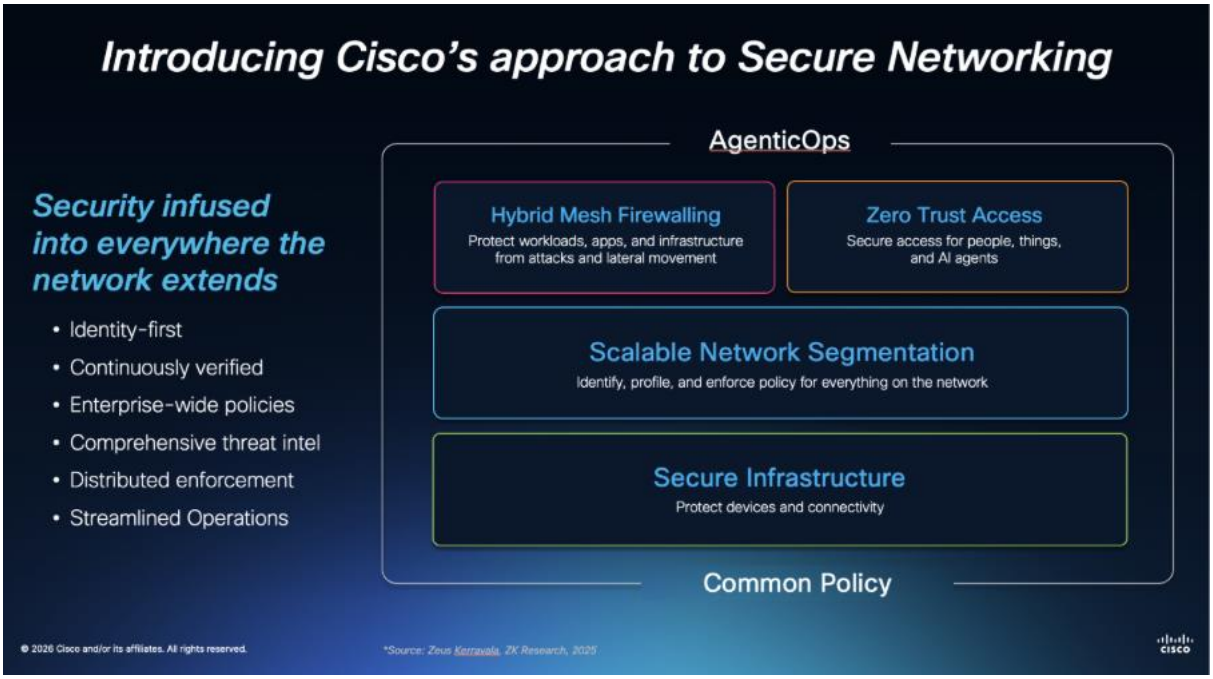

Figure 2. Cisco Secure Networking Reference Architecture

SNRA unifies networking, security, segmentation, policy, and operations into a modular architecture that supports cloud-managed, on-premises, and programmable deployment models. The architecture is organized across five domains that enable secure infrastructure, scalable segmentation, identity-driven policy, intelligent operations, and ecosystem integration.

This guide is intended for architects, engineers, partners, deployment teams, and customers evaluating secure networking strategies. It focuses on architectural guidance and validated design principles rather than configuration procedures.

Purpose of this guide

- Present the validated Secure Network Reference Architecture (SNRA) and associated Cisco design guidance
- Align Cisco Secure Networking customer outcomes to scalable and deployable architectural models
- Integrate networking, identity, segmentation, policy, visibility, and enforcement into a unified architecture
- Provide guidance for secure campus, branch, and distributed enterprise design decisions across multiple deployment models

Design objectives

The SNRA architecture is designed to help organizations build secure, scalable, and operationally efficient enterprise networks across campus, branch, and distributed environments.

Primary design objectives include:

- **Deliver resilient and secure** connectivity for users, devices, applications, and sites
- **Enable scalable segmentation** with consistent policy enforcement
- **Apply Zero Trust principles** using identity-aware access control and least-privilege design
- **Simplify deployment** and lifecycle operations through centralized management and automation
- **Improve** visibility, assurance, and operational response through integrated telemetry and analytics
- **Support** cloud-managed, on-premises, and programmable deployment models
- **Provide an extensible foundation** for future growth and ecosystem integration

Security integration considerations:

- Secure-by-design architecture across all layers
- Identity-driven policy enforcement for users, devices, and workloads
- Consistent policy enforcement across campus, branch, cloud, and remote access environments

Scope and limitations

This architecture guide provides high-level design guidance for implementing the Secure Network Reference Architecture across campus, branch, and distributed enterprise environments. It covers the architectural models, deployment approaches, design principles, and operational considerations for cloud-managed, on-premises, and programmable deployments.

SNRA is applicable to small, medium, and large-scale organizations and can scale across multi-site enterprise environments through modular design and validated deployment practices. Product selection, operational models, and feature adoption may vary based on customer requirements, compliance obligations, operational maturity, and existing infrastructure.

This document focuses on architecture and validated design outcomes rather than implementation detail.

The guide does not include:

- Step-by-step configuration procedures
- Command-line examples
- Migration runbooks
- Detailed low-level designs
- Exhaustive feature documentation

Product-specific deployment guidance should be referenced separately where required.

Final designs should account for site-specific constraints such as physical topology, application dependencies, operational maturity, compliance requirements, and performance objectives.

Business and technical drivers

Business drivers

Organizations continue to modernize infrastructure to support hybrid work, distributed operations, cloud adoption, and evolving security requirements. Modern enterprise architectures must improve agility, simplify operations, and maintain resiliency without increasing operational complexity.

Common business drivers include:

- Support hybrid workforces and distributed remote users
- Accelerate onboarding of new sites, users, and services
- Improve operational efficiency through centralized management
- Reduce downtime through resilient design practices
- Preserve existing investments while modernizing infrastructure
- Standardize user experience across campus and branch locations to applications everywhere from anywhere, both remote and on-premises
- Scale infrastructure in line with business growth

Technical challenges

Many enterprise environments evolved through using separate tools, inconsistent policy models, and manually intensive processes. These conditions can limit visibility, increase operational overhead, and expand security exposure.

Common technical challenges include:

- Fragmented networking and security management platforms
- Limited visibility across users, devices, applications, and sites
- Inconsistent segmentation and policy enforcement
- Manual provisioning and change management workflows
- Difficulty scaling legacy designs for new locations or modern workloads
- Complex troubleshooting across wired, wireless, WAN, and security domains
- Increased risk from flat or loosely controlled networks

Design requirements

The Secure Network Reference Architecture (SNRA) addresses these requirements through a modular, scalable, and deployment-flexible design model.

Functional requirements

The architecture should provide the following capabilities:

- Secure wired, wireless, and remote connectivity for users and devices
- Segmentation for business units, device classes, and trust zones
- Identity-aware access control with centralized policy enforcement
- Support for cloud-managed, on-premises, and programmable operations
- Integrated visibility for health, performance, and security monitoring

- Simplified onboarding of sites, devices, and services
- Extensibility through APIs and ecosystem integrations

Non-functional requirements

The architecture should also support:

- High availability and resilient service delivery
- Scalable growth across additional users, devices, and locations
- Consistent operational model across multiple sites
- Simplified lifecycle management and software maintenance
- Zero Trust–aligned security architecture
- Reduced operational overhead
- Flexibility for future technology and business requirements

Solution overview

The Secure Network Reference Architecture (SNRA) provides a validated framework for designing secure campus, branch, and distributed enterprise networks. The end-to-end architecture integrates networking, security, segmentation, policy, and operations into a unified model that simplifies deployment, improves resiliency, and enables consistent control across sites.

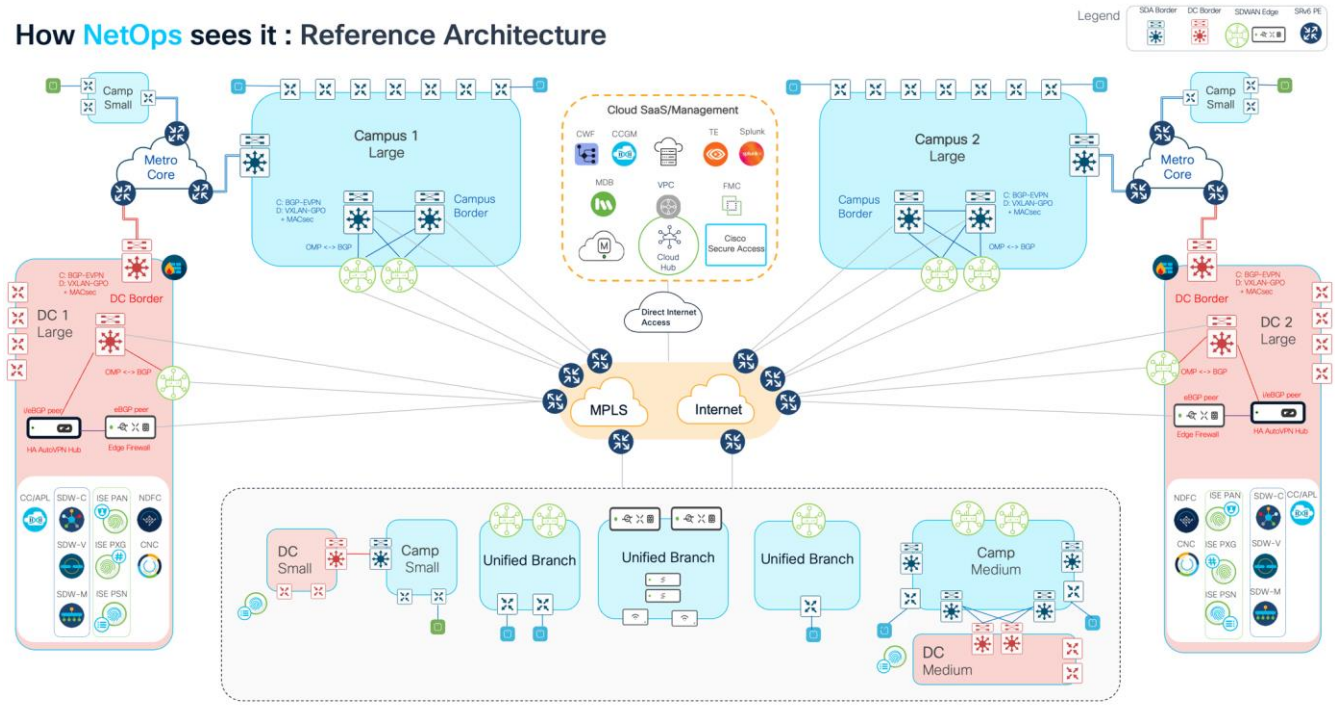

Figure 3. Cisco Secure Networking Reference Architecture

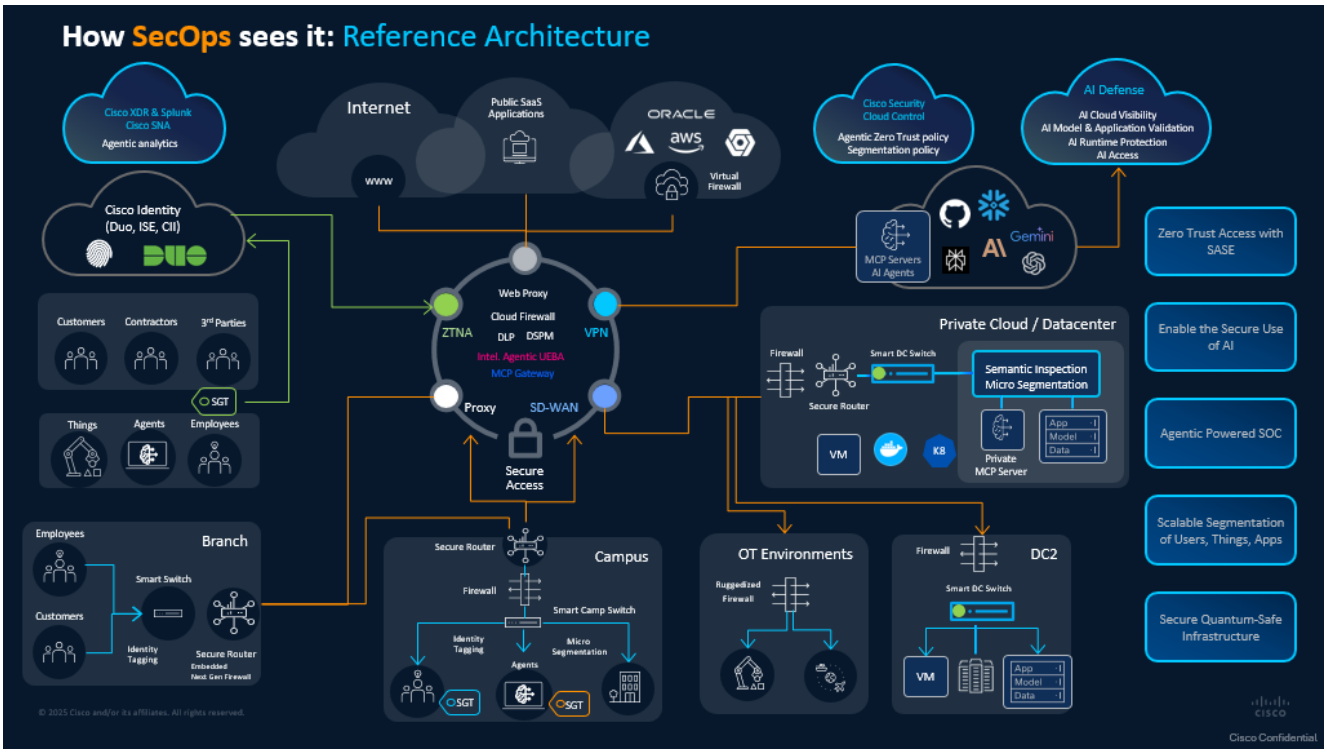

Figure 4. Security Operations Reference Architecture

SNRA supports multiple operational models while maintaining a consistent architectural foundation:

- Cloud-managed operations for centralized visibility and lifecycle simplicity
- On-premises operations for localized governance and customization
- Programmable operations for automation and orchestration-driven workflows

The architecture is organized across five separate design layers:

- Secure Network Infrastructure
- Scalable Fabric Segmentation
- Zero Trust Access and Hybrid Mesh Firewall
- Unified Management and Agentic Ops
- Extended Ecosystem

These layers provide a blueprint for secure connectivity, segmentation, identity-driven policy, intelligent operations, and extensibility.

SNRA reduces complexity by aligning wired, wireless, WAN, and security operations under a common architectural model. The result is a scalable and operationally consistent network architecture that supports evolving business, application, and security requirements.

Key capabilities

SNRA enables organizations to modernize networking and security operations through a consistent architecture that balances security, operational simplicity, and scalability.

The architecture delivers secure wired, wireless, and remote connectivity with consistent policy enforcement and reliable user experience. Integrated macro-segmentation and micro-segmentation help reduce risk, contain threats, and support organizational or regulatory separation requirements. Centralized visibility and automation simplify deployment, accelerate operational changes, and improve lifecycle management.

Key capabilities include:

- Secure connectivity across users, devices, and sites
- Scalable segmentation and Zero Trust policy enforcement
- Centralized visibility and simplified lifecycle operations
- Flexible cloud-managed, on-premises, or programmable deployment models
- Open ecosystem integration and automation readiness

Cisco solution components

SNRA can be implemented using Cisco technologies selected according to your operational requirements, deployment scale, and management preferences.

Reference solution components include:

- **Cisco Switches** for campus and branch LAN wired access, aggregation, and policy-enabled connectivity
- **Cisco Wireless** for secure mobility and wireless LAN user access
- **Cisco Routers** for branch connectivity, SD-WAN, and security services
- **Cisco Meraki Dashboard** for cloud-managed operations, monitoring, and lifecycle management
- **Cisco Catalyst Center** for centralized assurance, automation, and operational management in on-premises LAN environments
- **Cisco SD-WAN Manager** for centralized assurance, automation, and operational management in cloud or on-premises WAN environments
- **Cisco ISE** for identity services, device profiling, and policy-based access control
- **Cisco Secure Firewall Threat Defense (FTD)** for advanced threat protection, segmentation enforcement, and secure Internet and data center edge connectivity
- **Cisco Security Cloud Control (SCC)** for centralized security policy management, visibility, and multi-platform security operations
- **Cisco Hybrid Mesh Firewall** for distributed policy enforcement across campus, branch, cloud, and remote access environments
- **Secure Access** and related Cisco security services for secure user connectivity and policy enforcement
- **ThousandEyes** for digital experience monitoring and end-to-end visibility

Reference architecture for SNRA model

The Secure Network Reference Architecture (SNRA) organizes modern networking, security, segmentation, operations, and ecosystem integration into five coordinated design domains that provide a scalable and repeatable enterprise framework.

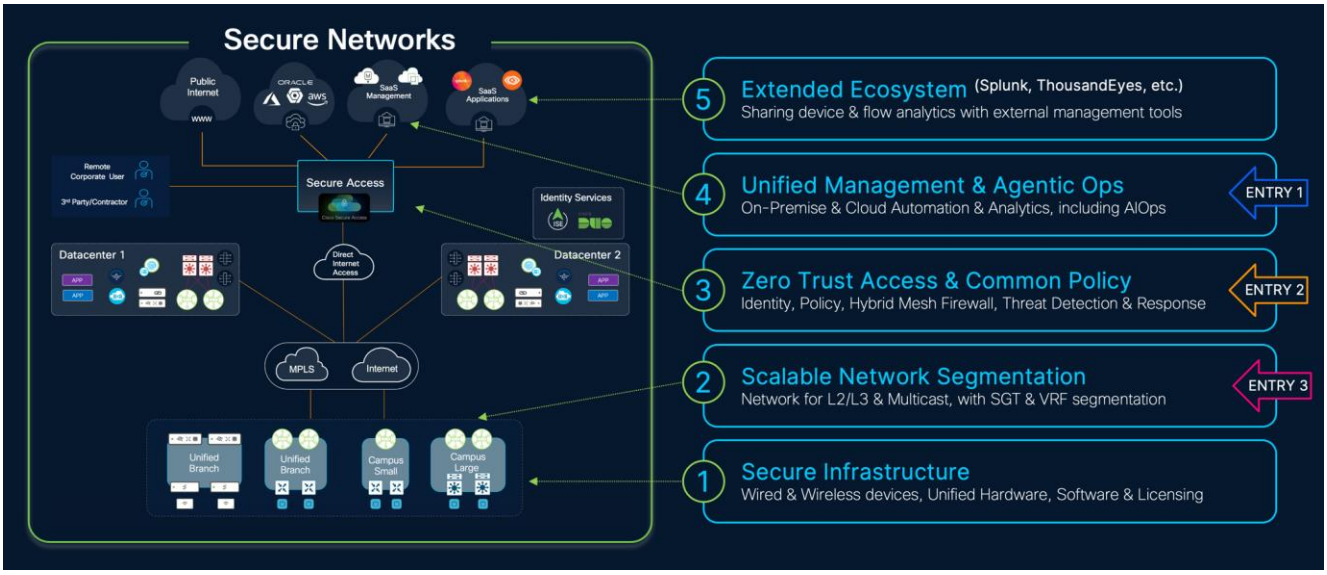

SNRA multi-layer architecture model

The SNRA framework maps architectural domains to operational components and deployment roles into five distinct reference layers. Each layer contributes to a coordinated architecture that aligns infrastructure, segmentation, policy, operations, and ecosystem integration.

SNRA Layer	Network Components	Purpose
Secure Network Infrastructure	Cisco Switches, Cisco Wireless Access Points, Cisco Secure Routers	Foundational physical connectivity, secure access, and resilient transport
Scalable Fabric Segmentation	Campus Fabric for LAN Switching and Wireless, SD-WAN for Secure Campus and Branch	Virtualized traffic forwarding for scalable end-to-end segmentation
Zero Trust Access & Hybrid Mesh Firewall	Cisco TrustSec (CTS), Cisco ISE, and Group Policies via Cisco SCC	Identity-based access control and consistent policy enforcement
Unified Management & Agentic Ops	Network & Security Dashboards, APIs, Telemetry, Automation tools	Centralized visibility, automation, and intelligent operations
Extended Ecosystem	Cisco Splunk, Marketplace integrations, third-party platforms	Integration with external tools for expanded capabilities

Layer 1: Secure Network Infrastructure

This layer provides the physical and logical foundation of the network through Cisco switching, wireless, and WAN edge platforms.

Note: Cisco infrastructure platforms incorporate integrated security capabilities across both hardware and

software generations. Newer Cisco secure infrastructure platforms further extend these capabilities with advanced protections such as post-quantum cryptography (PQC), PQC Secure Boot, PQC encryption, and eBPF-based runtime protection features.

For more information on Cisco Secure Infrastructure products, refer to the Validated software and hardware versions section for this document.

Cisco switching platforms provide a Layer 3-based underlay architecture using routed links to improve resiliency and reduce Layer 2 dependency. Cisco wireless platforms extend secure connectivity and policy consistency to wireless users and devices. Cisco WAN edge platforms provide external connectivity and integrate campus and branch environments with WAN and Internet services.

This layer provides secure, resilient, and predictable connectivity for users, devices, and locations so segmentation and policy enforcement can be applied consistently across the environment.

Layer 2: Scalable Fabric Segmentation

This layer implements scalable segmentation across the wireless and wired LAN, campus, branch, WAN, and extended enterprise environments. Cisco switches enforce macro-segmentation locally using VLAN and IP subnet boundaries, while VRF-aware and policy-aware VXLAN forwarding at the campus access or distribution layers provides scalable traffic isolation, segmentation consistency, and reduced Layer 2 dependency across the fabric.

Segmentation extends beyond the campus through Cisco routers and policy-aware forwarding constructs that preserve segmentation across WAN and external connectivity domains. This approach transforms segmentation into a scalable, network-wide architectural construct rather than a localized configuration model.

Layer 3: Zero Trust Access and Hybrid Mesh Firewall

This layer introduces identity-driven policy enforcement across the enterprise.

Cisco ISE authenticates users and devices and assigns identity attributes used for policy enforcement through Cisco TrustSec, Security Group Tags (SGTs), and group-based policy models. Policy enforcement is applied across wired, wireless, WAN, and security domains.

Cisco switching and wireless platforms enforce access policy at the edge, while firewalls and WAN edge platforms enforce policy between trust zones and forwarding domains. Hybrid mesh firewall principles allow enforcement to be distributed throughout the environment based on business, security, and operational requirements.

Layer 4: Unified Management and Agentic Operations

This layer centralizes network and security operations, visibility, policy management, automation, and lifecycle management through platforms such as Cisco Meraki Dashboard, Cisco Catalyst Center, Cisco SD-WAN Manager and Cisco Security Cloud Control (SCC). These platforms provide consistent interfaces for configuration, monitoring, assurance, and operational management across network infrastructure, Cisco security platforms, and supported third-party firewalls.

Guided workflows, zero-touch provisioning, template-driven deployment, programmable APIs, orchestration, and continuous telemetry collection simplify deployment and lifecycle operations while enabling integrated monitoring, analytics, operational response, and external system integration across the environment.

Layer 5: Extended Ecosystem

This layer integrates the network with external analytics, security, automation, and operational platforms through APIs and ecosystem integrations.

Platforms such as Cisco Meraki Dashboard, Cisco Catalyst Center, Cisco ISE, Cisco SD-WAN Manager and Cisco SCC provide integration interfaces for Cisco and third-party systems including Cisco Splunk, Cisco ThousandEyes, and operational analytics platforms.

These integrations extend visibility, automation, and coordinated operational workflows across the broader IT and security ecosystem.

End-to-end architecture flow

The SNRA model follows a logical flow with secure infrastructure, extends through segmentation and identity-based policy enforcement, and is continuously supported by centralized operations, telemetry, automation, and ecosystem integration.

Users, devices, and applications connect through a secure campus, branch, WAN, and cloud infrastructure. Traffic is segmented reliably across micro-segmentation policies, while identity-aware controls enforce Zero Trust access decisions across the environment.

Operational platforms provide assurance, visibility, automation, and lifecycle management across the environment, while ecosystem integrations extend coordination with external IT and security systems.

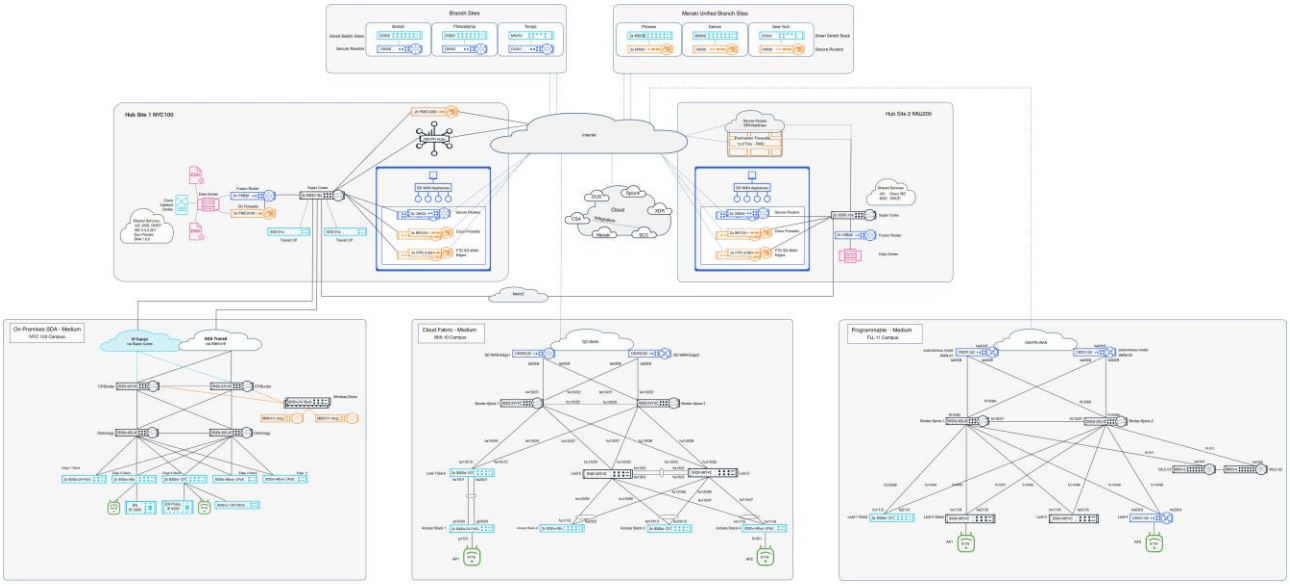

Figure 5. SNRA Traffic flow between layers

Architecture principles

SNRA is based on design principles intended to simplify operations, improve security outcomes, and support long-term scalability.

- Integrate security throughout the architecture
- Use segmentation to establish logical trust boundaries
- Apply identity-aware policy consistently across users, devices, and locations
- Standardize operations through centralized management and automation

- Design for resiliency, scalability, and lifecycle simplicity
- Preserve flexibility across cloud-managed, on-premises, and programmable models
- Support open integration for analytics, orchestration, and future extensibility

Deployment models

Supported deployment types

SNRA supports three operational deployment types:

- Cloud Managed
- On-premises
- Programmable (DIY)

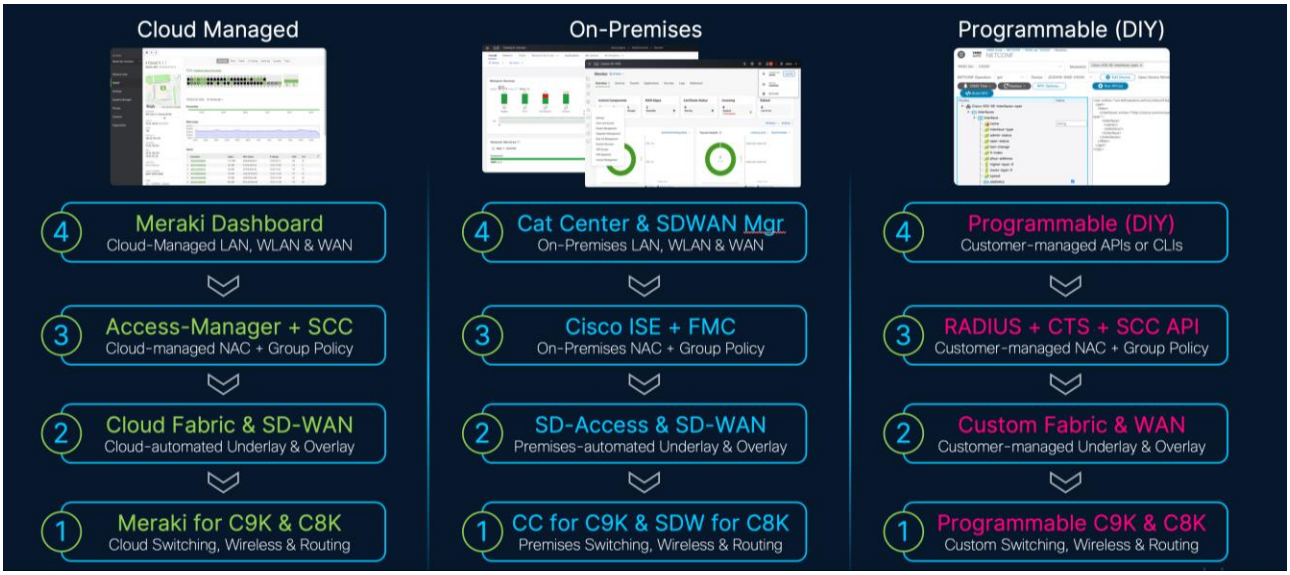

Figure 6. Supported Deployment Types

Each model follows the same architectural principles while providing different approaches to operations, automation, policy management, and lifecycle control.

SNRA supports environments ranging from smaller deployments with simplified operational requirements to large multi-site enterprises requiring advanced segmentation, resiliency, and scale. The architecture also supports specialized environments with requirements such as high availability, enhanced security controls, multicast scale, real-time media, and smart building integration.

Supported deployment scale

Typical deployment profiles include:

- **Small** environments using compact one-tier or two-tier LAN designs supporting up to approximately 5,000 endpoints
- **Medium** environments using two-tier or three-tier LAN designs supporting up to approximately 25,000 endpoints per site
- **Large** environments using three-tier or four-tier LAN designs supporting 50,000 or more endpoints across complex campuses or distributed locations

The modular scaling architecture also supports specialized operational requirements such as high-availability environments, high-security deployments, multicast-intensive operations, digital media workloads, and smart building infrastructure.

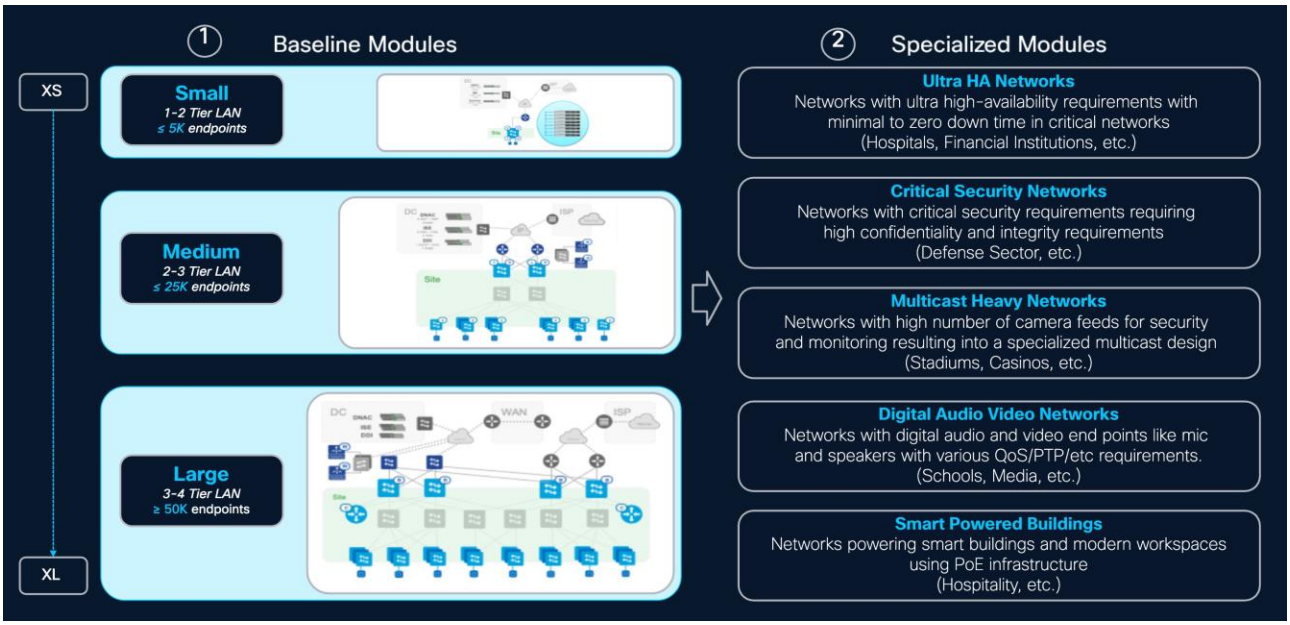

Figure 7. Secure Campus Network Modular Design

Deployment model comparison

Cloud-Managed

The cloud-managed model uses platforms such as Cisco Meraki Dashboard to simplify deployment, operations, monitoring, and lifecycle management for Cloud Fabric and SD-WAN environments. It is well suited for distributed enterprises requiring operational simplicity and centralized visibility.

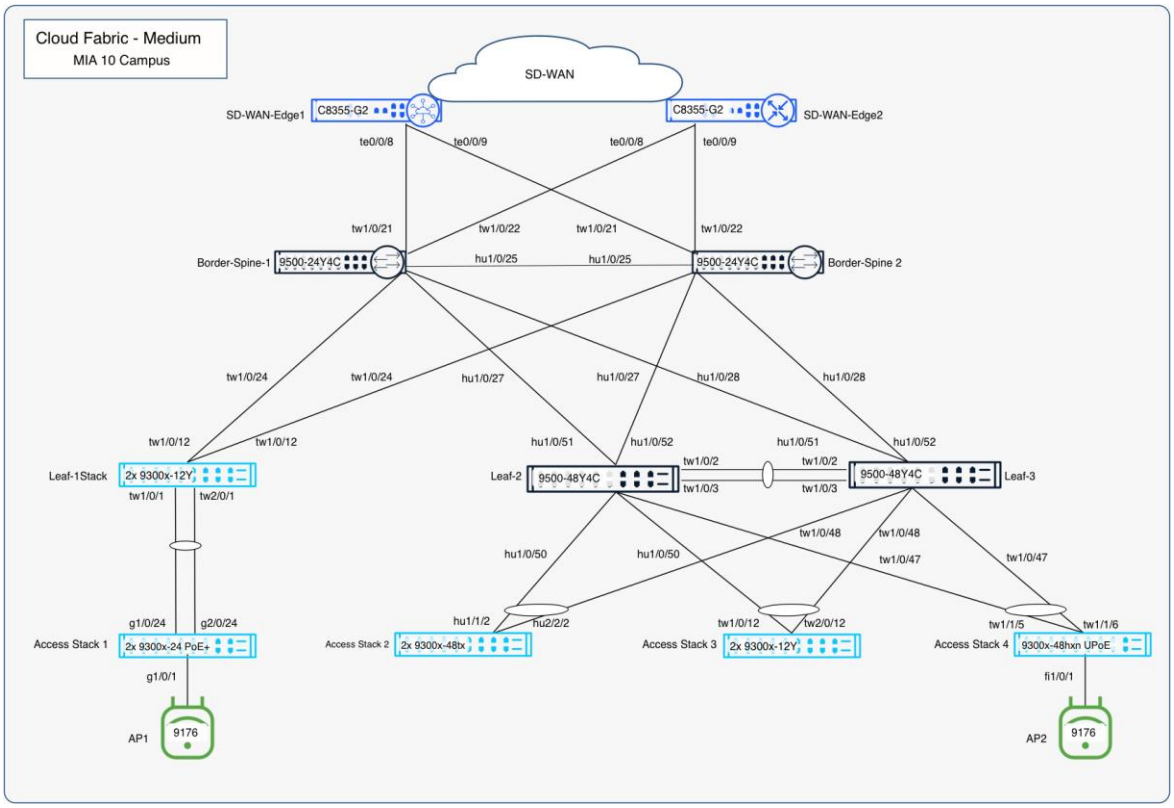

Figure 8. Cloud Fabric Medium Campus Simple Diagram

On-Premises

The on-premises model uses Cisco Catalyst Center, Software-Defined Access (SD-Access), and Cisco SD-WAN Manager to provide localized operational control, advanced segmentation, and deeper enterprise integration.

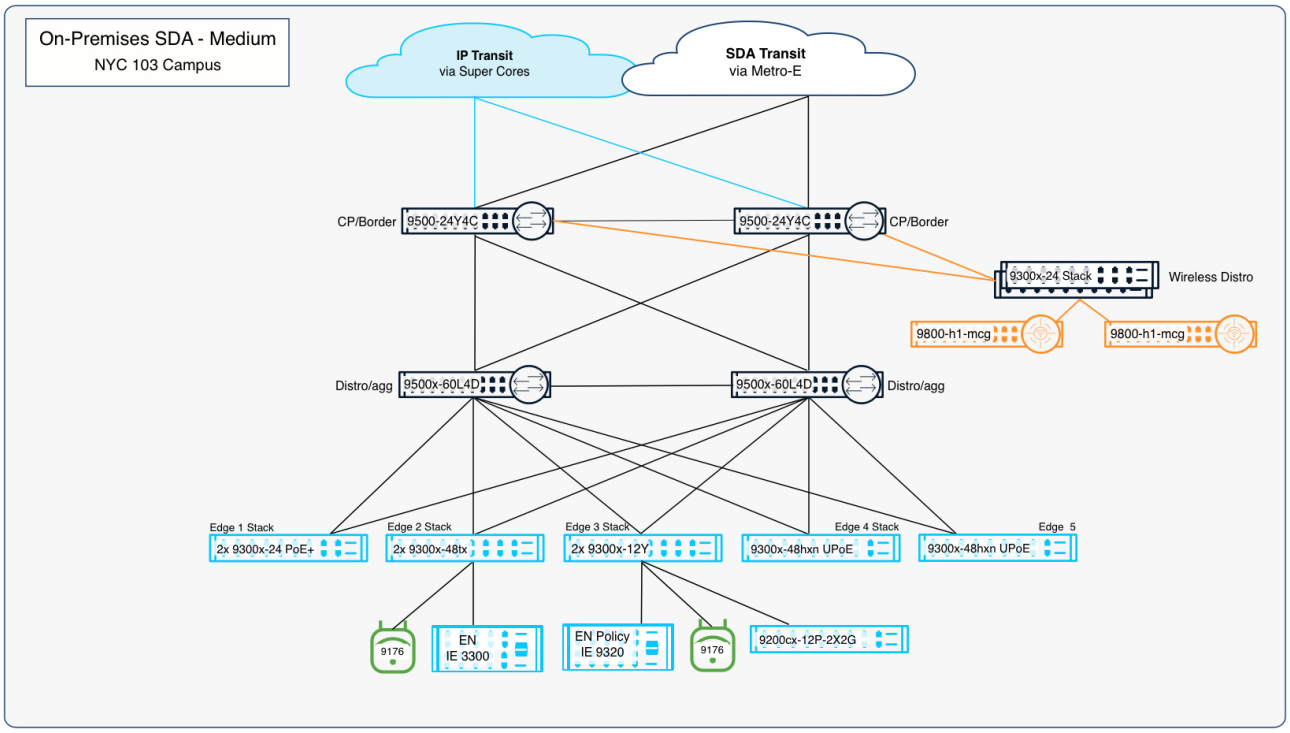

Figure 9. On-Premises Medium Campus Diagram

Programmable (DIY)

The programmable model emphasizes direct platform control through command-line operations, programmable APIs, and automation frameworks (e.g. Ansible, Terraform, etc.). It is intended for organizations requiring advanced customization and automation flexibility.

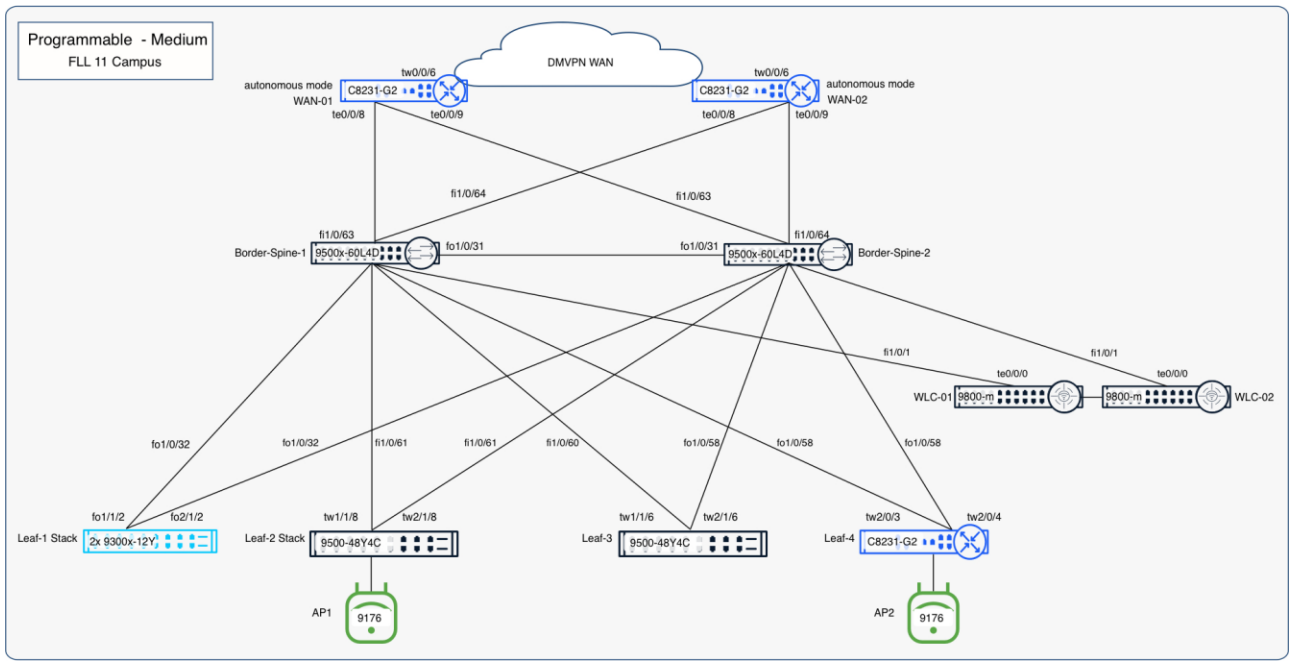

Figure 10. Programmable Medium Campus Diagram

Deployment Model	Primary Strength	Best Fit
Cloud-managed	Operational simplicity and centralized visibility	Distributed enterprises, lean IT teams, rapid growth
On-premises	Deep campus control and integrated enterprise policy	Large campuses, governance-focused organizations
Programmable (DIY)	Maximum customization and automation flexibility	Advanced engineering teams, bespoke environments

All three deployment models align to the same SNRA architectural outcomes.

Model selection guidance

Deployment model selection should align to operational priorities, governance requirements, internal skill sets, and long-term infrastructure strategy.

Key selection factors include:

- Operational simplicity versus customization needs
- Cloud-managed preference versus on-premises control
- Internal automation and engineering maturity
- Environment scale and geographic distribution
- Security, compliance, and governance requirements
- Existing platform investments and standards

Organizations may adopt a hybrid approach across different sites or business units while maintaining alignment to the overall SNRA architecture.

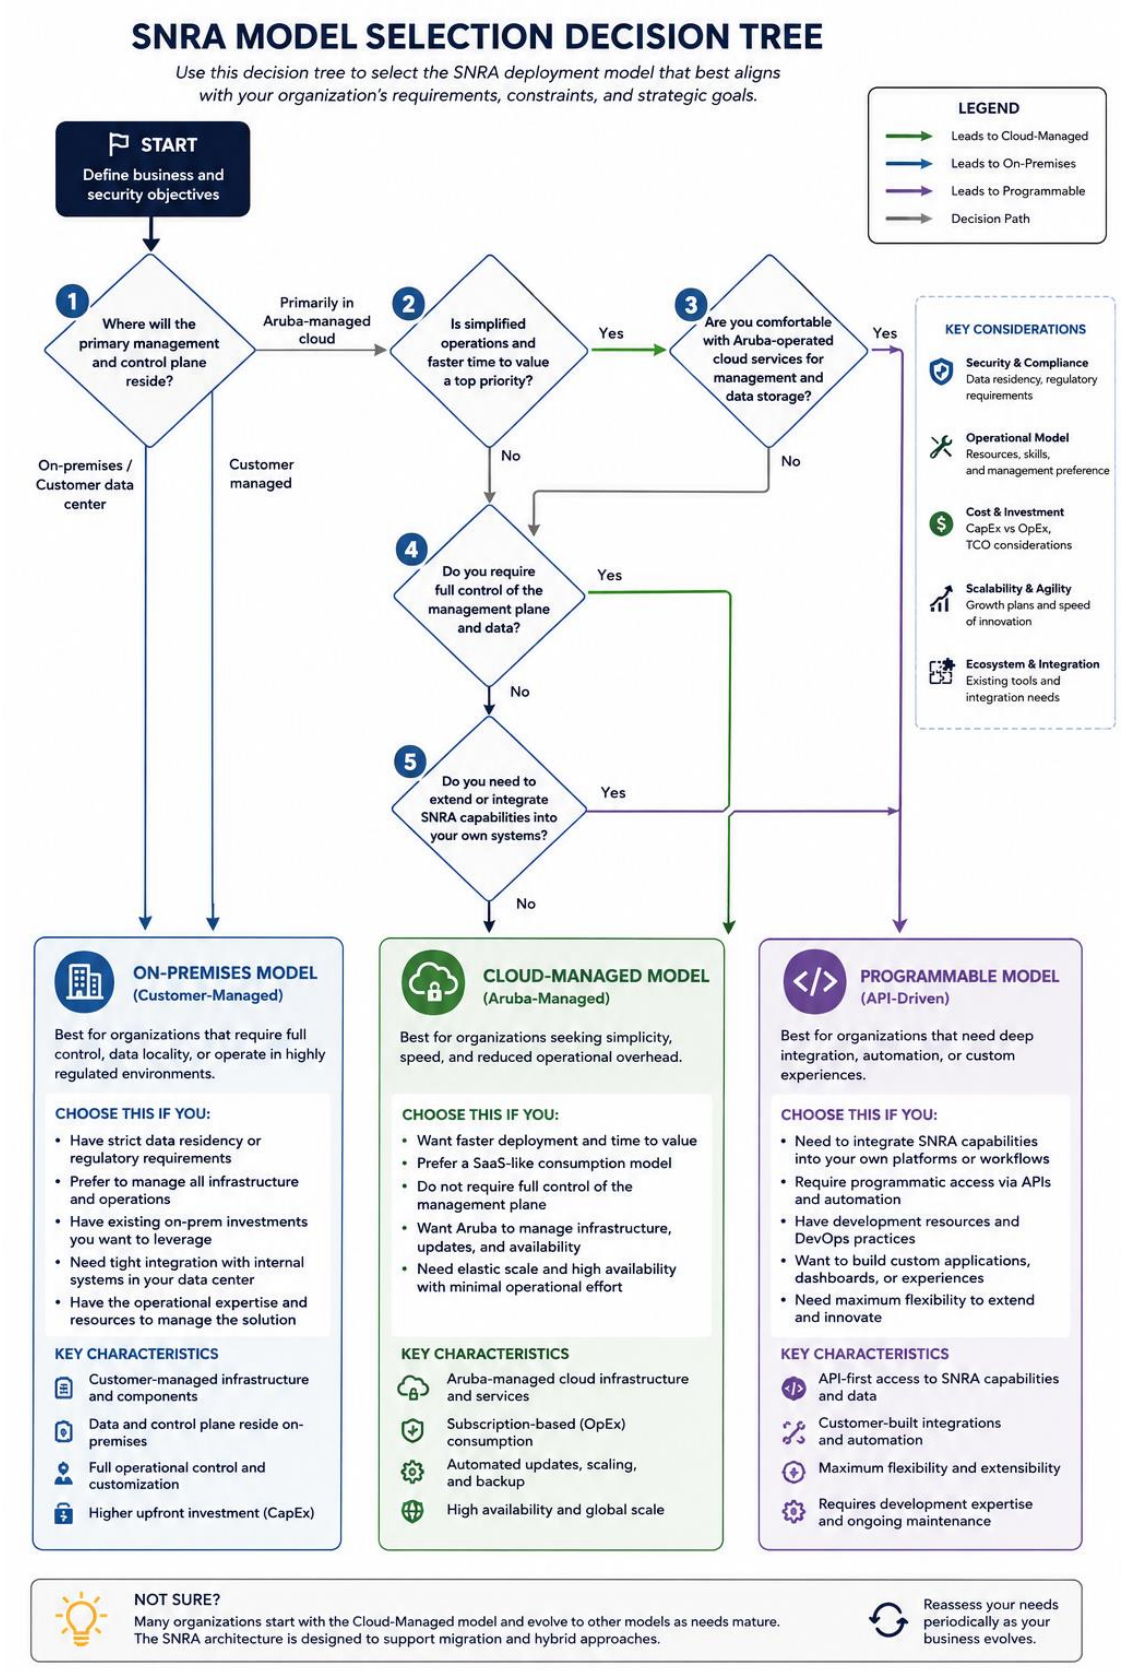

Figure 11. SNRA Decision Tree

Layer-by-layer design

The following section examines each of the five SNRA architecture domains using a consistent format, explaining what should be built, why it matters, and how security, scale, and operational outcomes are addressed within each layer of the overall design.

Layer 1: Secure Network Infrastructure

Secure Network Infrastructure provides the physical and logical foundation of the SNRA model. This layer delivers resilient connectivity, secure transport, and trusted access for users, devices, applications, and sites across campus, branch, remote, and cloud-connected environments. It includes Cisco switching, wireless, routing, WAN edge, and foundational network services that support all higher-layer segmentation, policy, security, and operational functions.

This layer establishes the secure transport foundation required for scalable segmentation, identity-aware policy enforcement, operational visibility, and lifecycle management across the enterprise. A well-designed infrastructure layer improves availability, simplifies expansion, and provides predictable performance for business-critical services.

Typical technologies used in this layer include:

- Cisco C9000 series switches for campus access, distribution and core connectivity
- Cisco C9100 and 9800 series wireless access points and wireless controllers
- Cisco C8000 series secure routers for branch, WAN, and Internet edge services
- Cisco Catalyst Center, Cisco Meraki Dashboard, Cisco SD-WAN Manager or programmable operational models - varies by deployment type
- Cisco Dynamic routing protocols such as OSPF, OMP and BGP
- DHCP, DNS, NTP, and foundational IP services
- High-availability mechanisms, resilient uplinks, and redundant forwarding paths

Core capabilities, resiliency, and security

Secure Network Infrastructure delivers resilient wired and wireless connectivity, Secure WAN, and consistent access services across campus, branch, and cloud environments. The layer provides the scalable transport foundation required for segmentation, policy enforcement, telemetry, assurance, and lifecycle operations.

Resiliency is achieved through redundant uplinks, diverse forwarding paths, modular topology design, and scalable routing boundaries that reduce operational complexity and limit single points of failure. Smaller deployments may use one-tier or two-tier designs, while larger environments often use two-tier, three-tier, or campus core architectures aligned to traffic patterns and endpoint scale.

Security controls include secure device onboarding, management-plane protection, role-based access control, encrypted management communications, infrastructure segmentation, and control-plane

protection. Cisco Catalyst switches, Cisco wireless platforms, Cisco routers, and WAN edge platforms also provide early enforcement points for identity-aware access control and higher-layer security services.

Design considerations

Key design decisions at this layer include topology selection, uplink redundancy, routing boundaries, wireless coverage models, WAN diversity, and operational ownership. Threats relevant to this layer include unauthorized access, device compromise, denial-of-service conditions, and exposure to management interfaces.

Organizations should also plan for telemetry generation such as interface health, client statistics, path quality, environmental alerts, and control-plane events to support assurance and troubleshooting.

- Prioritize cloud-managed operations when simplified administration, rapid deployment, and centralized multi-site visibility are strategic objectives.
- Select on-premises management when local governance, deeper campus control, or internal operational ownership is required.
- Adopt programmable operations when business success depends on advanced automation, custom integrations, or infrastructure-as-code workflows.
- Use modular multi-tier designs when endpoint growth, service expansion, or rising traffic demand requires scalable long-term capacity.
- Engineer for continuous availability through redundant paths, power diversity, and resilient platform roles when uptime is mission critical.
- Standardize templates, naming conventions, and addressing models when operating across geographically distributed sites to accelerate scale and reduce complexity.

Best practices

- Build resiliency into the architecture from the start rather than treating availability as a later enhancement.
- Standardize hardware roles and topology patterns to create repeatable deployments and lower operational overhead.
- Secure management access with strong authentication, role-based control, and least-privilege principles.
- Enable logging, telemetry, and proactive health monitoring from day one to improve visibility and accelerate issue resolution.
- Align wireless, switching, and WAN capacity plans to meet business growth and user experience demands.
- Maintain disciplined software lifecycle management and configuration governance.

Layer 2: Scalable fabric segmentation

Scalable Fabric Segmentation provides the logical separation framework of the SNRA model. This layer isolates users, devices, applications, and business services into trusted boundaries that improve security, simplify operations, and create predictable traffic flows across campus, branch, WAN, and extended enterprise environments.

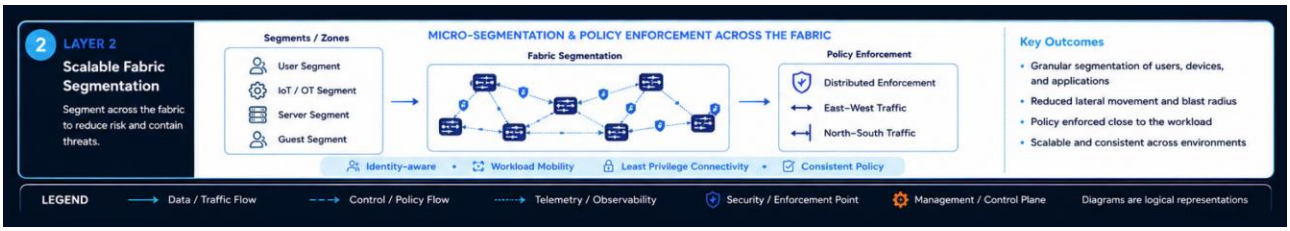

This layer transforms the network from a flat transport model into a scalable architecture where communications between segments are intentional, policy-driven, and operationally consistent. Logical segmentation enables organizations to expand services, sites, and user groups without redesigning the underlying infrastructure.

Typical technologies and platforms used in this layer include:

- Cisco Catalyst 9000 Series switches
- Cisco Catalyst Center fabric services
- VLAN and IP subnet-based macro-segmentation
- Virtual Routing and Forwarding (VRF) instances
- VXLAN EVPN overlays
- Cisco Software-Defined Access (SD-Access) segmentation constructs
- Cisco TrustSec and Security Group Tag (SGT)-based classification
- Inter-segment gateways and distributed policy enforcement nodes
- Cisco routing platforms for segmented WAN and external connectivity

Core capabilities, resiliency, and security

Scalable Fabric Segmentation delivers controlled traffic separation, scalable forwarding, and policy-aligned communication boundaries across the enterprise. Cisco switches enforce macro-segmentation using VLAN and IP subnet boundaries, while VXLAN EVPN and VRF-based forwarding provide scalable segmentation across campus access and distribution layers.

Macro-segmentation isolates trust zones such as employee, guest, IoT, voice, OT, management, and data center environments. Micro-segmentation extends control further through identity-aware and group-based policy enforcement between users, devices, workloads, and applications.

Segmentation extends beyond the campus through Cisco routers and segmented forwarding constructs that preserve traffic isolation across WAN and external connectivity domains. This approach reduces attack surface, limits unauthorized east-west movement, and aligns communications to business intent.

Resiliency is improved by reducing fault domains and limiting the operational impact of network or security events. Fabric-based segmentation also allows organizations to scale sites and services without significant infrastructure redesign.

Policy enforcement can be distributed across access layers, fabric boundaries, inter-segment gateways, or integrated security platforms. Distributed enforcement improves containment and scalability, while centralized inspection simplifies governance and operational consistency. Most enterprise environments use a blended enforcement model aligned to hybrid mesh firewall principles.

Design considerations

Key decisions include the number of segment counts, routing boundaries, policy ownership, enforcement location, and operational manageability. Over-segmentation can increase complexity, while under-segmentation can increase exposure and reduce visibility.

Threats relevant to this layer include unauthorized lateral movement, excessive trust between systems, inherited access permissions, and inconsistent policy across sites.

Telemetry should include inter-segment flows, denied communications, policy hit counts, trust group mappings, and anomalous east-west traffic patterns.

- Use VRFs and macro-segmentation when business separation, compliance boundaries, or service isolation are required.
- Apply micro-segmentation through identity-aware and group-based policy models when least-privilege access and threat containment are operational priorities.
- Use centralized gateways or security boundaries when unified inspection and simplified governance are strategic requirements.
- Distribute enforcement closer to access layers when scalability and rapid containment are required.
- Standardize reusable policy groups and centrally managed templates across multiple sites to simplify expansion and maintain consistent security posture.

Best practices

- Use macro-segmentation for primary trust boundaries and micro-segmentation for granular least-privilege enforcement.
- Enforce least-privilege communication by default.
- Standardize segment naming, policy intent, and operational models across sites.
- Remove unused or overly permissive communication paths regularly.
- Enable telemetry for inter-segment traffic, denied flows, and policy effectiveness.
- Align segmentation strategy to long-term business services and operational models.

Layer 3: Zero trust access and hybrid mesh firewall

Zero trust access and hybrid mesh firewall provide the identity-driven control layer of the SNRA model. This layer validates trust before connectivity is granted and applies consistent policy across users, devices, workloads, applications, and locations using identity, posture, role, device type, and contextual signals.

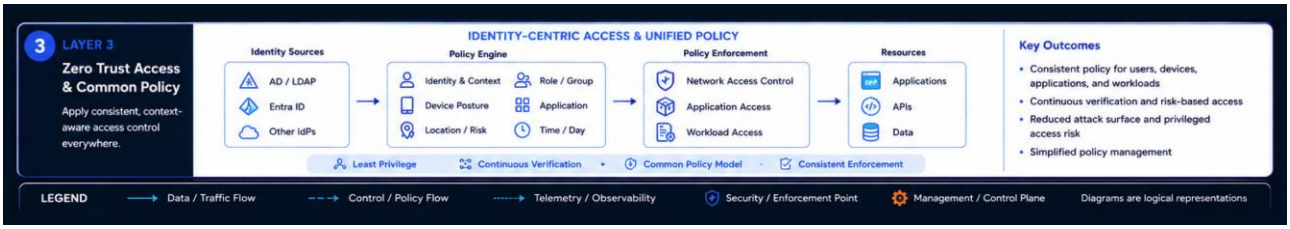

This layer shifts access control from implicit trust to continuous verification. By applying hybrid mesh firewall principles across wired, wireless, branch, remote, and application environments, the architecture improves security, operational consistency, and user experience while simplifying governance.

Identity-driven policy extends segmentation beyond traditional network boundaries. Macro-segmentation establishes broad trust zones, while identity-aware and group-based policies provide granular least-privilege enforcement between users, devices, workloads, and applications.

Typical technologies and platforms used in this layer include:

- Cisco Identity Services Engine (ISE)
- Cisco Secure Firewall Threat Defense (FTD)
- Cisco Security Cloud Control (SCC)
- Cisco TrustSec and Security Group Tags (SGTs)
- Authentication services for network, application, and Zero Trust access, including certificate-based trust and MFA integration
- Certificate-based authentication services
- Device profiling and posture assessment services
- Policy engines and authorization rules
- Cisco Secure Access
- Cisco XDR (, NDR, and endpoint telemetry platforms

Note: Cisco XDR supports Layer 3 through trust and policy context but is primarily aligned to Layer 5 for telemetry correlation and response orchestration.

Core capabilities, resiliency, and security

This layer enables identity-first access control with consistent policy enforcement across the enterprise. Users and devices can be authenticated once and governed through common policies regardless of connection method or location. Policies can be based on role, device type, security posture, time, location, or business context.

Zero Trust principles guide the architecture through least-privilege enforcement, continuous trust verification, and explicit validation. Access can be dynamically adjusted as conditions change, helping reduce exposure while maintaining productivity.

Resiliency is achieved through redundant identity services, distributed enforcement, survivable authentication methods, and fallback access strategies for critical operations. Centralized policy models and integrated Day-2 operational visibility help simplify lifecycle management, accelerate response, and maintain consistent control across the environment.

Security benefits include reduced unauthorized access, stronger credential governance, controlled device onboarding, and adaptive policy enforcement across campus, branch, remote, and application environments.

Design considerations

Key decisions vary based on Zero Trust use cases, access methods, and operational models. Architectures may use cloud identity providers, Cisco ISE, certificate-based trust, network access control systems, or application access brokers to deliver consistent policy across network and application environments.

Organizations must determine how identity-driven segmentation policies are applied across users, devices, applications, and workloads. Excessive policy complexity can increase operational overhead, while overly broad trust policies can weaken security outcomes.

Threats relevant to this layer include credential theft, compromised devices, unauthorized access, excessive trust permissions, policy drift, and abuse of trusted identities or sessions.

Telemetry should include authentication events, and authorization decisions, device trust state, identity mappings, and integrated threat detections.

- Use centralized identity services when consistent enterprise-wide policy and simplified governance are priorities.
- Apply posture-based access controls when device health must be validated before access is granted.
- Use role-based or group-based policy models when large user populations require scalable access governance.
- Apply identity-driven segmentation policies when least-privilege communication and east-west threat containment are required.
- Deploy distributed policy nodes when multiple sites require resilient local authentication performance.
- Integrate MFA or higher assurance controls when protecting sensitive applications or privileged access.

Zero Trust strategy and framework compliance

Following Zero Trust principles aligned to frameworks such as **NIST 800-207**, access decisions are continuously evaluated using identity, posture, device trust, application context, and environmental conditions rather than network location alone.

- **Policy Decision Point (PDP):** Cisco Identity Services Engine (ISE) and Cisco Secure Access evaluate identity, posture, and contextual trust signals.
- **Policy Enforcement Point (PEP):** Cisco Catalyst switches, Cisco wireless platforms, Cisco routers, Cisco Secure Firewall platforms, and Cisco Secure Access enforce identity-aware policy decisions across the environment.

Best practices

- Enforce least-privilege access by default.
- Standardize policy models across wired, wireless, and remote access domains.
- Use strong authentication methods for users, devices, and administrators.
- Automate onboarding and policy updates where practical.
- Enable detailed identity telemetry and logging for audit and response.
- Regularly review stale accounts, excessive privileges, and policy exceptions.

Layer 4: Unified management and agentic ops

Unified management and agentic ops provide the operational control layer of the SNRA model. This layer centralizes visibility, administration, automation, assurance, and lifecycle management across the networking and security domains. It enables organizations to replace fragmented toolsets and reactive support models with a coordinated framework for infrastructure, policy, monitoring, and change management.

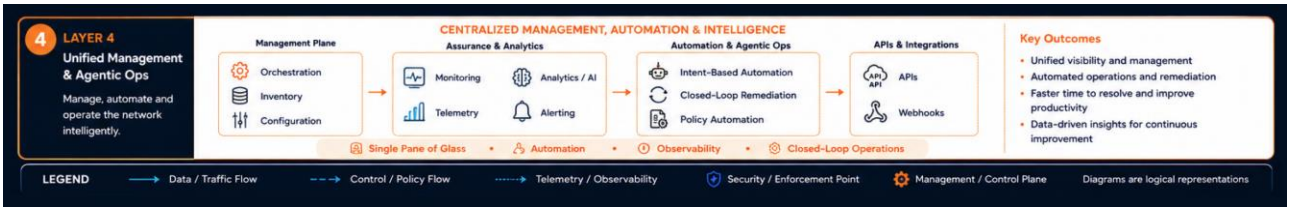

This layer of operational consistency, accelerates deployment, reduces manual effort, and enhances service quality through centralized operations and data-driven workflows across campus, branch, wireless, WAN, and security domains.

Typical technologies and platforms used in this layer include:

- Cisco Meraki Dashboard for cloud-managed operations
- Cisco Catalyst Center for on-premises network operations and assurance
- Cisco Security Cloud Control (SCC) for centralized security policy management
- Cisco Identity Services Engine (ISE) operational integrations
- Inventory, licensing, and asset management platforms
- Configuration templates and policy orchestration tools
- Workflow automation platforms and APIs
- Monitoring, assurance, and operational dashboards
- Event correlation, alerting, analytics, and reporting platforms
- AI-assisted operational workflows and guided remediation capabilities

Core capabilities, resiliency, and security

This layer provides a single operations plane for managing infrastructure lifecycle, policy consistency, software updates, assurance, and change control. Administrators gain centralized visibility into devices, users, applications, policy events, and service health through unified operational platforms.

Automation is a core capability. Standardized templates, zero-touch onboarding, policy orchestration, and secure API-driven integrations reduce repetitive tasks and accelerate operational changes. AI-assisted operations further enhance efficiency through analytics-driven troubleshooting, guided remediation, and operational insights.

Resiliency is strengthened through proactive monitoring, automated alerting, backup and recovery processes, distributed management architectures where required. Centralized lifecycle management improves operational consistency as the environment scales.

Security controls include role-based administrative access, change governance, secure API integrations, audit logging, encrypted management communications, and change-control workflows.

Telemetry, analytics, and operational workflows

Operational telemetry may include:

- Device and interface health
- Client experience metrics
- Path quality and application performance
- Configuration drift and capacity trends
- Policy events and security alerts
- Software lifecycle and compliance status

Centralized analytics improve troubleshooting, operational planning, assurance, and incident response across the environment.

Common operational workflows include:

- Site onboarding and provisioning
- Software lifecycle management
- Policy updates and orchestration
- Health monitoring and assurance
- Incident response and remediation
- Capacity planning and reporting
- Compliance validation and operational auditing

Design considerations

Key decisions include cloud-managed versus on-premises operations, automation maturity, API integration strategy, operational ownership, data retention requirements, and separation of duties. Excessive tool fragmentation can increase operational complexity, while over-centralization without governance can create bottlenecks.

Threats relevant to this layer include unauthorized administrative access, compromised credentials, insecure automation workflows, configuration drift, incomplete visibility, and delayed operational response.

Operational scale should be evaluated based on:

- Number of sites and devices
- Administrative scope
- Telemetry volume
- Automation frequency
- Lifecycle management requirements
- Use Cisco Meraki Dashboard when simplified administration and multi-site visibility are operational priorities.
- Use Cisco Catalyst Center when localized operational control, governance, or data residency requirements are key.
- Adopt API-driven automation when operational scale requires repeatable and programmatic workflows.
- Use AI-assisted operations when proactive troubleshooting and guided remediation are business priorities.
- Deploy distributed management architectures when survivability, latency, or scale require localized operational presence.
- Consolidate operational tools when fragmented workflows increase operational overhead or delay response times.

Best practices

- Standardize operational workflows across sites and technology domains.
- Automate repetitive tasks and preserve human oversight for critical changes.
- Enable rich telemetry, assurance, and analytics from initial deployment.
- Protect management platforms with strong authentication and role-based access.
- Maintain disciplined software lifecycle and configuration governance.

- Use analytics and automation to shift operations from reactive support to proactive optimization.

Layer 5: Extended ecosystem

Extended Ecosystem provides the integration and extensibility layer of the SNRA model. This layer connects the core network architecture to adjacent IT, security, cloud, observability, and business systems through APIs and operational integrations.

This layer transforms the network into a broader operational platform by integrating networking, security, analytics, automation, and enterprise workflows. These integrations improve operational visibility, accelerate response times, and support coordinated workflows across technology domains.

Typical technologies and platforms used in this layer include:

- APIs and webhook-based integrations
- Cisco Splunk and SIEM platforms
- SOAR and incident response platforms
- Cisco Secure Access and Cisco XDR
- Cisco ThousandEyes for digital experience monitoring
- Identity and directory services
- Cloud and SaaS integrations
- IT service management platforms
- Reporting, analytics, and data lake platforms
- Automation and orchestration frameworks

Core capabilities, resiliency, and security

This layer enables operational data exchange, workflow automation, coordinated security response, and cross-platform visibility across networking, security, cloud, and business systems. Network telemetry can be shared with analytics and security platforms, while external systems can trigger automated provisioning, policy updates, or operational workflows.

Cloud and SaaS integrations extend visibility and policy coordination across distributed users, applications, and remote operations. Ecosystem integrations also help organizations preserve existing investments while creating a more unified operational model.

Resiliency is improved through diversified integrations, redundant data paths where required, and automated failover or notification workflows. API-driven integration models also provide scalable connectivity between operational platforms without excessive manual processes.

Security controls should include secure APIs, token management, encrypted communications, role-based access control, audit logging, and governance over automation actions. Shared operational data should follow least-privilege and privacy principles.

Cloud, SaaS, and operational workflows

Common workflows include:

- Automated ticket creation
- Incident enrichment and threat response coordination
- User experience monitoring
- Cloud policy synchronization
- Executive reporting and analytics
- Asset reconciliation
- HR-driven provisioning workflows
- Cross-platform event correlation

Operational telemetry may include:

- Application path performance
- User experience metrics
- Security alerts and detections
- API transaction logs
- Automation success and failure events
- Cross-platform operational analytics

Design considerations

Key design considerations include integration priorities, API maturity, ownership boundaries, workflow criticality, data retention requirements, and vendor interoperability. Excessive integration complexity can increase operational overhead, while insufficient integration can preserve silos and reduce operational efficiency.

Threats relevant to this layer include insecure APIs, token compromise, excessive permissions, data leakage, automation failures, and third-party dependency risks.

Operational scale should be evaluated based on:

- Integration count
- Event volume
- Automation frequency
- Business criticality of connected systems
- Integrate security platforms first when faster detection and coordinated response are strategic priorities.
- Integrate IT service management platforms when workflow efficiency and operational consistency are primary objectives.
- Using cloud and SaaS integrations when supporting distributed users and cloud-hosted applications is operationally important.
- Adopt API-driven integrations when repeatability and scalability are required across multiple systems.
- Use event-driven automation when rapid response and reduced manual effort are priorities.
- Phase integrations gradually when governance models or operational maturity are still evolving.

Best practices

- Prioritize integrations that deliver measurable operational or security outcomes.
- Use secure API authentication, token rotation, and encrypted data exchange.
- Apply least-privilege access to integrations and automation accounts.
- Standardize data models and event handling where practical.
- Continuously monitor integration health and automation outcomes.
- Expand ecosystem integrations in phases to maintain operational control and reduce complexity.

Operations and lifecycle

This section explains how the SNRA is operated, maintained, and continuously optimized across the deployment lifecycle. Effective lifecycle operations translate architectural design into consistent service delivery, operational resiliency, and long-term scalability.

Day 0 / Day 1 / Day 2 model

SNRA aligns operational planning to a lifecycle model that simplifies deployment readiness, implementation, and ongoing optimization.

Day 0 activities focus on **planning and preparation**. This includes architecture selection, deployment model alignment, capacity planning, security policy definition, addressing standards, governance decisions, and readiness of identity, management, and integration platforms. Strong Day 0 discipline reduces deployment risk and creates a repeatable foundation for scale.

Day 1 activities focus on **deployment and service activation**. Infrastructure is introduced into production, management platforms are enabled, segmentation and policy are applied, and validation confirms that connectivity, security, and operational visibility meet design intent. Whether cloud-managed, on-premises, or programmable, Day 1 success depends on standardized execution and clear operational ownership.

Day 2 activities focus on **steady-state operations and continuous improvement**. This includes software lifecycle management, policy refinement, capacity expansion, troubleshooting, security response, compliance reporting, and user experience optimization. Day 2 maturity is where long-term operational efficiency and return on investment are realized.

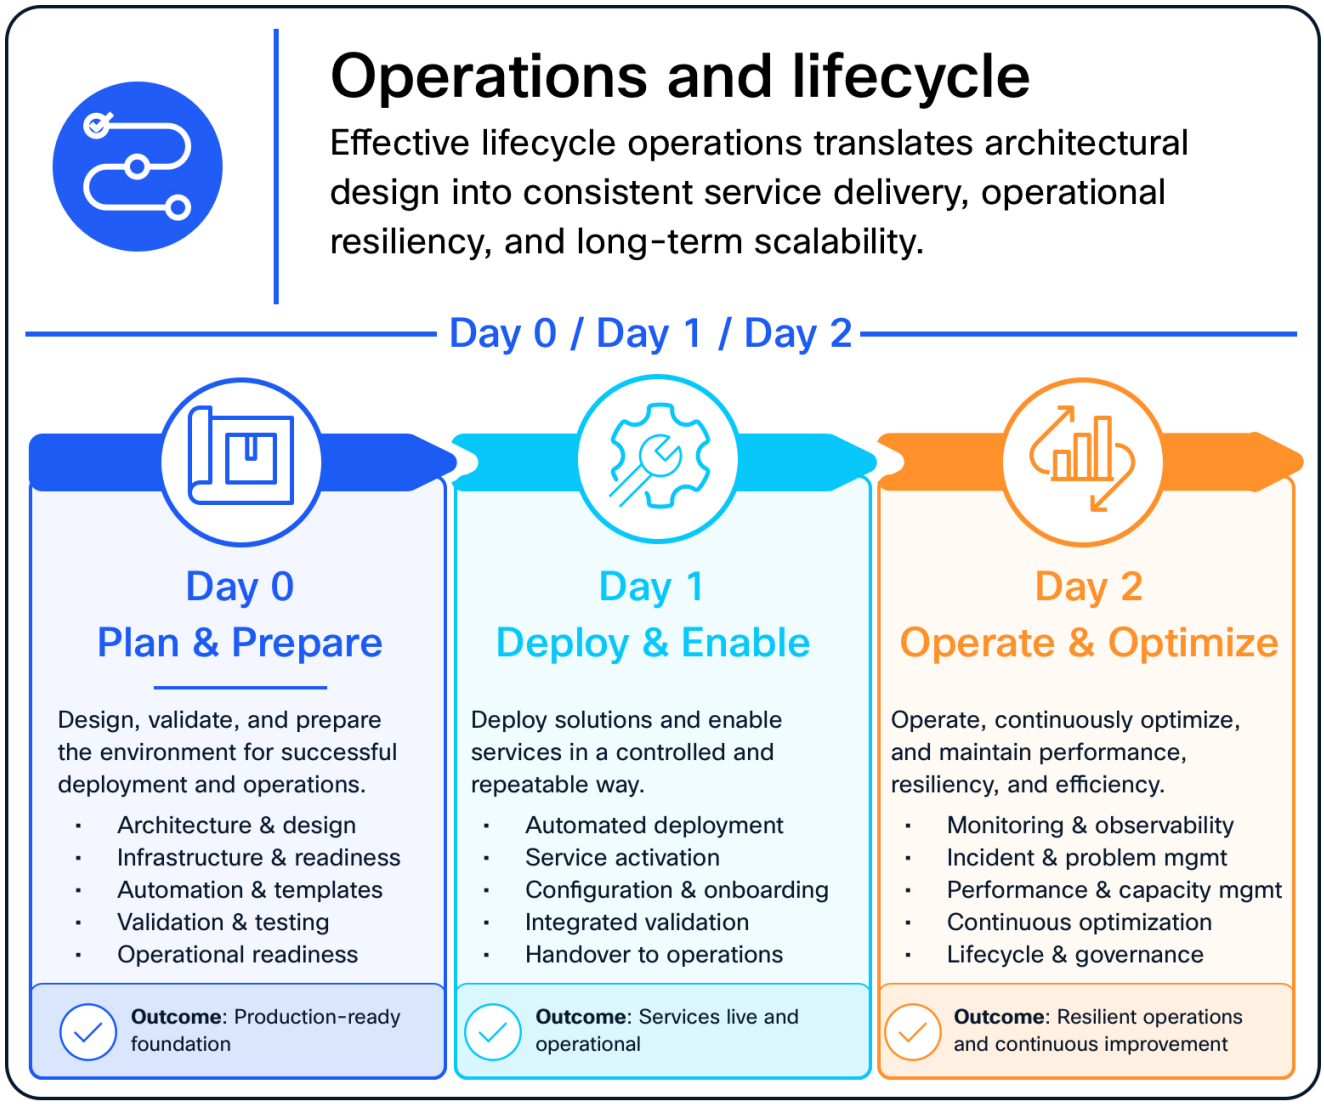

Monitoring and telemetry

SNRA adopts an observability-first operational model that extends beyond device status to include user experience, application performance, segmentation behavior, identity events, and end-to-end path visibility.

Telemetry is collected across infrastructure, wireless, routing, identity, and policy enforcement domains and may include:

- Device and interface health
- Client connectivity and wireless performance
- Authentication and authorization events
- Segmentation and policy outcomes
- Application response times
- WAN and path-quality metrics
- Security alerts and operational anomalies

Centralized platforms such as Cisco Meraki Dashboard, Cisco Catalyst Center, Cisco Security Cloud Control (SCC), Cisco Identity Services Engine (ISE), and Cisco ThousandEyes aggregate and correlate this telemetry to support assurance, analytics, troubleshooting, and operational response.

A mature observability strategy correlates telemetry across architecture layers to accelerate fault isolation and reduce mean time to resolution. User experience issues can be traced across connectivity, identity, segmentation, policy enforcement, and application performance through unified operational workflows rather than isolated troubleshooting processes.

Organizations should prioritize monitoring strategies aligned to service outcomes rather than infrastructure status alone. User experience, application availability, policy effectiveness, and operational trends should be monitored alongside traditional network health metrics.

Proactive alerting, anomaly detection, analytics, and trend analysis further enable operations teams to transition from reactive troubleshooting to predictive and preventative operations.

Automation and APIs

Automation is a core operating principle within SNRA and becomes increasingly valuable as environments scale across campuses, branches, remote users, and distributed services. Manual processes that work in small environments often become slow, inconsistent, and costly at enterprise scale. SNRA addresses this through programmable platforms, policy automation, and API-driven workflows that improve speed, accuracy, and operational efficiency.

Across all deployment models, automation can be applied to device onboarding, configuration standardization, policy deployment, software lifecycle management, inventory reconciliation, reporting, compliance validation, and incident response. These capabilities reduce repetitive effort while improving consistency across sites.

In **cloud-managed environments**, platforms such as Cisco Meraki Dashboard provide automation through templates, zero-touch provisioning, scheduled firmware workflows, APIs, and webhooks. Organizations can automate network creation, device onboarding, policy updates, alerting, and inventory synchronization.

In **on-premises environments**, platforms such as Cisco Catalyst Center provide controller-driven provisioning, policy orchestration, software image management, assurance workflows, and integration APIs. These capabilities are well suited for large campus environments requiring repeatable standards and coordinated lifecycle operations.

The **programmable deployment** model offers the highest level of customization through CLI automation, REST APIs, Python tooling, Ansible playbooks, Terraform workflows, and orchestration pipelines. This model is ideal for organizations with mature NetOps or DevOps practices that require version-controlled changes and automated operations.

Common programmability use cases within SNRA include:

- Automated site deployment for campus and branch expansion
- Bulk policy deployment across wired, wireless, and WAN environments
- Dynamic segmentation changes based on identity or operational events
- Scheduled software upgrades and maintenance workflows
- Compliance auditing against operational standards
- Automated ticket generation and operational alerting
- Automated remediation for common operational faults
- API-driven dashboards, analytics, and operational reporting

Organizations may begin with templates and targeted automation, then mature toward infrastructure-as-code and event-driven operations over time. When executed effectively, automation allows IT teams to shift focus from manual administration to innovation, resilience, and user experience.

Scalability and performance

SNRA is designed to scale predictably while maintaining consistent performance across campus, branch, and distributed environments. The architecture supports modular expansion at the Secure Network Infrastructure layer and flexible logical segmentation at the Scalable Fabric Segmentation layer, allowing organizations to add users, devices, services, and sites without requiring architectural redesign.

Architectural scaling and growth

Scalability in SNRA is achieved by separating physical infrastructure growth from logical segmentation. Infrastructure scales through hierarchical and fabric-based architectures using Cisco Catalyst switching platforms, Cisco routers, VXLAN EVPN fabrics, and modular access, distribution, and core designs. Logical segmentation scales independently through VRFs, VLANs, Security Group Tags (SGTs), and identity-based policy models that allow new services and user groups to be introduced without disrupting existing environments.

Growth considerations commonly include:

- Increasing endpoint and wireless client density
- Expanding east-west traffic flows
- Greater cloud and SaaS application usage
- Additional branch and remote connectivity requirements
- Expanded segmentation and policy domains

As environments scale, organizations often transition from compact one-tier or two-tier architectures to multi-tier or fabric-based designs with distributed policy enforcement and greater operational automation.

Performance boundaries and bottlenecks

Performance is influenced by infrastructure capacity, segmentation design, policy enforcement placement, and application traffic patterns. Key constraints often include uplink bandwidth, switching and routing throughput, wireless client density, WAN path quality, and processing overhead from inspection and policy enforcement.

Common bottlenecks include oversubscribed uplinks, centralized enforcement points, WAN congestion, and high-density access environments. These conditions can impact latency-sensitive applications such as voice, video, and real-time collaboration if not addressed early in the design.

As environments grow toward large-scale deployments, organizations may require:

- Higher-capacity Cisco Catalyst switching platforms
- Expanded VXLAN EVPN fabric architectures
- Distributed enforcement and segmentation models
- Greater automation and telemetry integration
- More granular traffic engineering and routing control

Mitigation strategies focus on distributing traffic across multiple paths, scaling uplink capacity, placing enforcement closer to access where appropriate, and aligning segmentation boundaries to actual traffic flows. Avoiding unnecessary traffic hair-pinning and reducing dependency on centralized choke points are critical to maintaining performance at scale.

Design guidance across SNRA layers

Scalability and performance are directly influenced by design decisions across both infrastructure and segmentation layers. Infrastructure determines how traffic moves and scales physically, while segmentation controls how traffic is isolated, secured, and forwarded logically across the environment.

Well-balanced architectures distribute traffic processing and policy enforcement, align segmentation to business intent, and use centralized operations and automation to maintain consistency as the environment grows. When these elements are properly aligned, SNRA enables organizations to scale efficiently while preserving user experience, operational simplicity, and security posture.

Validation and testing

The Secure Network Reference Architecture design is validated to confirm expected outcomes for connectivity, security, scalability, and operational performance across supported deployment models.

Validation methodology

SNRA validation follows a structured, scenario-based methodology aligned to real-world enterprise requirements across campus, branch, and distributed environments.

The methodology focuses on end-to-end architectural behavior rather than isolated feature validation. Each SNRA domain is validated based on its contribution to secure connectivity, scalable segmentation, identity-driven access control, operational visibility, and lifecycle management.

Testing incorporates functional validation, failure scenarios, scale conditions, and operational workflows. This includes validation of steady-state operations as well as dynamic events such as device onboarding, policy changes, failover conditions, and service disruptions. Observability platforms and telemetry are used to confirm expected system behavior and to ensure that issues can be detected and isolated efficiently.

Test scenarios

Validation scenarios are designed to reflect common enterprise operational conditions across SNRA layers, including:

- User and device onboarding across wired, wireless, and remote access environments
- Identity-based authentication, authorization, and posture validation workflows
- Segmentation enforcement across macro- and micro-segmentation boundaries
- ZTNA policy validation for application and remote access workflows
- Hybrid mesh firewall (HMF) and identity-driven micro-segmentation policy enforcement verification
- Inter-segment communication and policy enforcement verification
- Application access across campus, branch, and cloud environments
- High availability and failover across infrastructure components and control planes
- WAN path performance and application experience validation
- Centralized management operations including configuration updates and policy changes
- Telemetry collection, alerting, and observability across domains
- Integration validation with security, analytics, and IT service management platforms

These scenarios validate operational consistency under both normal and degraded conditions.

Key findings

Validation confirms that SNRA delivers consistent operational and security outcomes across all architecture domains when deployed according to validated design guidance.

At the **Secure Network Infrastructure layer**, testing validates resilient connectivity, stable routing behavior, and effective failover across Cisco Catalyst switching, wireless, and routing platforms. At the Scalable Fabric Segmentation layer, validation confirms consistent segmentation enforcement, controlled east-west traffic behavior, and scalable VXLAN EVPN and VRF-based forwarding.

At the **Zero Trust Access and Hybrid Mesh Firewall layer**, identity-driven access policies operate consistently across wired, wireless, remote access, and application environments with dynamic enforcement based on user, device, posture, and contextual trust signals.

At the **Unified Management and Agentic Ops layer**, centralized operational platforms provide consistent lifecycle management, telemetry visibility, policy orchestration, and operational assurance. At the Extended Ecosystem layer, integrations with analytics, security, observability, and IT service management platforms enable coordinated workflows and improved operational response.

Overall, validation demonstrates that SNRA supports scalable growth, resilient operations, predictable performance, consistent policy enforcement, and operational efficiency across all supported deployment models.

Validated software and hardware versions

This section defines the validated software versions and platform baselines used to support SNRA, ensuring predictable behavior, interoperability, and alignment to Cisco best practices.

These versions reflect validated interoperability across SNRA architecture domains, including Secure Network Infrastructure, Scalable Fabric Segmentation, Zero Trust Access and Hybrid Mesh Firewall, and Unified Management and Agentic Ops.

Validated software baseline (June 2026)

The following software versions represent the validated baseline for SNRA deployments and should be used as the reference point for design, testing, and production alignment:

- IOS-XE: 26.1.x or later (aligned to enterprise campus and routing platforms)
- Cisco Meraki dashboard: current production release (AC3 or later)
- Cisco Catalyst Center: 3.2.x release train
- Cisco SD-WAN: 20.18.x release train
- Cisco Identity Services Engine (ISE): 3.4 patch 4
- Cisco FTD 10.0.0

Validated hardware baseline (June 2026)

The following hardware models represent the validated baseline for SNRA deployments and should be used as the reference point for design, testing, and production alignment:

- Cisco Switches: C9200/9300/9400/9500 series, C9200CX/9300X/9500X series, C9350 and 9610 series, and MS150/250/350 series
- Cisco Wireless: C916x and C917x series, and C9800-H1-MCG
- Cisco Routers: C8200/8300 and C8500 series
- Cisco Firewalls: FTD 1200/3100 and 6100 series

Platform alignment by deployment model

Software selection should align to the chosen SNRA deployment model while maintaining consistency with the validated baseline.

Cloud-managed deployments leverage the Cisco Meraki Dashboard, where software lifecycle management is handled through the cloud, simplifying version control and operational overhead. In this model, customers benefit from automated updates, consistent feature delivery, and reduced dependency on manual software management.

On-premises deployments rely on IOS-XE-based infrastructure and Cisco Catalyst Center for lifecycle management, policy orchestration, and assurance. This model provides greater control over upgrade timing, software validation cycles, and operational governance.

Programmable deployments extend these platforms through APIs and automation frameworks, allowing customers to integrate software lifecycle processes into CI/CD pipelines, infrastructure-as-code workflows, and custom operational models.

Design guidance

When selecting and managing software versions within SNRA:

- Align Cisco platforms and services to validated software release trains to maintain interoperability and operational consistency
- Prefer long-lived and stable software releases for production deployments
- Standardize software versions across campus, branch, wireless, routing, and security environments to reduce operational complexity
- Use centralized lifecycle management platforms such as Cisco Meraki Dashboard, Cisco Catalyst Center, and Cisco Security Cloud Control (SCC) to maintain software consistency and operational visibility
- Plan software upgrades as part of ongoing Day 2 lifecycle operations rather than reactive maintenance activities

Consistent software baselines and disciplined lifecycle management practices help maintain stable, secure, and scalable SNRA deployments while improving long-term operational efficiency and resiliency.

Example use cases

Use case 1: Campus-Branch secure connectivity

This use case demonstrates how SNRA supports secure connectivity between campus and branch environments while maintaining consistent identity-based policy enforcement and segmentation across distributed locations.

Campus and branch sites connect through Cisco SD-WAN or secure routed infrastructure using encrypted IPsec tunnels, enabling protected communication across locations while maintaining resiliency and operational consistency. Users and devices authenticate through Cisco ISE, which provides identity, device profiling, posture validation, and authorization services across both environments.

Identity information is mapped to Security Group Tags (SGTs), allowing policy to follow users and devices independently of IP addressing or physical location. Macro-segmentation using VRFs separates major trust zones, while micro-segmentation applies least-privilege controls between users, devices, and applications.

For example, a branch user accessing a business application hosted at the campus location is authenticated and assigned an SGT-based policy. As traffic traverses the secure tunnel, identity context is preserved and evaluated against centralized policies, permitting access only to authorized applications and services.

This approach extends campus security policy consistently into branch environments, reduces lateral movement risk, and provides scalable Zero Trust policy enforcement across the enterprise.

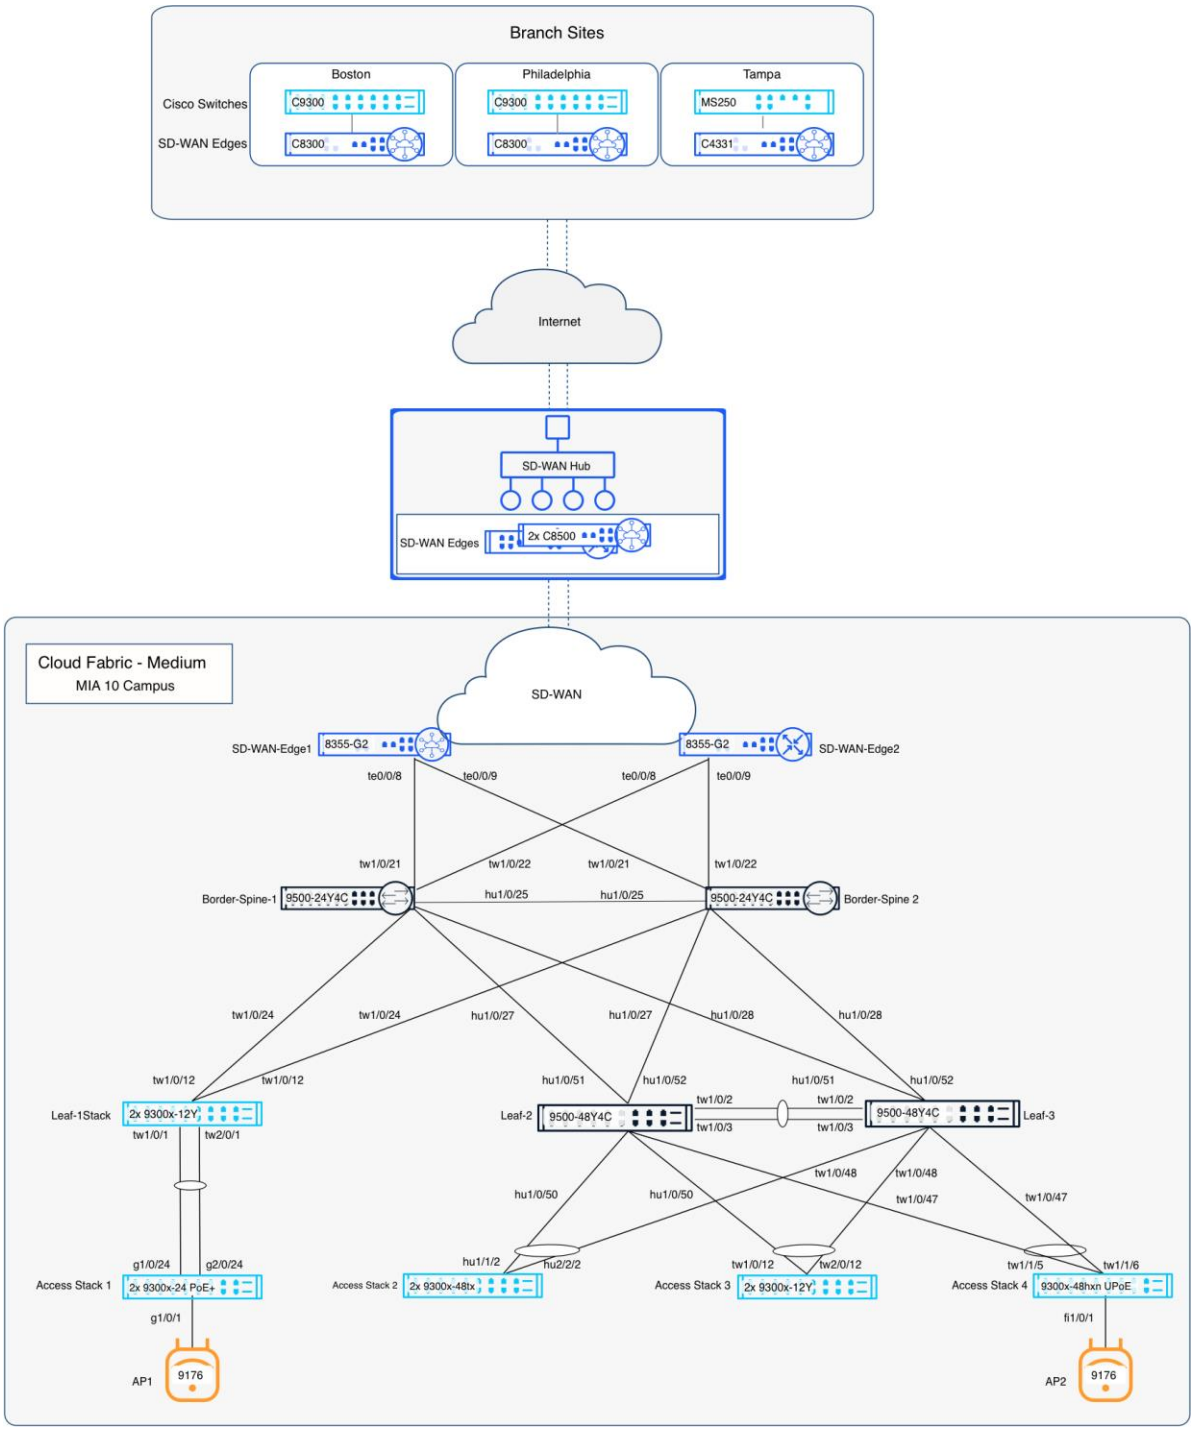

Figure 12. Use case 1: Campus-Branch VPN policy

Use case: The IoT ecosystem

This use case demonstrates how SNRA secures IoT and operational technology (OT) environments using identity-driven macro-segmentation, policy enforcement, and restricted communication models. Because many IoT devices cannot support endpoint agents or advanced security controls, SNRA uses network-based identity classification and segmentation to reduce risk and limit lateral movement.

Device identity is established at the access layer through Cisco Identity Services Engine (ISE) using profiling and MAC Authentication Bypass (MAB). Based on device classification, Cisco ISE assigns Security Group Tags (SGTs) that define policy scope and permitted communications across the environment. These identity attributes are propagated through Cisco switching and routed infrastructure to downstream policy enforcement platforms.

Typical IoT segmentation models include:

- Physical security devices restricted to network video recorder (NVR) systems and approved vendor update services
- Building automation systems limited to HVAC controllers and management platforms
- Imaging and printer services restricted to print servers and approved protocols
- Point of Sale (POS) systems isolated within PCI-aligned payment environments and permitted to communicate only with authorized payment gateways

Cisco Secure Firewall platforms managed through Cisco Security Cloud Control (SCC) enforce segmentation policies between trust zones and prevent unauthorized east-west communication between IoT, user, and application environments. This approach limits the ability of compromised devices to move laterally or access sensitive systems outside their authorized communication scope.

For IoT devices requiring Internet or cloud connectivity, Cisco Secure Access applies controlled egress policies using FQDN-based restrictions and identity-aware policy enforcement. These controls limit outbound communications to approved destinations and reduce exposure to command-and-control (C2) or unauthorized external communications.

This use case demonstrates how SNRA combines Cisco Identity Services Engine (ISE), Security Group Tags (SGTs), Cisco Secure Firewall platforms, Cisco Secure Access, and scalable segmentation to provide secure onboarding, controlled communications, and consistent policy enforcement across IoT and OT environments.

Use case 3: Unmanaged assets

This use case demonstrates how SNRA secures unmanaged devices, including guest endpoints, contractor-owned systems, and non-corporate hardware, using an isolation-by-default security model. Because these devices are outside centralized operational control, SNRA treats them as high-risk assets and restricts access to corporate resources, applications, and trusted network segments.

Device classification and access control are enforced through Cisco Identity Services Engine (ISE), which identifies unmanaged devices and assigns restrictive Security Group Tags (SGTs) aligned to limited-access policy groups. These devices are segmented from managed corporate users, applications, and infrastructure through identity-driven policy enforcement across Cisco switching, routing, and security platforms.

Unmanaged devices are denied access to private applications, internal services, and trusted corporate segments. East-west communication attempts toward managed endpoints or sensitive resources are blocked through distributed policy enforcement at Cisco Catalyst switching and security enforcement points.

Internet access for unmanaged assets is restricted through Cisco Secure Access using controlled egress policies and category-based filtering. Organizations can limit access to high-risk destinations such as:

- Peer-to-peer services
- File-sharing platforms

- Hacking and malicious-content categories
- Unauthorized external applications

Unmanaged assets are also excluded from Zero Trust Network Access (ZTNA) entitlement workflows and identity synchronization processes, preventing authentication to private applications or protected internal resources even when physically connected to the network.

This use case demonstrates how SNRA combines Cisco Identity Services Engine (ISE), Security Group Tags (SGTs), Cisco Secure Access, Cisco Catalyst switching platforms, and distributed policy enforcement to isolate unmanaged devices, reduce attack surface, and maintain separation between untrusted endpoints and corporate environments.

Design summary and recommendations

SNRA provides a unified architecture that integrates infrastructure, segmentation, identity-driven policy, operations, and ecosystem integration into a consistent and scalable design model. By aligning to the five SNRA architecture domains, organizations can build secure, resilient, and operationally efficient campus, branch, and distributed enterprise environments.

The architecture supports cloud-managed, on-premises, and programmable deployment models, allowing organizations to align operational approaches to business priorities, governance requirements, and technical maturity while maintaining a common architectural framework.

Across all deployment models, SNRA emphasizes modular design, consistent policy enforcement, operational simplicity, and lifecycle efficiency. These principles help organizations scale predictably, improve user experience, reduce operational complexity, and maintain consistent security posture across distributed environments.

SNRA also aligns closely to the Cisco networking and security portfolio. Platforms such as Cisco Meraki Dashboard, Cisco Catalyst Center, Cisco Identity Services Engine (ISE), Cisco ThousandEyes, Cisco Secure Access, Cisco Security Cloud Control (SCC), and Cisco Secure Firewall platforms provide integrated visibility, policy enforcement, observability, and lifecycle management across networking and security domains.

This integration enables organizations to extend operational visibility and coordinated security workflows beyond the network while improving troubleshooting, threat response, automation, and user experience monitoring.

When to use this architecture

SNRA is well suited for organizations:

- Deploying new campus or branch environments
- Modernizing legacy network architectures
- Standardizing operations across multiple sites
- Implementing Zero Trust and segmentation strategies
- Improving operational visibility and lifecycle management
- Supporting distributed users, applications, and hybrid work environments

The architecture is designed to scale from smaller environments to large, distributed enterprises through modular expansion, fabric-based architectures, distributed policy enforcement, and automation-driven operations.

Key operational and architectural benefits include:

- Improved security posture through Zero Trust access and scalable segmentation
- Consistent policy enforcement across wired, wireless, remote access, and cloud environments
- Simplified lifecycle operations through centralized management and automation
- Scalable infrastructure aligned to user, device, and application growth
- Enhanced visibility and faster troubleshooting through integrated telemetry and observability
- Flexible deployment models aligned to operational, governance, and business requirements

References

Additional guidance and detailed implementation information can be found in related Cisco Validated Designs, product documentation, and industry standards.

Public SNRA References

- Secure Networking Overview
- Secure Campus & Branch Networking
- Secure Campus Networks – Solution Brief

Existing Cisco Validated Designs

- Cisco Cloud Fabric Validated Case Study
- Cisco SD-Access Design Guide
- Cisco Unified Branch Design Guide
- Cisco Common Policy Integration Guide

Related References

- Cisco product documentation for switching, wireless, security, and management platforms
- Observability and assurance documentation including ThousandEyes
- Relevant industry standards for networking, security, and identity frameworks

By following the guidance in this architecture, organizations can build a secure, scalable, and future-ready network that supports both current requirements and long-term digital transformation initiatives.